Università degli Studi di Padova Corso di Laurea in Scienze Statistiche



ANALISIÆ PREVISIONI DE LA PICCHI DI PREZZO ANEIA MERCMITIMILETTRICIM

Analysis and prediction of Price spikes in Electricity Markets

Relatore: Pott. Francesco Risi	
DipartimentoRdidhienzedStStistih	ned
d	
d	
	LSureSndozơGiSnhSrւoc֍imiond
d	
d	
d	
d	
d	
A110AAhhSAemił	noz&011/2012A
А	

- A A A A

Contents

				Pa	age
Tε	able o	f Conte	ents		i
Li	st of	Tables			v
Li	st of	Figures	• • • • • • • • • • • • • • • • • • • •		vii
Cl	hapte	$\mathbf{e}\mathbf{r}$			
In	trodu	ction			1
1	Ene	rgy Ma	rkets and Spot Prices		3
	1.1	Energ	y Markets: a Short Overview		3
		1.1.1	Liberalization and regulation		3
		1.1.2	The marketplace		4
		1.1.3	Other relevant energy aspects		7
	1.2	Featur	res of energy spot prices		10
2	The	Price S	Spikes		13
	2.1	Introd	luction		13
	2.2	What	is a price spike?		13
		2.2.1	Spikes causing factors		14
		2.2.2	Why should we study price spikes?		15
	2.3	Price	spikes detection: a literature review		15
		2.3.1	Models for price forecasting		16
		2.3.2	Depuration of electricity prices time series		17
		2.3.3	Detection of spikes		20
	2.4	A pric	e spikes identification method		21
		2.4.1	Model for electricity prices		21
		2.4.2	Method for cleaning the time series		22
		2.4.3	Testing model accuracy		25
		2.4.4	Some final observations		25
	2.5	Applie	cation: spike identification		26
		2.5.1	1st example		28
		2.5.2	2nd example		31
		2.5.3	Results		32

3	The	Model 3	7	
	3.1	Introduction	7	
	3.2	The standard ACD Model	8	
		3.2.1 The model and its assumption	8	
	3.3	The ACH Framework	0	
		3.3.1 Why ACH?	0	
		3.3.2 The model and its assumption $\ldots \ldots \ldots \ldots \ldots 4$	0	
		3.3.3 $ACH(1,1)$ framework summary $\ldots \ldots \ldots \ldots \ldots 4$	4	
	3.4	Parameters estimation	4	
		3.4.1 Numerical aspects	5	
	3.5	Forecasting	6	
4	App	lication	7	
	4.1	Introduction: Christensen et al. (2011)	7	
	4.2	The Australian National Electricity Market	0	
		4.2.1 Data	1	
		4.2.2 Estimating $ACH(1,1)$ parameters	3	
	4.3	The Italian electricity market	8	
		4.3.1 Data	2	
		4.3.2 Spike detection $\ldots \ldots \ldots$	5	
		4.3.3 Estimating $ACH(1,1)$ parameters	5	
	4.4	Comparing results: a benchmark Logit model	9	
		4.4.1 New South Wales EM analysis	1	
		4.4.2 Italian EM analysis $\ldots \ldots \ldots$	2	
		4.4.3 Comparison between Italian and Australian forecasts 74	4	
		4.4.4 Applying the Logit to the whole NSW spot prices TS \cdot 7.	5	
	4.5	Yet another $ACH(1,1)$ estimation	8	
5	Con	cluding Remarks	3	
	5.1	Short summary	3	
	5.2	Significance of the result	3	
	5.3	Future work	6	
Aj	Appendix			
Α	Defi	nitions, Models and Formulas	9	
	A.1	Definitions	9	

	A.2	Models
		A.2.1 Seasonal ARIMA model
		A.2.2 GARCH model
	A.3	Formulas
	A.4	Logit estimates
	A.5	Overall results: Australian and Italian cases
В	R co	le
	B.1	ACH likelihood and estimation functions
Bi	bliogr	aphy

List of Tables

2.1	EACF for the remaining time series (British spot prices 1st load	
	period)	30
2.2	Spike detection summary (British EM)	34
4.1	ACH(1,1) estimation results (Christensen)	49
4.2	Estimates for the ACH(1,1) model (NSW data) $\ldots \ldots \ldots$	56
4.3	$\operatorname{ACH}(1,1)$ spike occurrences analysis (NSW data)	59
4.4	Estimates for the ACH(1,1) model (Italian data) $\ldots \ldots \ldots$	68
4.5	ACH(1,1) spike occurrence analysis (Italian data)	69
4.6	Spike detection capability of $\operatorname{ACH}(1,1)$ and Logit (NSW data) .	71
4.7	Spike detection capability of $ACH(1,1)$ and Logit (Italian data)	73
4.8	Comparison between Italian and Australian forecasts	74
4.9	Spike detection capability of $ACH(1,1)$ and Logit (whole NSW	
	spot prices TS)	77
4.10	Spike detection capability of ACH(1,1) (British data)	82
A.1	Estimates for the Logit model (NSW data)	91
A.2	Estimates for the Logit model (Italian data)	92
A.3	Summary of the overall ACH(1,1) results (NSW data) $\ . \ . \ .$	94
A.4	Summary of the overall $ACH(1,1)$ results (Italian data)	94

List of Figures

1.1	The PX two-side auction mechanism $\hdots \hdots \$	7
1.2	The supply stack	8
2.1	The British spot prices	27
2.2	The APX-PUK spot prices ACF and PACF	27
2.3	GCV and smoothing splines applied to the British spot prices	
	1st load period \ldots	29
2.4	Analysis of the remaining time series (British spot prices 1st	
	load period)	30
2.5	Spike detection with S–ARMA GARCH model (British spot	
	prices 1st load period) \ldots	31
2.6	Spike detection with S–ARMA GARCH model (British spot	
	prices 24th load period)	32
2.7	Spike detection with robust procedure (British spot prices 24th	
	load period)	33
3.1	ACD and ACH timelines	41
4.1	The NEM New South Wales spot prices	51
4.2	The NEM New South Wales spot prices ACF and PACF	52
4.3	The NSW spot prices boxplots and spike counts	55
4.4	The $\operatorname{ACH}(1,1)$ spike analysis (NSW spot prices 33rd load period)	60
4.5	The Italian spot prices	63
4.6	The Italian EM spot prices ACF and PACF	64
4.7	The Italian detrended spot prices boxplots	66
4.8	The $\operatorname{ACH}(1,1)$ spike analysis (Italian spot prices 9th load period)	70
4.9	ACH spikes detected boxplots (NSW data)	72
4.10	The $\operatorname{ACH}(1,1)$ and Logit spike analysis (NSW, 1st dataset) $$.	79
4.11	The $\operatorname{ACH}(1,1)$ and Logit spike analysis (NSW, 2nd dataset)	80
5.1	Analysis of the hazard function barrier (NSW data) $\ . \ . \ .$.	87

Introduction

The energy market liberalization

Over the past two decades, several countries have decided to take the path of the power sector liberalization. The liberalization has been possible by disrupting the traditional vertically integrated monopoly, which involved the whole power dispatch system, from the generation to the sell side, through the energy transportation and distribution.

Chile was the first country in the world to make the power market liberalization, in the year 1982. After that, in the 1990s, the Chilean leading example was followed by a lot of nations, especially in Europe. The number of liberalized electricity markets is still steadily growing.

As liberalizations of other sectors, the motivation behind electricity liberalization is, in the long run, promoting efficiency gains, stimulating technical innovation and leading to efficient investments (Weron, 2006). Nevertheless, in contrast to the described benefits, it is not clear whether the power liberalization can generate losses in the short time yet. Furthermore, we can easily observe that there might not be sufficient incentives for investing in the new generation energy sources within a totally liberalized power marketplace. These are two of the most important issues, which are nowadays influencing the speed of the power liberalization process in the world and conducting different countries to have new ways to approach the liberalization. This is also one of the reasons why the electricity liberalized markets are different from one to another.

Since the energy market (EM) liberalization is a recently evolved and still ongoing process, it has inspired plenty of researchers to study it. Concerning the statistics science field, a lot of research articles have been written and many models have been developed, mostly focusing on prices and loads forecasting. Since the EMs scientific literature is quickly growing, being of great interest for a lot of governments, we will focus on a special feature of the EMs: the price spikes. Also known as price jumps.

Purpose of the document

The purpose of this essay is to develop a model for electricity price spikes occurrences forecasting. In order to do that, at first we provide a short overview of the EMs in general and the stylized facts of electricity loads and prices (e.g. how the liberalized EM works, what price spikes are). A specific method for detecting spike occurrences is then proposed. In the second part, inspired by the research of Christensen et al. (2011), we implement a model for spikes forecasting. At last, we apply the previously developed model on the data coming from real EMs.

Chapter 1

Energy Markets and Spot Prices

Although there is not just one single market model which can be used as a benchmark (also due to the economic and technical characteristics of a given power system), we can explain how the EM generally works. In the first part of this chapter, we briefly introduce the elements upon which an energy liberalized market is based on, while in the second part we focus on the price spike concept.

1.1 Energy Markets: a Short Overview

According to Edwards (2009), the EM is a collection of interrelated businesses focused on delivering electricity and heating fuel to consumer, generating power and actually distributing it. In this section we explore the EM features more precisely and analyze how the market exactly works.

1.1.1 Liberalization and regulation

With the advent of EM liberalization, it was necessary to create an authority which has both the capability and the commitment for regulating them, in order to protect consumer rights and avoid oligopolies. While competitive markets face challenges, it should be acknowledged that competition in wholesale power markets is a national policy. Therefore, each government has to choose the regulator, which seeks to challenge against the prices' extreme volatility, reform markets if needed and search for evidence of anti-competitive behaviour of the market players. In a nutshell, a regulator must identify and assess solutions for making existing wholesale markets more competitive (Hogan, 2009). Examples of regulators are the European Energy Community (EEC) in Europe, the Australian Energy Market Commission (AEMC) in Australia and the Federal Energy Regulatory Commission (FERC) for the United States of America.

Since the transition from regulation to competition has to guarantee an economic and secure system operation, the EM regulation has been often a rather complicated process, which is still being modified in the largest part of the existing EMs.

1.1.2 The marketplace

Two main kinds of EM have emerged for organizing markets at the wholesaler level: *power pools* and *power exchanges*. Since both of them are sharing many aspects, it is often complicated to distinguish one from another. Nevertheless, they can be explained by using two criteria of initiative and participation (Weron, 2006).

- Power pool (PP) is the natural ever existing market. It can be split into two types: technical (or generation) pools and economic pools. The first type was used by the vertically-integrated utilities to optimize generation with respect to cost minimization and optimal technical dispatch. In such a system the power plants were ranked on merit order, based on production costs. Hence, generation costs and network constraints were the determining factor for dispatch. Trading activities were limited to transactions between utilities from different areas. International trade was limited, due to a low level of interconnection capacity.

Economic pools have been established during the liberalization process in order to facilitate competition between generators. The bid is based on the price at which they are willing to run their power plants. The market clearing price (MCP) is established through a one-sided auction as the intersection of the supply curve (constructed from aggregated supply bids) and the estimated demand, defining also the market clearing volume (MCV). Once the technical constraints have to be satisfied, the PP bids can be very complex, leading to a low level of transparency.

Participation in an economic pool is mandatory (i.e., no trade is allowed outside the pool) and the participants can only be generators.

- Power exchanges (PX) are commonly launched on a private initiative, usually by a combination of generators, distribution companies, traders and large consumer. Unlike the power pools, participation in the exchange is not mandatory. The genuine role of a PX is to match the supply (bid) and demand (ask) of electricity to determine a transparent publicly announced *market clearing price* (MCP).

The most of the recently developed markets are based on the PX model, including IPEX market for Italy, ASX for Australia and UKPX market for United Kingdom (nowadays part of the APX Group).

The spot market

The electric power spot market usually includes three different spot markets: the *day-ahead market*, the *intraday market* and the *dispatch services market*.

- The day-ahead market is a 24 hours open market. It exists because a classical spot market would not be possible, since the transmission system operator (TSO) needs advanced notice to verify that the schedule is feasible and lies within transmission constraints (Weron, 2006). The day-ahead market concerns the day after power transactions and its trading activity is based on a unique day after auction session.
- The intraday market (also known as *balancing market*) allows the TSO to operate in the very short horizons before delivery. Thus, it is possible to slightly modify the power amount traded in the day-ahead market.
- The tasks of the dispatch services market (also called *ancllary services market*) are network flow administration and control, congestion solving and the network real time balance guaranteeing, minimizing the reaction time in case of deviation in supply and demand.

Matching supply and demand in day-ahead market

In order to understand how the MCP is established in the day-ahead market, according to the Keynesian economic rule, we will define the supply and demand sides. On the supply side a supplier assures to sell an amount of energy, which has to be not higher than the amount specified in the selling offer, for a unit price not lower than the one specified. On the other side, a customer assures to buy an amount of energy which has not to be higher than the amount specified, for a unit price not higher than the indicated one.

The bid implicitly contains both fixed operation costs and start-up costs of all generation units. The MCP is then calculated by means of the market clearing algorithm, differently implemented by the authority and the TSO, for each PX.

The day-ahead auction mechanism

As previously introduced, competition in day-ahead PX market has been established through auctions where mainly generators, rather than operators, bid energy prices and quantities.

The MCP is generally not established on a continuous basis but rather in the form of a conducted once per day two-sided auction. For instance, assuming the hour as the market time unit, buyers and suppliers submit bids and offers for each hour of the next day and each hourly MCP is set so that it balances supply and demand (Weron, 2006).

The simple bid format consists of a pair of values: quantity, expressed in MWh, and price (e.g., \in /MWh). Each selling or buying partecipant is able to propose several pairs of values for the same generation or demand unit (Contreras et al., 2001). The intersection of the supply and demand curve, constructed from aggregated supply and demand bids, allows to calculate the MCP. Figure 1.1 shows graphically how the two curves can be traced and how the intersection point, which fixes the MCP and the MCV, can be obtained.

Since the actual EMs may have several further extra conditions which can be differently defined for each market¹, in addition to the simple bids, we have complex bids. In such a case, the market clearing algorithm is a modification of the simple matching algorithm respecting the extra conditions and the tech-

¹Examples of extra conditions are: the cheapest bid quantity has to be designated as non-divisible and a minimum daily income amount, respecting the nature of the generating unit, has to be considered.



Figure 1.1: The Power Exchange two-sides auction mechanism. *MCP* indicates the Market Clearing Price, while *MCV* indicates the Market Clearing Volume.

nical constraints, involving the advanced methods concerning network flowing optimization.

1.1.3 Other relevant energy aspects

Before proceeding with our research, we have to remark some further relevant aspects which have to be considered in order to get a more complete view of the energy prices.

Marginal production cost

For a better understanding of how the EM spot prices are determined, the concept of marginal production cost becomes essential. Let the *supply stack* be the ranking of all generation units of a given utility or of a set of utilities in a given region (Weron, 2006). The supply stack intuitively influences the supply side of the market and is based on the marginal cost of production. Basically, the utility dispatches firstly the energy from nuclear and renewable plants (e.g., hydro, solar, wind, etc.), followed by fossil-fuel power stations. Some EMs, especially where nuclear energy sources are available, exhibit low flexibility within the supply stack, due to the fact that their plants can provide a lot of energy with low or moderate marginal costs. A schematic example of the supply stack is shown in Fig. 1.2.



Figure 1.2: A schematic supply stack (source Weron (2006)).

Normal commodities vs energy: the non-storability

Energy is clearly a commodity, since it is supplied without qualitative differentiation across the EM. However, it differs from other commodities; electricity can not be stored (unless hydro) and it is very expensive to transmit over long distances. Thus, there can be neither a benefit from holding electricity nor a storage cost. For this reason, the energy price is usually not determined by the level of inventories. This property influences the extreme volatility marked by the electricity spot prices, which is higher than any other commodity price volatility.

Technical constraints and traded products

Trading has to account for numerous technical limitations. The bids are accepted solving an optimization problem which maximizes the transactions value and guarantees both the aggregated supply and demand balance and the transit limits between each pair of transmission nodes respect (i.e., according to the minimum and maximum transportation capacity of every single route of the electricity network). When there is no congestion, MCP is the only price for the entire system. On the contrary, there might be a different local price for each region, known as *locational marginal price* (LMP).

Another technical constraint is due to the fact that plenty of electrical plants require a minimum power to run. Such a necessity influences the whole network balance. Furthermore, to guarantee a lower black-out risk, the entire network has to constantly provide a high voltage and face the several bottle necks which might be present in the power transmission grid. These limitations lead to the problem that in presence of a low aggregate demand (often noticed during the off-peak hours), power is sold by the supplier for a lower level required to stimulate new entry. This issue is known as *scarcity value* of energy, which deals in the PX two-sides auction mechanism to a scarcity of stimulating suitable new investments.

Concerning the trading products, the commodification of electricity has led to the development of novel types of contracts for electricity trading. These contracts can be sold either over-the-counter or on organized market. They can be physical or financial contracts (Weron, 2006). Both of the types are strongly necessary to keep supply and demand in balance. The first type has been established to cover the utilities future consumption, while the second allows operators to hedge and speculate on the energy prices.

An interesting class of long-term contracts present in electricity markets are the CO_2 emissions allowances, also known as green allowances. A generator polluting more is obliged to buy extra allowances to covering the pollution produced in a whole given year. On the other side, a cleaner generator can sell the excess allowances and therefore gain an extra profit. This kind of contracts have been established to change the marginal production costs and consequently press operators to promote the renewable energies by investing in the 'green economy', respecting to the Kyoto Protocol.

Inelasticity of demand

Within the largest part of the liberalized EMs the demand side response contribution to establish electricity prices is rather poor. Electricity transmission, unlike other transportation networks, requires coordinated behavior to ensure that injections and withdrawals of electricity are continuously balanced. Consequently, an efficient competitiveness in the generation sector was limited by the government regulation. One of the inefficiencies generated by the regulation is the gap between regulated retail prices and wholesale market prices.

In general, EMs do not handle a demand elastic response as a consequence

of the vertically-integrated era, when the monopoly regime were used to set fixed basic service tariffs. Still, nowadays few measures exist to promote a mature price-responsive demand side (Siddiqui, 2003).

1.2 Features of energy spot prices

One of the important consequences of liberalization is that prices are now determined according to the fundamental economic rule of supply and demand. As we have previously described in §1.1.3, in contrast to stock and bonds, electricity prices are equally spaced high frequency data which are affected by transmission constraints, nature of the generation stack and non-storability constraint. Furthermore, they are characterized by seasonality and weather conditions, as power demand correlates with temperature (Geman and Roncoroni, 2006).

A brief description of the main energy prices stylized facts follows.

Stationarity and mean-reversion

Mean-reversion is widely used in the field of finance as a property of commodities and it can be applied to assets as energy prices. Exley et al. (2004) give the following mean-reversion definition: an asset model is mean reverting if interest rates (and volatilities), yields or growth rates are stationary. Essentially, when the price is high, supply tends to increase thus putting a downward pressure on the price, and vice versa (Deng, 2000).

Therefore, the mean reverting incorporates the tendency of energy prices to gravitate towards a 'normal' equilibrium price level that is usually governed by the cost of production and level of demand (Blanco and Soronow, 2001).

Mean-reversion in energy prices is well supported by the empirical studies of energy price behavior, as well as by the basic microeconomic theory.

High volatility

The previously described characterizing factors add to the energy spot prices an extreme high volatility component. Yet, because of the properties of electricity transmission, an imbalance of supply and demand at any one location on an electricity grid can threaten the stability of the entire grid. The supply and demand matching between any customer and supplier is just part of the overall grid balancing and any mismatch could disrupt the delivery of the product for all suppliers and consumers on the grid (Borenstein, 2001). In concrete terms, applying the standard notion of volatility –namely the standard deviation of returns²–, daily prices usually exhibit volatility up to 50%.

Seasonality and calendar effects

In addition to the intuitively desirable yearly seasonality, due to the consequent high demand during the warmer seasons, weekly and daily periodicity is observed from the ACF. During the week the prices show a business daysweekend structure, with a higher level registered during business days. The intra-day variability is due to the lower demand during the night hours, called *off-peak hours*. Clearly, the energy demand and price increase during the working hours, namely *on-peak hours*.

Furthermore, a calendar event, such as the National Day or a public holiday, may usually move the demand side, causing an energy price shift. As the natural effect of a holiday is to reduce the active businesses, the electricity demand is subjected to a flexion.

Distribution of electricity prices

Since the electricity returns diverge significantly from the log normal distribution, as well as the financial asset returns, they are not normally distributed. The empirical observed distributions exhibit excess kurtosis. We can say that the electricity prices are heavy-tailed or leptokurtic. Electricity prices also show positive asymmetry.

Jumps presence

One of the most pronounced features of electricity markets is the abrupt and generally unanticipated extreme change in the spot prices, known as *jump* or *price spike*. A jump occurs as, within a very short period of time, the system price significantly increases and then drops back to the previous level.

 $^{{}^2}r_t = \ln P_{t+1} - \ln P_t$

Chapter 2

The Price Spikes

2.1 Introduction

Although the widest part of the statistical studies concerns energy spot prices modelling, the research has recently evolved focusing also towards the price spikes estimation issue. Such issue doubtlessly leads researchers to face the most critical and characterizing feature of energy prices.

In the beginning of this chapter we provide a general definition of the term price spike. After that, by means of a brief literature review, starting from the spot price existing models, we outline the different emerged approaches to detect the spikes within an energy spot prices time series. The several issues that may emerge during the spike detecting process are outlined. At the end, we propose a specific framework for detecting spikes and we apply such framework on a real electricity spot prices time series.

2.2 What is a price spike?

Many definitions of a price spike have been given since the electricity markets have been quantitatively analyzed. Although there is no commonly accepted definition, we may in general define a spike as a temporary very high (or low) shift compared to the normal fluctuation regime, noticed within the electricity prices TS.

More precisely, according to Weron (2006), who considers only high shifts, price spikes are "prices that surpass a specified threshold for a brief period of time". At the same time, Geman and Roncoroni (2006) define a spike as "a cluster of upward shocks of relatively large size with respect to normal fluctuations". Later, the starting points of our analysis towards the spike detection are both of the previous definitions. These definitions clearly need to be extended before analyzing the spot prices time series, since they do not provide enough knowledge for coming out with a specific price spike definition.

2.2.1 Spikes causing factors

Firstly, we must specify that there are markets where no spikes are present. For instance, the Polish organized electricity exchanges (i.e. PoIPX and POEE) do not register any extreme event for an incredibly long period of time. Spikes are mainly observed in the largest and most deregulated markets, such as American, Australian and Western European electricity markets.

Electricity spikes are caused by several factors, also depending upon features described in Chapter 1, which can be classified as follows:

- inelasticity of demand curve, mostly attributable to the non storability electricity property (which leads to an absence of cash-and-carry arbitrage);
- severe weather conditions, often in combination with generation outages or transmission failures;
- network features, such as supply grid composition with its constraints and marginal production costs, related to the demand amount;
- composition of electricity sources influencing variables (e.g., cost of fuel, gas, etc.);
- energy policy followed by the authority, such as market incentives (e.g., green allowances), investing plans and amortization schedules;
- trading activities of the market players.

Generally, the 'spiky' behavior is attributable to the fact that a typical regional aggregate supply function of electricity almost always has a kink at a certain capacity level and the supply curve has a steep upward slope beyond that capacity level (Deng, 2000).

However, there are some disagreements. For instance, Weron (2006) argues that the primary spike causing reason is explained by the market players' bidding strategies. Indeed, some players are willing to pay almost any price to be safe by securing a sufficient and continuous supply of power. From an EM point of view, it means that some power amount bids would have been set at the maximum allowed price level. On a large scale, such agent behavior can cause price spikes.

2.2.2 Why should we study price spikes?

Since a price spike is usually followed by a sharp return to normal price level, there is apparently no matter to think that spikes have a huge effect for the system stability. However, the violent spike behavior of electricity prices constitutes a significant risk for EM players (Klüppelberg et al., 2010). Besides, a spike is frequently followed by several further spikes, forming a sort of cluster, with the possibility –as in the Western U.S. Energy Crisis of 2000 and 2001– of a structural temporary system break (i.e., black out). We can therefore affirm that the price spikes influence any EM and even any network risk measure.

From a market point of view, spikes may bring to agents both benefits and losses. On one hand, they may provide profitable opportunities to sell electricity at higher prices in the spot market. On the other hand, they can be a burden if a player has contracts to supply electricity at low, predetermined price (Kanamura and Ohashi, 2007). For instance, an option market player and a supply bilateral contract keeper may easily take the advantages of price spikes.

2.3 Price spikes detection: a literature review

At the beginning of this chapter (see §2.2), two electricity price spike definitions have been given and spikes causing factors have been outlined. Our target now shifts to detecting spikes within an electricity spot prices time series. Consequently, we must firstly define what we quantitatively mean by 'threshold'.

A good strategy to do that is by focusing on the existing electricity price modelling literature and pull out what we believe is relevant for detecting price spikes.

Despite their rarity, price spikes are the very motive for designing insurance protection against electricity price fluctuations. Therefore, any actually developed model must be taking into account the discontinuous component. Following a widespread practice in literature, we ideologically work on a panel dataset, where the statistical units are the intraday different load periods (also called *time slots* in the remaining part of this document) and time represents the days. The main reasons why we handle each load period separately is that it permits us both to avoid the problem of modelling the intra-daily periodicity and not to care about clusterization issue, that is whether we classify closed spikes as one single spike or not. In fact, since the 'spiky' regime does not usually last for more than one day (Klüppelberg et al., 2010), we are indeed able to ignore the clusterization issue introduced by the Geman and Roncoroni (2006) price spike definition (see §2.2).

Let $\{P_{j,t}\}, j \in \{1, \ldots, J\}, t \in \{1, \ldots, T\}$ be an electricity prices time series (TS), where J is the number of daily load periods, fixed by the EM, and T is the number of days, meaning that $P_{j,t}$ is the electricity spot price at day t and load period j.

According to Lisi and Nan (2012), the starting point of our literature study is the assumption that the dynamics of electricity log prices can be represented by the following model, which is itself a subset of the 'industry standard' model defined by Janczura and Weron (2010):

$$\log P_{j,t} = D_{j,t} + v_{j,t}, \tag{2.1}$$

where $D_{j,t}$ is a nonstationary deterministic component and $v_{j,t}$ is a residual stationary stochastic component. The choice of the logarithmic transformation for electricity prices is largely justified by the EM literature – Klüppelberg et al. (2010), Schindlmayr (2005). In short, logarithmic function guarantees a progressive smoothing treatment for extreme values, directly proportional to the price value without loss of TS generality.

2.3.1 Models for price forecasting

Aggarwal et al. (2009) classify the price-forecasting models into three classes:

 Parsimonious stochastic models, such as ARMA models (e.g., SARI-MAX model, which tries to capture both the non-stationarity of the prices and the exogenous variables effect), heteroskedastic class models (e.g., GARCH, which models the conditional variance as time-changed), Markov regime switching models (e.g., 2-regime MRS model for normal and spike regime – Janczura and Weron (2010), Bierbrauer et al. (2005)).

- 2. *Regression or causal models*, based on the relationship between electricity prices and a number of independent variables. Such variables, known as explanatory variables, are identified by means of correlation analysis.
- 3. Artificial intelligence models, which adopt the modern data-mining techniques, such as neural networks and closest k-neighborhood categorization.

Hybrid models, which involve more than one of the previous classes, have been largely tested by different authors – Lora et al. (2002), Zhou et al. (2006). For instance, Zhou et al. (2006) extend the ARIMA approach by using a robust procedure to include error correction for improving accuracy of California EM spot price forecasting, supporting that such forecasting approach is very effective with satisfactory accuracy.

Almost every single approach for electricity price modelling agrees that it is possible to detect the spikes only once the predictable component is taken away from the TS. Indeed, we wonder how we can identify D_t and how we can detect spikes in $v_{j,t}$.

Although a lot of authors include some additional factor for electricity spot prices modelling (e.g., TS of loads – Kanamura and Ohashi (2007)), we assume for rest of this chapter that every information we are looking for stands within the TS. Such assumption means also that our perspective for detecting spikes does not comply with the market specifications, implying a strong limitation for the subsequently proposed framework.

2.3.2 Depuration of electricity prices time series

In a nutshell, *deputation* is the challenge of finding the deterministic component D_t , by using the adoption of various techniques.

Electricity spot prices exhibit the well known seasonality (at the annual, weekly and daily horizons) and mean-reverting behavior. Thus, the estimation of a component to deal with trend and seasonality is a first important issue. Unfortunately, electricity spot prices additionally show both extreme volatility and spikes. Hence, classical regression analysis estimation routines (e.g. OLS and GLS) should be carefully used, since they are very sensitive to extreme observations and outliers (Trueck et al., 2007). An additional note about seasonality is that some deregulated markets show a more marked seasonality than others because of the seasonal electricity consumption fluctuation patterns, which strictly depend upon the climatic characteristics of a certain area.

One of the most used parametric model for detrending and deseasonalizing electricity prices comes from ACF analysis evidences, leading to the use of a linear 'sinusoidal' model, which may also include a trend component (Klüppelberg et al., 2010, Weron, 2006, Schindlmayr, 2005). Since the least squares modelling lack of gaining efficient estimates due to the presence of extreme values is evidet, a data preprocessing procedure is largely adopted to gradually drop extreme price values. On one hand, such procedure permits to use the least squares estimators for better calibrating the parameters values, but on the other hand it must be handled carefully since it introduces further problems. Some connected relevant questions are following listed.

- How to drop the extreme values? Some authors adopt a fixed price threshold. However, Trueck et al. (2007) observes that "the choice of the levels themselves is non-trivial and rather arbitrary".
- How to replace the outlayers? A lot of authors replace the observed outlayers with the value of the threshold chosen for detecting them (Weron, 2006). Others replace them by one of the neighboring prices instead (Geman and Roncoroni, 2006). However, such procedure may lead to complications when the presence of consecutive outlayers is observed.
- When the chosen preprocessing procedure must stop? Several iterative procedures have been implemented to identify price spikes, such as recursive filters. Still, since a straightforward application of iterative techniques may bring to an overestimation of the number of price spikes, it is not clear whether better results in finding $D_{j,t}$ can be reached by stopping such procedures before the last iterations are performed.

One may use the Least Trimmed Squares (LTS) estimator or another robust procedure such as Huber methods instead, but it usually brings to unsatisfying results with EM time series, if compared to the combination of preprocessing and traditional least squares procedures.

In other cases, the different seasonalities are modeled with other techniques, such as dummy variables regression analysis (Schindlmayr, 2005) and combinations of robust estimators with non-parametric procedures (Trueck et al., 2007).

Other variants for estimating the nonstationary deterministic component $D_{j,t}$ are largely adopted, depending on the specific EM data set, the purpose of research and the forecast horizon of the developed model. For instance, some authors avoid the problem of intra-daily periodicity by focusing on the daily average prices, without considering each daily load period separately (Bierbrauer et al., 2005). Others do not take into account electricity prices which are observed during the non-business days, as they are not used to exhibit special futures like spikes and high volatility and they may consequently not be assumed as object of study.

The last issue concerns the estimated residual component $\hat{v}_{j,t}$. Its ACF and PACF may still show the presence of strong patterns. Although some authors include an additional stochastical trend component (Geman and Roncoroni, 2006), assuming a brownian motion presence in $v_{j,t}$ (which can influence the second-order component of $\hat{v}_{j,t}$), it is mostly assumed to be a first moment stationary process. As $v_{j,t}$ may show high time varying volatility, an intuitive possibility is to take the advantages of either regime switching models or non-parametric variance estimators. Another possibility consists in a further separation of $v_{j,t}$ in some components, assuming a specific stochastic nature for each of them. For instance, Klüppelberg et al. (2010) adopt an additive model for $v_{j,t}$, identifying a 'spiky' component $X_1(t)$ and a stationary sum of Ornstein-Uhlenbeck processes component $X_2(t)$.

2.3.3 Detection of spikes

Knowing that spikes are outliers within the electricity spot price TS, after the TS is 'depurated', a threshold has to be defined.

Basically, from our literature review (see §2.3), three different standpoints for detecting electricity price spikes time of occurrence have emerged:

- 1. Empirical standpoint: it is the simplest technique, which allows to not estimate the component $D_{j,t}$ for detecting price spikes. A threshold is being fixed on a certain either technical or empirical basis. Thus, every time the energy spot price oversteps such a treshold it is interpreted as a price spike. Christensen et al. (2011) use this procedure on the Australian EM data, fixing the threshold to 100 AUD/MWh.
- Financial standpoint: derived from the classical financial perspective which considers asset returns, once electricity spot price variations exceed a fixed threshold (e.g., 30% – Bierbrauer et al. (2005)), they are classified as spikes.
- 3. Quantile-based standpoint: a threshold is calculated by means of a quantilebased approach. Basically, a specific quantile defines whether a price value is classified as spike or not. Therefore, if we let τ be an arbitrarily small percentage, the largest $(1 - \tau)\%$ price observations can be identified as outliers (i.e. spikes). Two methods of quantile-based approach have emerged.
 - (a) The first method is the simplest: price spikes are directly detected from the empirical cumulative distribution functions of spot prices {P_{j,t}}_{t=1,...,T}, ∀j = 1,..., J.
 - (b) Once the nonstationary deterministic component D_{j,t} is identified by adopting a specific procedure (see §2.3.2), the second method aims to detect spikes from the residual stationary stochastic component v_{j,t} either considering the time-varying volatility or not. In the first case, the previously described 'first method' is adopted on the empirical cumulative distribution functions of {v_{j,t}}_{t=1,...,T}, ∀j = 1,...,J.. In the second case, the time varying volatility of v_{j,t} needs to be modeled by using either parametric procedures such as

GARCH modelling, or non-parametric techniques such as rolling volatility – Weron (2006).

2.4 A price spikes identification method

Finally, after the previous literature study (see §2.3), we have enough tools develop a method to identify D_t and consequently price spikes occurrences. Since each method has some advantages, but also drawbacks, we decide to propose a procedure for spike detection, which (empirically) produces the most satisfying trade-off between flexibility and results.

Our objective is to detect spikes, by firstly removing the whole predictable component from the electricity prices TS.

2.4.1 Model for electricity prices

Coherently with the literature review, we assume the following additive structure of Eq. 2.1:

$$p_{j,t} = LT_{j,t} + W_{j,t} + v_{j,t}, \ v_{j,t} \sim (0, \sigma_{j,t}^2), \tag{2.2}$$

where $\log P_{j,t} = p_{j,t}$ and $D_{j,t} = LT_{j,t} + W_{j,t}$, referring to the notation of Eq. 2.1. $LT_{j,t}$ is the long term component and $W_{j,t}$ is the weekly periodic component. Specifically, $LT_{j,t}$ includes the long run trend and the yearly periodic component; $\{v_{j,t}\}_{t=1,...,T}$ is assumed to be a mean-stationary process of uncorrelated variables $v_{j,1}, \ldots, v_{j,T}$ with time-varying variances, $\forall j = 1, \ldots, J$. Such process assumption is stronger than that for $v_{j,t}$ in Eq. 2.1, but rather necessary for the spike identification method that we are going to propose.

Definition 2.1. For each $j \in \{1, ..., J\}$, a *spike occurrence* is defined as every load period at which $v_{j,t}$ exceeds the threshold, that is fixed at the 95th percentile value of the $v_{j,t}$ distribution for off-peak hours and 90th for on-peak hours (see Def. 2.2).

Consequently, Def. 2.1 leads us to take into account the time-varying variance problem.

2.4.2 Method for cleaning the time series

The method adopted for cleaning the TS and consequently for isolating $v_{j,t}$ is based on a two-steps batch processing. By means of different techniques, we sequentially remove each component of Eq. 2.2 present within the time series. Once a substantial difference between off-peak hours and on-peak hours detection process is noticed, the process differs for each load period typology.

Definition 2.2. The *on-peak hours* are the load periods, where spikes are particularly extreme and occur with the highest rate. Spiky behavior is less marked in the *off-peak hours*.

Despite the literature usually labels on-peak hours and off-peak hours with two fixed daily intervals sets (e.g. from 7:00 AM to 8:00 PM as on-peak hours – Weron (2006)), we prefer to identify a load period as an on-peak hour whenever the off-peak hour procedure produces a particularly unsatisfactory estimation of the component $D_{j,t}$. Therefore, we need to carefully analyze the results of each single load period considered, instead of running the spike detection procedure in one single step.

First step The first step of our procedure aims mainly to remove the long term component $LT_{j,t}$. Many of the previously described techniques, such as the sinusoidal model proposed by Klüppelberg et al. (2010), have been tested on a real EM time series but the best results are produced by the non-parametric approaches.

Basically, local regression (LOESS), moving average and smoothing splines techniques show a good effectiveness in removing the long term components. As they lead to the nearly same results, for operational aspects¹, smoothing splines are chosen for off-peak hours.

Smoothing spline is a well known non-parametric procedure, based on the minimization (over the class of the twice differentiable functions) of the pe-

¹both local regression and moving average techniques do not allow the estimation of $LT_{j,t}$ for the first and the last part of a time series.

nalized least squares criterion

$$D_j(\mathcal{F}_j, \lambda_j) = \sum_{t=1}^T \left[p_{j,t} - \mathcal{F}_j(t) \right]^2 + \lambda_j \int_{-\infty}^{+\infty} \left\{ \frac{\partial^2 \mathcal{F}_j(x)}{\partial x^2} \right\}^2 \mathrm{d}x; \quad \lambda_j \ge 0, \ j = 1 \dots J_j$$

$$(2.3)$$

where λ_j is the penalization parameter, which influences the irregularity of the function \mathcal{F}_j . The minimizer is a natural cubic spline (Azzalini and Scarpa, 2004). To define λ_j , for each load period $j \in \{1, \ldots, J\}$, we computationally calculate the equivalent degrees of freedom df_j , minimizing the GCV (Generalized Cross Validation) criterion

$$GCV_j = \frac{s_j^2(df)}{(1 - \frac{df}{T})^2}; \ s_j^2(df) = \frac{1}{n} \sum_{t=1}^T (p_{j,t} - \hat{\mathcal{F}}_j(t; df))^2.$$

In short, our algorithm chooses the equivalent degrees of freedom as follows:

$$\hat{df}_j = \lfloor k_w \cdot \operatorname*{arg\,min}_{df} \{GCV_j\} \rfloor,\tag{2.4}$$

where k_w is a constant term, set to 2^{-1} , which guarantee the right long term smoothing, keeping the weekly seasonality component.

Concerning the on-peak hours, following the monthly smoothing proposed by Trueck et al. (2007), a 31 days rolling median procedure is chosen in order to remove the long term components, where the 2×15 prices on the tails are modeled with a progressive window-decreasing rolling median.

Second step The second step focuses on the remaining time series

$$p_{j,t} - \widehat{LT}_{j,t},\tag{2.5}$$

which still exhibits weekly seasonality, serial correlation and time-varying variance.

Differently than the first step, after many modelling attempts, two methods have emerged for their effectiveness. Although they bring to different results in spike detection, once they differ approaching philosophies, we decide to keep one for the off-peak hours and the other for the on-peak hours. The first one is based on a traditional parametric approach, while the second is supported by robust estimation procedures, following the modern data mining philosophy. Before describing the two methods, we point out that applying the Dickey–Fuller test for unit root presence on the residuals $p_{j,t} - \widehat{LT}_{j,t}$, the null hypothesis is largely rejected.

- By the parametric approach, we study the application of the seasonal autoregressive moving average (S–ARMA) model with generalized autoregressive conditional heteroscedasticity (GARCH) errors (see §A.2.1 and §A.2.2). The appropriate S–ARMA model is identified both by using the modern Tsay–Tiao procedure (Tsay, 2010) and by always considering the Akaike information criterion (AIC). A GARCH(1,1) model is chosen in any case. Reasonable variants of GARCH modelling, such as seasonal (S–GARCH) and threshold (T–GARCH), are not taken into account. In this approach, we model seasonality without considering any national holiday effect.
- By the robust approach, we model the seasonal effects by means of dummy variables, including each day of the week and also national holidays.

As in practice the dummy-based estimation does not directly produce the expected results, leading to a strong patterns presence in the ACF and PACF functions of $\hat{v}_{j,t}$, we must firstly smooth $p_{j,t} - \widehat{LT}_{j,t}$ in order to remove the time-varying volatility effect. In order to do that, the 31days rolling variance is estimated and $p_{j,t} - \widehat{LT}_{j,t}$ is rescaled by dividing it by the smoothed standard deviation. The 31-days rolling variance is arbitrary chosen because it is a good compromise, since a spiky regime lasts for at most two days (Klüppelberg et al., 2010).

The weekly seasonality is then estimated and removed as follows.

(a) The daily effect is calculated with a robust procedure, which excludes the national holidays. For each daily time series

$$\{p_{j,i} - \widehat{LT}_{j,i}\}_i; \ i = \lfloor \frac{t}{7} \rfloor + (t \% D), \ t = 1, \dots, T, \ D = 1, \dots, 7,$$

the arithmetic mean of is computed without the lowest value and the values that surpass the 95th percentile threshold, representing the Dth day effect ².

(b) When the weekly seasonality is subtracted from $p_{j,t} - \hat{L}\hat{T}_{j,t}$, a linear regression on the holiday dummy variable is computed to calculate

²For each day of the week D = 1, ..., 7, we have a time series as follows: $\{p_{j,i} - \widehat{LT}_{j,i}\}_i = \{p_{j,D} - \widehat{LT}_{j,D}, p_{j,D+7} - \widehat{LT}_{j,D+7}, p_{j,D+14} - \widehat{LT}_{j,D+14}, ...\}$

the holidays effect. Such effect is then also subtracted, leading to the estimation of $v_{j,t}$.

2.4.3 Testing model accuracy

Firstly, we should check whether or not the time series $p_{j,t} - \widehat{LT}_{j,t}$ exhibits weekly seasonality, autocorrelation and time-varying variance. Furthermore, according to Eq. 2.2 the time series $p_{j,t} - \widehat{LT}_{j,t}$ should be a mean-stationary process. Time-varying variance may be checked with the Engle's Lagrange Multiplier (LM) test for ARCH effects. Generally, by analyzing the ACF and PACF functions, we could have a good indicator of the long-term smoothing effect. Eventually, we could repeat the first step, changing the parameter k_w value in Eq. 2.4.

There are several methods to validate an ARMA model, such as examining the autocorrelation function of the estimated residuals and calculating the Ljung–Box portmenteau statistic for the estimated residuals. After the S– ARMA model estimation, the existence of a GARCH effect can be checked with the MCLeod–Li statistic for the squared estimated residuals.

The same tests may be used for validating the results of the robust approach for on-peak hours. However, since the robust procedure usually leads to reject the null hypothesis in almost every of the previous tests, we keep its result as a good compromise, without validating it through the tests but only looking at the ACF and PACF functions.

2.4.4 Some final observations

Before applying the framework to a real sampled electricity spot prices time series, some observations are pointed out:

- from our point of view, the estimation of $v_{j,t}$ is a crucial factor for spikes detecting;
- the effectiveness of a model upon estimating $v_{j,t}$ may strongly depend upon the market features;
- since some markets exhibit a more accentuated seasonal component, the best model to identify $v_{j,t}$ may be different, depending on the market;

- off-peak and on-peak hours are discriminating the detection procedure.
 At the same time, they are arbitrarily defined by using two procedures.
 Such method for defining on-peak and off-peak hours could have more effectiveness variants;
- concerning the spike identification procedure for the off-peak hours, the choice of estimating a SARMA–GARCH model for $p_{1,t} \widehat{LT}_{1,t}$ is rather arguable, as it could be an overparameterization;
- some markets are affected by a high frequency of extreme weather events, which can indirectly cause extreme spikes within the EM. In such cases we believe that the only use of robust procedures should be considered more reliable;
- the best fitting model for prices is not necessarily the best spike detecting one;
- the threshold chosen (both for on-peak and off-peak hors) for detecting price spikes is a justified but rather arbitrary choice. Clearly, alternative options are available in the literature, as written in §2.3.3.

2.5 Application: spike identification

The last part of this chapter concerns the application of the previously defined framework (see $\S2.4$) to a real electricity prices time series.

The data come from the British EM (APX-PUK), which is a half hourly day-ahead market (implying a value of 48 for the parameter J). The time series includes prices from the 1st of April 2005 to 31st of December 2010 and shows all the remarked features of electricity prices time series (e.g. seasonality, jumps presence, etc.). Specifically, a first inspection based on graphs and ACFs indicates that the spot prices TS show neither a well-defined long-run behavior nor a clear annual dynamics, while a marked common characterisic is a weekly periodic component and a very persistent autocorrelation function (see Fig. 2.1 and Fig. 2.2). Concerning spot price values, in spite of a median of £35.83/MWh, the maximum registered price is £553.30/MWh, which is more than 15 times higher.
Electricity prices (British EM)



Figure 2.1: The British EM spot prices time series, from the 1st of April 2005 to the 31st of December 2010. Prices are expressed in \pounds/MWh .



Figure 2.2: The APX-PUK spot prices ACF and PACF (load periods 12, 24, 36).

2.5.1 Applying the framework to the daily time series of the 1st half-hour

Our analysis now concentrates on the time series of the first load period, which has midnight and half past midnight as endpoints. Such TS has 2101 observations.

First step

The function of GCV_1 indicates 77 as the value of df_1 , which satisfies Eq. 2.4. The smoothing splines technique produced results are shown in the bottom panel of the Fig. 2.3. As we can see, the estimation of the long term component $\widehat{LT}_{1,t}$ seems to fit well to the $p_{1,t}$ time series. If we look at the remaining time series $p_{1,t} - \widehat{LT}_{1,t}$, as defined in Eq. 2.5, we can see that it is a mean-stationary process and it shows a strong correlation and time-varying volatility (Fig. 2.4).

Specifically, both the Ljung–Box and the Box–Pierce tests for the presence of autocorrelation strongly reject the null hypothesis (p-value $< 10^{-15}$). Concerning the heteroskedasticity presence, the null hypothesis of absence of ARCH effects is strongly rejected by the Engle's LM test at any lag (pvalue $< 10^{-15}$). Furthermore, we can notice how the ACF diagram remarks the weekly seasonality, as the values at lags 7,14,... are constantly and significantly far from zero. Before proceeding with the second step we test the presence of a stochastic trend component by means of the Augmented Dickey– Fuller (ADF) test. The test without constant and trend components gives the approximated value of -25.28, which falls outside the 99% region. The null hypothesis is therefore rejected.

Second step and spike detection

As written in §2.4.2, the seasonal ARMA order identification follows the Tsay–Tiao procedure that uses the extended autocorrelation function (EACF). EACF is based on the simple idea that if we can obtain a consistent estimate of the AR component, we can derive the MA component, which can be used to identify MA order by means of its ACF (Tsay, 2010).

Table 2.1 suggests the estimation of a S-ARIMA($(2, 0, 1) \times (1, 0, 1)_7$). The Ljung-Box test at lag 1 (p-value = 0.99) and 7 (p-value = 0.92) indicates that





Figure 2.3: GCV_1 function and smoothing splines applied to $p_{1,t}$ (\hat{df}_1 is set to 77). The top panel shows the GCV function for the spot prices TS of the 1st load period, which sets the minimum as 153.5. The grey line in the bottom panel indicates the smoothing splines.





Figure 2.4: The remaining time series $p' = p_{1,t} - \widehat{LT}_{1,t}$ (top panel,), with its ACF (middle panel) and PACF (bottom panel) functions.

Table 2.1: EACF table for $p_{1,t} - \widehat{LT}_{1,t}$, where 'x' denotes nonzero, 'o' denotes zero, and bold font remarks the upper left vertices of the o-triangles identified. AR indicates the AR order, MA indicates the MA order.

							Μ	Α					
\mathbf{AR}	0	1	2	3	4	5	6	7	8	9	10	11	12
0	x	х	х	0	0	0	х	0	х	х	х	х	х
1	x	0	х	0	0	0	0	0	х	х	0	х	х
2	x	0	0	х	0	0	0	х	0	х	0	х	х
3	x	х	0	0	0	0	0	х	0	х	0	0	0
4	x	х	0	0	0	0	0	х	0	х	0	0	0
5	x	х	0	х	0	0	0	0	х	х	0	0	х
6	x	х	х	х	0	0	х	х	0	х	0	0	х
7	x	х	0	х	х	х	х	х	0	х	х	х	х
8	x	х	0	х	х	х	х	х	0	0	0	0	0
9	x	х	0	х	х	х	х	х	х	0	0	0	0
10	x	х	х	х	0	х	х	0	х	0	х	0	0
11	x	х	х	х	0	0	х	0	х	х	х	0	0
12	v	v	v	v	v	0	v	v	v	0	v	v	0



Figure 2.5: Spike detection with S–ARMA model with GARCH errors (1st load period, threshold = 95%). The circles indicates a price value recognized as spike.

the model works perfectly and removes the weekly seasonality. However, the null hypothesis is still largely rejected by the Engle's LM test, as expected.

The model is completed with the GARCH(1, 1) part, necessary to estimate the features of $v_{1,t}$. The estimated parameter values are $\omega = 1.9 \cdot 10^{-4}$ (p = 0.0037), $\alpha_1 = 7.7 \cdot 10^{-2}$ ($p = 3.5 \cdot 10^{-8}$) and $\beta_1 = 9.1 \cdot 10^{-1}$ ($p < 10^{-15}$). The stationarity condition is satisfied by the parameter values.

The residuals time series

$$\epsilon_{1,t} = \frac{\hat{v}_{1,t}}{\hat{\sigma}_{1,t}}$$

shows no correlation and the absence of ARCH effects (the McLeod–Li and LM tests accept the null hypothesis at any lag). Therefore, we can detect spikes from $\epsilon_{1,t}$, following Def. 2.1. The results are shown in Fig. 2.5.

The time interval we have considered is technically always associated to the non-peak hours set. It means that the spikes have a relative low amplitude, as we can see in Fig 2.5.

2.5.2 Applying the framework to the daily time series of the 24th half-hour

Once the procedure for detecting spikes differs for on-peak hours, we believe that it is interesting also to show an application of our framework upon a load period which is typically subjected to a marked spiky behavior.



Figure 2.6: Spike detection with S–ARMA model with GARCH errors (24th load period, threshold = 90%). The circles indicates a price value recognized as spike.

As expected, the results of applying the off-peak hours version of our framework are unsatisfactory. As we can see in Fig. 2.6, even if we set the threshold to 90%, there are some evident spikes to the naked eyes which are not computationally detected.

Applying the second procedure for deseasonalizing (i.e. the on-peak hours one), we achieve a different and more reliable result, as we can see in Fig. 2.7.

2.5.3 Results

The framework explained in 2.4 is applied for each daily time series. 48 different TS are analyzed. Table 2.2 summarizes the reached results.

During our analysis, we find an exception within the 13th load period. Apparently, there is no matter to classify such load period as an on-peak hour. However, the first procedure is heavily influenced by the presence of a cluster of downward shocks. Therefore, we decided to apply the robust procedure, keeping the threshold set to 95%.



ROBUST procedure for spike detection (24th time slot - British EM)

Figure 2.7: Spike detection by means of the robust procedure (24th load period, threshold = 90%). The circles indicates a price value recognized as 'spike' from the procedure.

Table 2.2: British EM spike detection summary, where J indicates the load period, *peak* whether a load period is on-peak or not, ADF, L-B test and LM test the augmented Dickey–Fuller, Ljung–Box and Engle's LM test p-values, S–ARMA and GARCH the model orders, and Thr. the threshold

j	peak	ADF	S-ARMA	GARCH	GARCH L–B test		LM test	Thr.
		(p-value)	(order)	(order)	(p-value)	(p-value)	(p-value)	
1	no	$< 0.01^{***}$	$(2,1)\times(1,1)_7$	(1, 1)	0.071^{*}	0.275	0.763	95%
2	no	$< 0.01^{***}$	$(1,3)\times(1,1)_7$	(1, 1)	0.195	0.479	0.708	95%
3	no	$< 0.01^{***}$	$(1,3)\times(2,2)_7$	(1, 1)	0.308	0.697	0.902	95%
4	no	$< 0.01^{***}$	$(1,3)\times(2,2)_7$	(1, 1)	0.790	0.566	0.589	95%
5	no	$< 0.01^{***}$	$(1,1) \times (1,2)_7$	(1, 1)	0.877	0.477	0.528	95%
6	no	$< 0.01^{***}$	$(2,1) \times (1,3)_7$	(1, 1)	0.874	0.338	0.304	95%
7	no	$< 0.01^{***}$	$(1,3)\times(0,1)_7$	(1, 1)	0.894	0.261	0.149	95%
8	no	$< 0.01^{***}$	$(1,3)\times(1,1)_7$	(1, 1)	0.598	0.122	0.417	95%
9	no	$< 0.01^{***}$	$(1,1)\times(1,1)_7$	(1, 1)	0.486	0.047^{**}	0.265	95%
10	no	$< 0.01^{***}$	$(1,1)\times(1,1)_7$	(1, 1)	0.353	0.040**	0.261	95%
11	no	$< 0.01^{***}$	$(2,4) \times (2,1)_7$	(1, 1)	0.661	0.065^{*}	0.188	95%
12	no	$< 0.01^{***}$	$(2,3)\times(1,4)_7$	(1, 1)	0.796	0.352	0.059^{*}	95%
13	yes	$< 0.01^{***}$	_	_	_	_	_	95%
14	yes	$< 0.01^{***}$	_	_	_	_	_	90%
15	yes	$< 0.01^{***}$	_	_	_	_	_	90%
16	yes	$< 0.01^{***}$	_	_	_	_	_	90%
17	yes	$< 0.01^{***}$	_	_	_	_	_	90%
18	yes	$< 0.01^{***}$	_	_	_	_	_	90%
19	yes	$< 0.01^{***}$	_	_	_	_	_	90%
20	yes	$< 0.01^{***}$	_	_	_	_	—	90%
21	yes	$< 0.01^{***}$	_	_	_	_	_	90%
22	yes	$< 0.01^{***}$	_	_	_	_	—	90%
23	yes	$< 0.01^{***}$	_	_	—	_	_	90%
24	yes	$< 0.01^{***}$	_	_	_	_	—	90%
25	yes	$< 0.01^{***}$	_	_	—	_	_	90%
26	yes	$< 0.01^{***}$	_	_	_	_	—	90%
27	yes	$< 0.01^{***}$	_	_	—	_	_	90%
28	yes	$< 0.01^{***}$	-	_	_	_	_	90%
29	yes	$< 0.01^{***}$	_	_	_	_	_	90%
30	yes	$< 0.01^{***}$	_	_	_	_	_	90%
31	yes	$< 0.01^{***}$	_	_	_	_	_	90%

32	yes	$< 0.01^{***}$	_	—	-	—	-	90%
33	yes	$< 0.01^{***}$	_	_	_	_	—	90%
34	yes	$< 0.01^{***}$	_	_	_	_	—	90%
35	yes	$< 0.01^{***}$	_	_	_	_	—	90%
36	yes	$< 0.01^{***}$	_	_	_	_	_	90%
37	yes	$< 0.01^{***}$	_	_	_	_	_	90%
38	yes	$< 0.01^{***}$	_	_	_	_	_	90%
39	yes	$< 0.01^{***}$	_	_	_	_	_	90%
40	no	$< 0.01^{***}$	$(1,1)\times(1,0)_7$	(1, 1)	0.373	0.678	0.985	95%
41	no	$< 0.01^{***}$	$\left(1,1\right)\times\left(0,1\right)_{7}$	(1, 1)	0.273	0.241	0.992	95%
42	no	$< 0.01^{***}$	$(2,1)\times(1,2)_7$	(1, 1)	0.392	0.783	0.994	95%
43	no	$< 0.01^{***}$	$(2,1)\times(1,2)_7$	(1, 1)	0.535	0.818	0.999	95%
44	no	$< 0.01^{***}$	$(2,1)\times(1,2)_7$	(1, 1)	0.276	0.016^{**}	0.991	95%
45	no	$< 0.01^{***}$	$(1,0)\times(1,2)_7$	(1, 1)	0.898	0.320	0.946	95%
46	no	$< 0.01^{***}$	$\left(2,1\right)\times\left(1,0\right)_{7}$	(1, 1)	0.703	0.274	0.874	95%
47	no	$< 0.01^{***}$	$(1,2)\times(1,1)_7$	(1, 1)	0.181	0.712	0.783	95%
48	no	$< 0.01^{***}$	$(1,2)\times(1,2)_7$	(1, 1)	0.822	0.318	0.982	95%

As we can see in Table 2.2, the off-peak hours procedure exhibits surprising results on deseasonalizing the TS. In fact, the largest part of the statistical tests made marks how the SARMA–GARCH modelling is appropriate for off-peak hours, with a high *p*-value.

Chapter 3

The Model

3.1 Introduction

Electricity spot prices, loads, production figures, etc., are sampled 24 hours a day, 365 days a year. This gives us a unique opportunity to apply statistical methods for time series in the way they were meant to be used.

In this chapter we introduce the duration modelling starting with the ACD model, which has necessarily to be known in order to understand the ACH model. The cases study in Chapter 4 adopt the latter model.

Duration Models

The duration models are born to model and forecast time intervals, which occur between two particular market events. For instance, events may be financial transactions, price fluctuations and shifts, trading volumes, etc. (Lisi, 2010)

Let us consider an orderly marked point process (see §A.2), where events occur at random times $t_0 < t_1 < \cdots < t_n < \ldots$, where t_i represents the time at which the (i + 1)th event occurred.

Definition 3.1 (Duration). Given two events E_i and E_{i-1} , which occur at times t_i and t_{i-1} , the time interval $x_i \in \mathbb{R}^+$ between E_i and E_{i-1} is called *i*th duration, such that $x_i = t_i - t_{i-1}$.

The main pillar, on which the duration models stand, is the concept that the time intervals $\{x_i\}_{i=1,...}$ can be handled as random variables. One of the most popular duration model is the ACD model, which is explained in the following section.

3.2 The standard ACD Model

Suppose we have an event counter N(t), which is a function of time and counts the events that have occurred in the interval $(t_0, t]$. Let \mathcal{F}_t represent the whole information available in the history of the process over $[t_0, t]$, which comprises past durations up to and including x_t , but also some pre-determined variables suggested by the process microstructure. The conditional intensity function is defined as

$$\lambda(t \mid \mathcal{F}_t) \equiv \lim_{\Delta t \to 0^+} \frac{\mathbb{P}[N(t + \Delta t) > N(t) \mid \mathcal{F}_t]}{\Delta t}.$$
(3.1)

The point process econometric analysis typically deals with an appropriate parametrization of $\lambda(\cdot)$, with the aim of determining which exogenous variables (if any) drive the intensity of the process, and the extent to which this intensity is influenced by its history (Christensen et al., 2011).

In absence of memory, we may assume that the TS is modeled by an inhomogeneous Poisson process, which implicates that the durations are assumed to be determined uniquely by the exogenous variables.

3.2.1 The model and its assumption

The Autoregressive Conditional Duration model (ACD) uses the idea of Bollerslev's GARCH model to study the dynamic structure of the (adjusted¹) durations. The process is therefore assumed not to be memoryless.

Defining $\{z_0, z_1, \ldots, z_n, \ldots\}$ as the sequence of marks associated with the random times $\{t_0, t_1, \ldots, t_n, \ldots\}$, let the *conditional expected duration* be

$$\psi_{N(t)} \equiv \mathbb{E}\left[x_{N(t)} \mid \mathcal{F}_{t-1}\right] = \psi_t(\{x_0, x_1, \dots, x_{N(t)-1}\}, \{z_0, z_1, \dots, z_{N(t)-1}\}).$$
(3.2)

All the temporal dependence of the duration process is thus captured by the conditional expected duration (Pacurar, 2006).

Then, Engle and Russell (1998) define ACD model as

$$x_{N(t)} = \psi_{N(t)} \epsilon_{N(t)}, \quad \epsilon_i \sim \text{i.i.d. s.t. } \mathbb{E}\left[\epsilon_i\right] = 1.$$
(3.3)

¹The original framework of Engle and Russell (1998) was born to model high-frequency financial durations, which need to be adjusted.

The multiplicative error structure and the non-negativity of the duration sequence require that the density function of ϵ (with parameters θ_{ϵ}) has a nonnegative support.

Given $p_0(t) = p_{\epsilon}(t; \theta_{\epsilon})$, let S_0 be the associated survival function (see A.3). The Eq. 3.1 of the conditional intensity function can be expressed as

$$\lambda(t \mid \mathcal{F}_t) = \lambda_0 \left(\frac{t - t_{N(t)}}{\psi_{N(t)+1}}\right) \frac{1}{\psi_{N(t)+1}}; \ \lambda_0(t) = \frac{p_0(t)}{S_0(t)},$$
(3.4)

where λ_0 is called the *baseline hazard*. Therefore, the past history influences the conditional intensity by both a multiplicative effect and a shift in the baseline hazard (Engle and Russell, 1998). For instance, taking ϵ as exponentially distributed means that $\lambda_0 = 1$ everywhere, so that the expression for conditional intensity simply becomes

$$\lambda(t \mid \mathcal{F}_t) = \frac{t - t_{N(t)}}{\psi_{N(t)+1}^2} = \psi_{N(t)+1}^{-1}.$$
(3.5)

The parametrization proposed by Engle and Russell (1998) for standardized duration, known as ACD(m, q), relies on a linear parametrization on Eq. 3.2:

$$\psi_{N(t)} = \omega + \sum_{j=1}^{m} \alpha_j x_{N(t)-j} + \sum_{k=1}^{q} \beta_k \psi_{N(t)-k}, \qquad (3.6)$$

which depends on m past durations and q past expected durations. Note that some authors introduce pre-determinated variables in the ACD Eq. 3.6. This change in specification is primarily made in order to test for some hypotheses on the market's microstructure (Weisang, 2008).

The model parameters are subject to the following constraints: $\omega > 0$, $\alpha \ge 0$ and $\beta \ge 0$. Once the autoregressive nature of the ACD process allows us to formulate it as an ARMA process (see §A.3), we have the sufficient condition for x_i to be covariance-stationary:

$$\sum_{j=1}^{m} \alpha_j + \sum_{k=1}^{q} \beta_j < 1.$$
(3.7)

Furthermore, ARMA representation allows durations forecasting. It also permits to find both the unconditional mean of x_i and its conditional variance:

$$\mathbb{E}\left[x_{i}\right] = \frac{\omega}{1 - \sum_{j=1}^{m} \alpha_{j} + \sum_{k=1}^{q} \beta_{j}} , \text{ Var}[\mathbf{x}_{i} \mid \mathcal{F}_{i-1}] = \psi_{i}^{2} \text{Var}[\epsilon_{i}].$$
(3.8)

3.3 The ACH Framework

Let the energy prices time series to be given and spikes be the event of interests. The econometric model outlined in this section specifies, for our case study, the probability of observing a spike as a function of the process history and a set of exogenous variables (e.g. loads and weather variables, such as temperatures, etc.).

Engle and Russel's ACD specification posed the question, How much time is expected to pass before the next event occurs? Hamilton and Jordá's ACH specification reframed differently the question as, How likely is it that the target will change tomorrow, given all that is known today?

3.3.1 Why ACH?

As already mentioned, the ACD model and its variants assume that events can occur at any instant in time, binding the modelling only to continuous underlying processes. In the context of the incidence of electricity price spikes, the interval of interests is fixed, and all transactions within this fixed interval are settled at the pool price for that interval (Christensen et al., 2011). At most one event can occur within an interval. Therefore, slightly modifying the reframed question of Hamilton and Jordá (2002), we would like to know whether or not an event occurs at a certain time interval tomorrow, given all that is known up today.

Fig. 3.1 outlined the different stylized timelines of both ACD and ACH models.

3.3.2 The model and its assumption

Since we are going to analyze fixed time intervals in the timeline, instead of adopting the conditional intensity function defined in Eq. 3.1, in the ACH model it is necessary to think in terms of conditional hazard.

Definition 3.2 (Conditional Hazard). Christensen et al. (2011) define the *conditional hazard* as

$$h_{t+1} = h(t+1 \mid \mathcal{F}_t) = \mathbb{P}[N(t+1) > N(t) \mid \mathcal{F}_t],$$



Figure 3.1: The ACD (top panel) and ACH (bottom panel) timelines.

which represents the probability of an event occurring in a given interval (e.g. a given load period of a given day), conditional on \mathcal{F}_t . The past history of events is now interpreted in terms of the discrete process.

Definition 3.3. Let y_t be a random variable defined as

$$y_t = \begin{cases} 1, & \text{there is a spike occurring at the } th time \\ 0, & \text{otherwise} \end{cases}$$

Using $\{y_t\}_{t=1,\dots}$ we can rewrite the conditional hazard of Eq. 3.2 as

$$h_{t+1} = \mathbb{P}[y_{t+1} = 1 \mid \mathcal{F}_t]. \tag{3.9}$$

Rewriting the ACD conditional expected duration

We aim now to rewrite Eq. 3.6, so that it is indexed by calendar time t rather than by a count of the cumulative number of target changes i (Hamilton and Jordá, 2002). For motives of simplification we will use Hamilton and Jordá's notation. **Definition 3.4.** Let $w_{j,t}$ be the time of the *j*th most recent spike time starting from time *t*. Then we have:

$$w_{1,t} = ty_t + (1 - y_t)w_{1,t-1}$$
$$w_{2,t} = y_t w_{1,t-1} + (1 - y_t)w_{2,t-1}$$
$$\vdots$$
$$w_{j,t} = y_t w_{j-1,t-1} + (1 - y_t)w_{j,t-1}.$$

Therefore, $w_{1,t-1}-w_{2,t-1}$ would correspond to the length of the most recent duration that has been completed prior to time t (i.e. $x_{N(t)}$).

In discrete time, Eq. 3.6 can then be written as follows:

$$\psi_t = \tilde{\omega} + \sum_{j=1}^m \alpha_j (w_{j,t-1} - w_{j+1,t-1}) + \sum_{k=1}^q \beta_k \psi_{w_{k,t-1}}.$$
 (3.10)

Eq. 3.10 makes the value of ψ_t change only when a new spike has been observed within the previous load period (i.e., if $y_{t-1} = 1$).

Assumption 3.3.1. The standardized durations $\epsilon_i = \frac{x_i}{\psi_i}$ are i.i.d. (independent and identically distributed), with $\epsilon_i \stackrel{\text{iid}}{\sim} Exp(1)$.

Defining hazard function If the times of the previous spikes were the only informations available in \mathcal{F}_t , the hazard rate would not change until the next price spike. In that case, under the Assumption 3.3.1, the expected length of time until the next price spike occurrence would be

$$\psi_t = \sum_{j=1}^{\infty} j(1-h_t)^{j-1} h_t = \frac{1}{h_t}$$

The hazard rate that is directly implied by the ACD model (Eq. 3.5) would then be

$$h_t = \frac{1}{\psi_t}.\tag{3.11}$$

However, we have seen in Eq. 3.2 that there could be a sequence of further marks associated with the spikes times. So, let \mathbf{z}_t denote the vector of (exogenous) variables that are known at time t; the hazard rate in Eq. 3.11 can be generalized as follows:

$$h_t = \frac{1}{\Lambda(\psi_{N(t-1)} + \boldsymbol{\gamma}' \mathbf{z}_{t-1})},$$
(3.12)

where $\Lambda(\cdot)$ is chosen to ensure that h_t is a probability.

We assume that the first element of \mathbf{z}_{t-1} is a constant term and normalize the first element of the vector $\boldsymbol{\gamma}$, namely γ_1 , relative to unity and likewise normalize $\tilde{\omega}$ in Eq. 3.10 to zero. Therefore, we work with a specific version of Eq. 3.10, without the constant term ω :

$$\psi_t = \sum_{j=1}^m \alpha_j (w_{j,t-1} - w_{j+1,t-1}) + \sum_{k=1}^q \beta_k \psi_{w_{k,t-1}}.$$
(3.13)

Since h_t is a probability, it is important to ensure that a numerical search procedure does not select a value of h_t outside of the interval [0, 1] (Hamilton and Jordá, 2002). Therefore, taking into account Eq. 3.13, the Hamilton and Jordá's form of the function $\Lambda(\cdot)$ is chosen, such that

$$\Lambda(v_t) = \begin{cases} 1.0001, & v_t \le 1\\ 1.0001 + \frac{2\Delta_0(v_t - 1)^2}{\Delta_0^2 + (v_t - 1)^2}, & 1 < v_t \le 1 + \Delta_0 \\ 0.0001 + v_t, & v_t \ge 1 + \Delta_0 \end{cases}$$
(3.14)

which ensures that $h_t \in [0, 1)$ and that the resulting expression is always differentiable indeed.

Assumption 3.3.2. The order chosen for the ACH specification is (1,1), meaning that m = 1 and q = 1 in Eq. 3.13.

Final parametrization of durations

Since there might be a phenomenon of persistence of the effect of an observation long after it (which is characteristic of processes with unit roots), when the ACF of the spike durations is statistically significant for a long term, we can occur in estimation and specification problems (Weisang, 2008). To extend the use of the ACH model, Christensen et al. (2011) propose to use the Box-Cox transformation (see §A.4), defined by Fernandes and Grammig (2006) as follows:

$$\psi_{N(t)+1}^{v} = \alpha_1 \left[|\epsilon_{N(t)} - b| + c(\epsilon_{N(t)} - b) \right]^{\lambda} \psi_{N(t)}^{v} + \beta_1 \psi_{N(t)}^{v}; \ v > 0,$$
(3.15)

where b is the shift parameter, c is the rotation parameter, λ is the shape parameter and v is the Box-Cox transformation parameter. Such representation is also known in literature as *augmented ACH* (AACH) representation.

The main advantage of augmented representation is that Eq. 3.15 nests the original ACH specification in Eq. 3.10 as a special case, providing a more flexible model for the conditional expected duration.

The Box-Cox transformation parameter v determines the shape of the transformation, with $v \ge 1$ representing a convex transformation and $v \le 1$ representing a concave transformation. Asymmetric responses in duration shocks are permitted through the shift parameter b and the rotation parameter c. The shape parameter λ assumes a similar role as v, with $\lambda \ge 1$ iducing convexity and $\lambda \le 1$ inducing concavity (Tse and Tao, 2010).

The generalized AACH representation of Eq. 3.15 allows us to use the same specification of $\psi_{N(t)+1}$ originally adopted by Christensen et al. (2011). In fact, since Eq. 3.3 implies that $\psi_{N(t)} = \frac{x_{N(t)}}{\epsilon_{N(t)}}$, if $\lambda = -v$, b = 0 and c = 0, Eq. 3.15 becomes

$$\psi_{N(t)+1}^{v} = \alpha_1 x_{N(t)}^{v} + \beta_1 \psi_{N(t)}^{v}.$$
(3.16)

Relation to continuous time model

Hamilton and Jordá (2002) demonstrate that the ACH model is the discrete time equivalent of the ACD model. In fact, once the time interval used to discretize, calendar time becomes arbitrary small, the ACH model includes the ACD model as a special case.

3.3.3 ACH(1,1) framework summary

The ACH model, which fits the probability of a price spike occurrence, comprises Eqs. 3.12, 3.14 and 3.15, with parameters vector $\boldsymbol{\theta} = [\alpha_1, \beta_1, \boldsymbol{\gamma}, v, b, c, \lambda]$, where $\alpha_1, \beta_1 \ge 0, v > 0$ and $\alpha_1 + \beta_1 < 1$.

3.4 Parameters estimation

The parameters vector $\boldsymbol{\theta}$ is estimated by the maximum likelihood, based on the log-likelihood function, with standard errors computed using the typical sandwich procedure.

The conditional probability density function of the variable y_t (see Def. 3.3)

can be written as

$$\mathbb{P}[y_t = i \mid \mathcal{F}_{t-1}; \boldsymbol{\theta}] = h_t^i (1 - h_t)^{1-i}.$$

Therefore, in a sample of $T \to \infty$ time intervals, the log-likelihood function is

$$\ln \mathcal{L}(\boldsymbol{\theta}) = \sum_{t=1}^{T} \ell_t(\boldsymbol{\theta}) = \ln \ell(y_1 = i_1, \dots, y_k = i_k \mid \mathcal{F}_{k-1}; \boldsymbol{\theta}) + \sum_{t=k}^{T} \ell_t(\boldsymbol{\theta}) \simeq$$

$$\simeq \sum_{t=k}^{T} (i_t \ln h_t(\boldsymbol{\theta}) + (1 - i_t) \ln (1 - h_t(\boldsymbol{\theta}))),$$
(3.17)

where the term $k \in \mathbb{N}$ is the smallest value of time between 1 and T, at which a spike is observed (i.e. N(k) = 1).

The result obtained in Eq. 3.17 is due to the fact that the component $\ln \ell(y_1, \ldots, y_k \mid \mathcal{F}_{k-1}; \boldsymbol{\theta})$ does not influence the log-likelihood for $N(T) \to \infty$. Robustness of numerical maximization routines likely requires $\alpha_1 \ge 0, v > 0$ and $\beta_1 \in [0, 1]$, while the mean stationarity condition is $\alpha_1 + \beta_1 < 1$. The quasi-maximum likelihood estimator is then

$$\tilde{\boldsymbol{\theta}} = [\tilde{\alpha_1}, \tilde{\beta_1}, \tilde{\boldsymbol{\gamma}}, \tilde{v}, \tilde{b}, \tilde{c}, \tilde{\lambda}] = \arg\max_{\boldsymbol{\theta}} \Big\{ \sum_t \ell_t(\boldsymbol{\theta}) \Big\}.$$
(3.18)

3.4.1 Numerical aspects

R is well-suited for programming maximum likelihood routines. Indeed, there are several procedures for the numerical optimization of the likelihood functions. The core of our estimation procedure relies on the maxLik command from the maxLik package. Such command computes a numerical optimization and does not need the derivatives declaration. Linear constraints are embedded in the log-likelihood function, returning NA value as the constraints are not satisfied (see §B.1). The Hessian is automatically computed and used to obtain the observed Fisher information matrix and the standard errors of $\tilde{\boldsymbol{\theta}}$ (Steenbergen, 2006).

Following Hamilton and Jordá (2002) original ACH specification, the recursion in Eq. 3.15 starts by setting the initial values $x_{-1} = \bar{x}$ and $\psi_{-1} = \bar{\psi}$, where

$$\bar{\psi} = \frac{\alpha_1 \bar{x}}{1 - \beta_1}.\tag{3.19}$$

Furthermore, the value of Δ_0 in Eq. 3.14 is set to 0.1.

3.5 Forecasting

Given a load period, predicting the presence of a spike for any given day requires answering the question: Is the MCP going to show a spike this day or leave within a normal regime?

One advantage of the ACH framework is that it generates a closed-form expression for the one-period-ahead forecasting of the target y_{t+1} (Hamilton and Jordá, 2002). Specifically

$$\hat{y}_{t+1} = \begin{cases}
1, & 0.5 \le \hat{h}_{t+1} \le 1 \\
0, & 0 \le \hat{h}_{t+1} < 0.5
\end{cases},$$

$$\hat{h}_{t+1} = \frac{1}{\Lambda(\psi_{N(t)} + \tilde{\gamma}' \mathbf{z}_t)}.$$
(3.20)

We use the same procedure of Christensen et al. (2011) for spike forecasting. Once the ACH model parameters are estimated using the market data, without including the last part of the TS for the estimation procedure, the model is used to provide one-step-ahead forecasts. For a number of reasons, however, the model parameters are not re-estimated step by step, mostly because the sheer size of the estimation sample makes the model estimation a rather complex task which requires a long runtime.

Still, since the barrier value for hazard function h_t (set to 0.5 in Eq. 3.20) is an arbitrary choice. A brief analysis on the barrier value is outlined in Chapter 5.

Chapter 4

Application

In this chapter we present some empirical results upon which this document is centered, namely the electricity price spikes occurrences modelling and forecasting. Three EMs are basically considered: the Australian National Electricity Market (NEM), the Italian electricity market and the British electricity market (APX-PUK).

Before applying the ACH modelling by using the load period approach (i.e. by differently modelling each load period time series of spot prices) in §4.2, §4.3 and §4.5, an ACH estimation on the whole spot prices TS of NEM is proposed in §4.1, following the approach adopted by Christensen et al. (2011). Concerning Italian EM and NEM modelling, a further comparison between the ACH and a benchmark Logit model is proposed in §4.4.

4.1 Introduction: Christensen et al. (2011)

The application part of this document is introduced by trying to replicate the research of Christensen et al. (2011). Since ACH model has been specified in Chapter 3, before applying it to other EMs by using a different load period based approach, it is important to verify whether we can achieve the same results by considering the same data that were used by the authors who have inspired our essay. Therefore, we study the same EM (i.e. Australian NEM) by using the same method for spike detection, the same set of exogenous variables¹ and the same time period for forecasting (i.e. third quarter of 2007), withour separating every load period (i.e., by estimating the parameters with the whole spot prices TS). However, the time period for estimating ACH(1,1) parameters is slightly different: it starts from 1st January 2003 instead of 1st

¹Since Christensen et al. (2011) do not allow the exact construction of the exogenous variables (i.e. some informations are missing within the paper), we construct such variables as accurate as we can by using the available information.

March 2001.

Since the AACH specification of conditional expected duration (Eq. 3.15) includes a wider class of models that the one used by the authors (Eq. 3.16), Table 4.1 reports the results of different estimations, considering both specifications. Still, since the hazard rate specification in Eq. 3.12, originally proposed by Hamilton and Jordá (2002), has been modified by Christensen et al. (2011) to

$$h_t = \frac{1}{\Lambda(\psi_{N(t-1)} + \exp\left(\boldsymbol{\gamma}' \mathbf{z}_{t-1}\right))},\tag{4.1}$$

the estimation results shown in Table 4.1 consider both of such hazard rate specifications.

Specifically, our attempts are following listed and descripted.

- 1. Attempt 1: the ACH and hazard rate specifications proposed by Christensen et al. (2011) are adopted for the numerical estimation of $\boldsymbol{\theta}$.
- Attempt 2: the ACH, the hazard rate specifications and the estimates of θ, which have been proposed and reported by Christensen et al. (2011), are adopted.
- Attempt 3: the ACH, the hazard rate specifications and the estimates of α₁ and β₁ proposed by Christensen et al. (2011) are adopted, while the other parameters (i.e. γ and v) are numerically estimated.
- 4. Attempt 4: the ACH and hazard rate specifications proposed in Chapter 3 are adopted for the numerical estimation of $\boldsymbol{\theta}$ (see §4.4.4 for further analyses).

As we can see in Table 4.1, although Christensen et al. (2011) reports that the ACH model forecasts about 48% of the spikes with a relatively low number of false alarms (approx. 19%), in our analyses, since the achieved forecasting results differ from the considered research, the lack of forecasting is clear. Possibly, it may be caused both by the different time period chosen for estimating the ACH(1,1) parameters and by the different exogenous variables adopted (the exact method we used to construct such variables is described in §4.2.1). Therefore, our analyses induce us to believe that the method chosen by the Australian authors in order to forecast spike occurrences strongly depends on the time period analyzed.

Table 4.1: ACH(1,1) estimation results. *Estimation* indicates the estimation period, *Forecast* indicates the forecast horizon, *Model* indicates the stardardized duration ψ and hazard rate h models. *Spike detected* indicates the number of spikes detected by the ACH(1,1) model and the number of spikes observed within the TS, while *False alarm* indicates the number of spikes detected which are not observed for real over the total number of estimated spikes. *Christensen* row shows the original results obtained by Christensen et al. (2011).

Description	Dat	ta	Mo	del	Forec	ast
	Estimation	Forecast	ψ	h	Spike detected	False alarm
Christensen	2001-2007	3 months	Eq. 3.16	Eq. 4.1	144/299	34/178
Attempt 1	2003-2007	3 months	Eq. 3.16	Eq. 4.1	131/299	549/680
Attempt 2	2003-2007	3 months	Eq. 3.16	Eq. 4.1	66/299	83/149
Attempt 3	2003-2007	3 months	Eq. 3.16	Eq. 4.1	89/299	103/192
Attempt 4	2003-2007	3 months	Eq. 3.15	Eq. 3.12	15/299	22/37

Even though Table 4.1 shows that the model specifications of Eqs. 3.16 and 4.1 adopted by Christensen et al. (2011) work better in forecasting than the original specifications of Eqs. 3.15 and 3.12 proposed by and Hamilton and Jordá (2002) and Fernandes and Grammig (2006), in the next part we adopt the latter specifications. Basically, our choice is justified by four facts:

- the conditional expected duration specification of Eq. 3.16 is included in Eq. 3.15 as a special subcase (see §3.3 for further details);
- 2. the original hazard rate specification of Eq. 3.12 permits to evaluate the real effect of the chosen exogenous variables upon the probability of a spike occurring at a given time t;
- 3. the spiky regime in year 2007, with 1,495 observed spikes from 1st January to 30th September, is a totally anomalous year (an average of about 200 spikes per year is normally observed). We believe that the extraordinary fitting of the ACH specification chosen by Christensen et al. (2011) is partially due to the anomalous spiky activity of the year 2007.

Thus, although the number of parameters is higher, by using the generalized specifications for ψ_t and h_t we should theoretically reach better fits and forecasts.

4.2 The Australian National Electricity Market

The reason why we start with the Australian EM analysis, is that it is the same market studied by Christensen et al. (2011). Although the specified ACH framework is not exactly the same adopted by the authors, since an approach which models every single load period is applied and the original ACH and hazard specifications are adopted (see §3.3 and §4.1), we are going to study how the ACH model works with a different dataset. Moreover, in the following part we gather that the Australian EM spot prices TS is particularly suitable for spike occurrences analysis, as it exhibits a singular spiky behavior than other EMs such as Italian EM end APX-PUK.

Since 1998, the Australian National Electricity Market (NEM) operates as a wholesale market for the supply of electricity to retailers and end-users in Queensland, New South Wales, the Australian Capital Territory, Victoria and South Australia (Weron, 2006). In 2005 the NEM grew further with the entrance of Tasmania. The regions are connected in the electricity network, so that if the local demand exceeds the local supply or the electricity in a neighboring region is sufficiently inexpensive to warrant transmission, electricity is imported or exported between regions, subject to the physical constraints (Christensen et al., 2011).

The NEM spot market is a day-ahead market, which operates with half-hourly load periods (i.e., J = 48). Precisely, prior to 12:30 pm on the day before production, the spot prices are established matching supply and demand sides, subject to a cap of 10,000 AUD/MWh.

In terms of the supply stack, coal-fired generators and hydroelectric plants have the lowest marginal cost of production, covering about 90% of the whole NEM capacity.

Supply-side influencing variables According to Christensen et al. (2011), as the temperature shifts may lead to a higher consumption of energy, due to air conditioning demand, temperature is a possible spikes influencing variable. Furthermore, as load represents the contemporaneous demand directly, because of the inelasticity of demand in the EMs (see §1.1.3), it may be regarded



Figure 4.1: The NEM New South Wales spot prices time series from the 1st January 2003 to the 31st December 2010 (140,256 total observations).

as spikes influencing exogenous variable as well.

4.2.1 Data

Focusing on the region of New South Wales (NSW), the data for the estimation process consist of a series of 122,736 half hourly observations of the spot prices and loads, starting from 1st January 2003 up to and including 31st December 2009, while year 2010 is kept for evaluating the forecasts. Although such a time interval is not the same originally considered by Christensen et al. (2011) (i.e. from 1st March 2001 to 30th June 2007 for estimation, with a forecasting horizon of three months), we investigate a more recent time series in order to test the ACH model reliability. In fact, since year 2007 shows a singular spiky behavior, with an extremely high number of observed spikes, we need to know if the ACH model has the capability to fit and forecast spikes within a different time period.

The data exhibit the typical liberalized EM stylized properties (see §1.2). The spot price median is 25.14 AUD/MWh, while the maximum value observed is 10,000 AUD/MWh. The whole TS of the NSW spot prices available can be seen in Fig. 4.1.

Before proceeding, a first inspection on graphs and related ACFs points out that each load period time series does not show a well defined long run



Figure 4.2: The NEM New South Wales spot prices ACF and PACF (load periods 12, 24, 36).

behavior, while Fig. 4.2 remarks a strong weekly periodic component and a persistent autocorrelation function. Furthermore, the ADF test for stochastic trend presence rejects the null hypothesis at 5% significance level for all the load periods.

Spikes issue Although we developed a specific model for spike detection by means of a quantile-based approach (see §2.4), the NSW spikes can be easily identified without any complicated procedure. Therefore, in this case we adopt the same threshold-based method chosen by Christensen et al. (2011). Specifically, the threshold which defines an extreme price event (i.e. a spike) is fixed and set to 100 AUD/MWh, which corresponds approximately to the 98th spot prices percentile.

Exogenous variables The exogenous variables are obtained as follow:

- $T_{\max,t}$ represents the daily absolute deviation of the maximum temperature above its average over the preceding year. Specifically, let $T_{M,t}$ be the maximum temperature observed on the tth day, $T_{\max,t}$ is defined as

$$T_{\max,t} = \begin{cases} 0, & T_{\mathrm{M},t} - \bar{T}_{\mathrm{M},t} \le 0\\ T_{\mathrm{M},t} - \bar{T}_{\mathrm{M},t}, & \text{otherwise} \end{cases},$$
(4.2)

where $\bar{T}_{M,t} = \frac{1}{365} \sum_{i=1}^{365} T_{M,t-i}$.

- $T_{\min,t}$ represents the daily absolute deviation of the minimum temperature below its average over the preceding year. Therefore, similarly to the constructed variable $T_{\max,t}$ (see Eq. 4.2), letting $T_{m,t}$ be the minimum temperature observed on the *t*th day, $T_{\min,t}$ is defined as

$$T_{\min,t} = \begin{cases} 0, & T_{m,t} - \bar{T}_{m,t} \ge 0\\ T_{m,t} - \bar{T}_{m,t}, & \text{otherwise} \end{cases},$$
(4.3)

where $\bar{T}_{m,t} = \frac{1}{365} \sum_{i=1}^{365} T_{m,t-i}$.

- Load_t is constructed by detrending the load at time t, using the mean and the standard deviation of the previous year's worth of data, since the TS of loads exhibits non-stationarity in mean and in variance.

Finally, referring to Eq. 3.12, $\mathbf{z}_t = [1, \text{Load}_t, T_{\max,t}, T_{\min,t}]$ constitutes the set of exogenous variables for the ACH model.

Sources Spot prices and loads were provided by the Australian Energy Market Operator (AEMO), while the daily temperature data were provided by the Australian Government Bureau of Meteorology (BOM). Specifically, the climate statistics historical dataset comes from the Sydney Observatory Hill, which is located in the biggest city in NSW.

4.2.2 Estimating ACH(1,1) parameters

Since both the exogenous variables set and the spike detecting method have been chosen, we should directly proceed with computing the estimates $\tilde{\theta}$. However, we need to face two more issues before starting the procedure.

 The first one is about facing the NAs presence within the climate dataset. Within the temperatures time series, NA probably indicates a temperature value that was not recorded in the dataset available on the BOM website. Since they are not so many (11 missing values over about 2,900 observations) and they do not form any cluster, we replace every single NA temperature value of the tth day with the most recent value available up to the tth day.

2. The second one concerns the load periods, upon which we want to fit the ACH model. As we can see in Fig. 4.3, some load periods do not show a particularly marked spiky behavior. Our choice is to study the more 'interesting' load periods, from the 25th (12:30-13:00) to the 38th (19:00-19:30), which exhibit more than 85 spikes between the year 2003 and 2010 (an average of about 10 spikes per year). Since the whole spot prices TS should be considered for spike forecasting, this is a disputable choice. However, such choice is justified by the need to have a large number of spikes observed to obtain a better calibration of the ACH model. More precisely, the estimation of parameters α₁, β₁, v, b, c, λ is more reliable as the value of N(T) is particularly high (see §3.4).

Anyway, the results of the ACH models applied to the whole set of load periods are summarized in Table A.3.

Finally, we are able to estimate $\tilde{\theta}_j$, $\forall j = 25, \ldots, 38$, where j represents the jth load period.

Estimation results

Table 4.2 shows the ACH(1,1) estimates $\tilde{\theta}_j$ for each of the 14 load periods considered. The presence of NAs is due to the optimization process partial failures of maxLik routine in calculating the Hessian matrix. Such presence indicates that the procedure is particularly stressed for calculating the Hessian matrix for such a high number of parameters. Perhaps, the use of other routines permits the numerical calculations of the entire Hessian matrices, without any presence of NA values.



NSW EM: Half-hourly Prices Boxplot and Spike Counting

Figure 4.3: The NSW spot prices boxplots and spike counts, $\forall j = 1, ..., J$. The top panel shows the log $(P_{j,\cdot})$ boxplots. The bottom panel shows the number of spikes observed within the TS.

Table 4.2: Estimates for the ACH(1,1) model, with the NSW data. The table shows, for each j, the values of the parameter vector $\tilde{\theta}_j$ and the standard errors; significant codes are '* * *' (0.001), '**' (0.01), '*' (0.1). NAs are due to the optimization process partial failures of maxLik routine in calculating the Hessian matrix.

j	j ACH(1,1) Parameter (Variable)									
	α_1	β_1	γ_0	γ_1	γ_2	γ_3	v	b	c	λ
			(const.)	(Load)	(T_{\max})	(T_{\min})				
25	0.394	0.594^{*}	0.542	-1.010	-1.547^{***}	1.627	3.395	1.816	-0.257	0.013
	(0.317)	(0.310)	(2.221)	(0.688)	(0.201)	(1.447)	(2.677)	(1.447)	(0.601)	(0.010)
26	0.0831	0.914***	0.803	-0.947	-1.444***	1.883	2.311	0.814	1.264	0.0879
	(0.056)	(0.051)	(4.437)	(1.277)	(0.322)	(1.412)	(1.673)	(0.633)	(1.795)	(0.074)
27	0.305**	0.689***	1.447	0.114	-1.341	0.461	2.713***	0.750***	0.448***	0.167^{*}
	(0.103)	(0.099)	NA	NA	NA	(0.689)	(0.345)	(0.088)	(0.171)	(0.067)
28	0.122	0.868	3.138	-1.506^{*}	-1.092^{***}	3.280^{*}	4.167	-2.441	0.225	0.078
	NA	NA	(2.214)	(0.746)	(0.182)	(1.47)	NA	NA	NA	NA
29	0.330	0.639	2.390	-1.295	-1.135^{***}	2.823*	1.645***	-2.070	0.438	0.421
	NA	NA	(2.176)	(1.056)	(0.095)	(1.250)	(0.087)	NA	NA	NA
30	0.560	0.427	-1.486	-0.774	-0.984^{**}	2.073^{*}	3.529**	-0.404	-1.897	0.131
	NA	NA	(3.356)	(0.853)	(0.310)	(0.847)	(1.212)	NA	NA	NA
31	0.674	0.299	1.350	-1.989^{***}	-0.979^{***}	2.353^{*}	1.984***	-1.344	0.149	0.231
	NA	NA	(2.275)	(0.351)	(0.140)	(1.115)	(0.187)	NA	NA	NA
32	0.228***	0.756***	1.401^{*}	-1.570^{***}	-0.947^{***}	2.870^{*}	1.965	-2.326^{***}	0.431***	0.197***
	(0.014)	(0.012)	(0.791)	(0.157)	(0.079)	(1.193)	NA	(0.084)	(0.021)	(0.012)
33	0.935	0.057	-0.170	-1.656^{***}	-1.158^{***}	1.767	1.369***	-0.195	-2.481	0.051
	NA	NA	(2.898)	(0.434)	(0.052)	(1.280)	(0.124)	NA	NA	NA
34	0.451	0.542	0.288	-1.603	-1.474	1.825	2.314	-0.811***	-0.219	0.260
	(0.549)	(0.546)	(2.127)	NA	NA	(1.735)	NA	(0.104)	(0.157)	(0.414)
35	0.900	0.076	0.183	-1.192	1.883	-0.356	2.015	-0.314^{***}	-1.863^{***}	3.051
	NA	NA	NA	(0.934)	(1.241)	(0.303)	NA	(0.024)	(0.224)	NA
36	0.384***	0.217	-0.536	-0.867	2.738^{*}	-0.120	1.708	-0.284^{***}	-2.562^{***}	1.883***
	(0.063)	NA	NA	NA	(1.403)	NA	NA	(0.018)	(0.284)	(0.244)
37	0.826***	0.162	0.102	-0.188	2.744**	-0.331	1.544***	-0.092***	-3.436	0.729

(0.126)	(0.149)	NA	(1.079)	(0.864)	(0.291)	(0.044)	(0.004)	(0.133)	NA
$38 0.644^{***}$	0.338*	0.871	-0.572	$5,000^{***}$	-0.299	1.716^{***}	-0.234	-1.879	0.668***
(0.135)	(0.165)	(1.623)	NA	(1.287)	NA	(0.202)	NA	NA	(0.177)

The significance (where standard errors are available) of the ACH model parameters α_1 and β_1 shows that the memory is rather important to model spike occurrences. As we could expect, the coefficients of Load and T_{max} are both tendentially strongly significant for each of the 14 model developed, while the T_{min} coefficient is not.

Specifically, the coefficient of Load assumes negative values, telling us that an upward shift of the demand at time t induces the probability of a spike occurrence in t + 1 to be higher.

On the other hand, the interpretation of the coefficients of the climate variables (i.e., T_{\min} and T_{\max}) is more difficult: the effect of a higher value of maximum temperature (i.e. above the average) is to increase the spike probability, while a lower value of minimum temperature (i.e. below the average) has the effect to decrease such probability. The described situation is observed unless we look at the load periods from 17:30-18:00 to 19:00-19:30, where the influence of the variables is the opposite. It is hard to explain such a behavior, which could be hidden in the office working schedule. In fact, when offices and businesses in general are running, there should be an overuse of air conditioning, which should not be so marked outside the working schedule. Besides, since the constructed variable T_{\min} is related to the minimum temperature observed within the *t*th day, it is expected that its role becomes more critical during the evening hours, when the temperature usually falls down towards the lowest daily temperature value, as well as the role of T_{\max} getting insignificant.

At last, we remark that the NSW temperatures rarely fall below 10°C. As a consequence, a further interpretation of the mostly non-significance of $T_{\rm min}$ coefficient is that the energy demand induced by the minimum temperature changes does not suffer from sharp growth since the minimum temperatures are never particularly harsh.

Still, the significance of the calculated coefficients v and λ , for many of the estimated models, suggests that the augmented form specified in Eq. 3.15 is fairly necessary to model the time intervals between spike occurrences.

The remaining parameters estimated b and c can not be really interpreted, as they are ACH structural shift and rotation parameters which can also not be omitted from the model specification.

Forecast results

Table 4.3 shows the results of spike detection within estimation and forecast by the ACH models, obtained by computing the hazard rate h_t for each of the estimated parameters $\tilde{\theta}_j$, $j = 25, \ldots, 38$ (see §3.20). Fig. 4.4 shows an example of the spike analysis upon the 33rd load period, which compares the observed spikes with the estimated spikes obtained with the value of $h_{33,t}$.

Concerning the forecast results (year 2010), the ACH model predicts 18 spikes over the 76 observed, with a number of 131 false alarms. It means that about 12% of the model predicted spikes are actually real spikes. These results are rather not exciting, and lead us to suppose that both the exogenous variables chosen and the history of the process based on the past and expected durations are not enough to completely explain why a spike occurs. However, we should compare the obtained results with a memoryless model for having a wider view of the ACH power (see $\S4.4.1$).

At the end, we may wonder whether or not an appropriate transformation of the exogenous variables, which can lead the model to produce a better spikes occurrences estimation, exists.

For instance, the combinations of the exponential transformation of Load, $T_{\rm max}$ and $T_{\rm max}$ have been tested². However, such attempts was unsatisfying compared to the previous results.

4.3 The Italian electricity market

The second EM we consider for testing the ACH model based on the load period approach is the Italian electricity market. Since the Italian EM actually

 $^{^{2}}$ We could believe that the effect of the considered exogenous variables fluctuations on the spike occurrence was exponential, rather than linear.

Table 4.3: ACH(1,1) spike occurrences analysis for NSW data. The estimated spikes are split between estimation (2003-2009) and forecast (2010). *Spike detected* indicates the number of spikes detected by the ACH(1,1) model and the number of spikes observed within the TS, while *False alarm* indicates the number of spikes detected which are not observed for real over the total number of estimated spikes.

	Estima	tion	Forecast			
j	Spike detected	False alarm	Spike detected	False alarm		
25	14/85	33/47	3/6	4/7		
26	4/101	21/25	0/5	0/0		
27	3/116	8/11	0/5	0/0		
28	18/124	34/52	2/6	1/3		
29	8/131	13/21	0/5	0/0		
30	72/135	293/365	6/6	57/63		
31	18/135	48/126	1/5	2/3		
32	14/132	46/60	2/5	1/3		
33	26/107	100/126	4/5	66 / 70		
34	1/87	7/8	0/2	0/0		
35	0/89	0/0	0/4	0/0		
36	87/222	69/156	0/11	0/0		
37	91/234	70/161	0/9	0/0		
38	0/160	0/0	0/2	0/0		
Tot.	356/1858	742/1098	18/76	131/149		
	(19.2%)	(67.6%)	(23.7%)	(87.9%)		



Figure 4.4: The NSW ACH(1,1) spike analysis of the 33rd load period. The observed extreme price events are indicated by the grey bar in the top panel. The bottom panel shows the ACH estimated probabilities, with the horizontal grey bar indicating the threshold of 0.5.

differs from the NEM, previously studied in §4.2, in the beginning of this section we explain the main characterizing features of the Italian market, before proceeding by testing our framework.

Before liberalization, the Italian electricity market was under the monopoly of a single vertically integrated and state-owned company (Enel), which basically had full control of the generation, transmission and distribution networks. In March 1999, with the Bersani legislative decree n.79, the Enel privatization process started, the Italian wholesale electricity market was born and the Italian Electricity Market Operator (GME) was instituted.

The Bersani decree implemented a new structure of electricity system, organized as follows:

- The activities related to distribution were subject to license.
- The activities related to transmission and dispatching were subject to the natural monopoly of the Transmission System Operator (TERNA) and the GME.
- The generation, import, export and supply activities were liberalized.

In addition, the liberalization of the demand was gradually introduced from 1999 to the 1st July 2007, when all the electricity consumers became eligible, meaning that they were all able to choose their own supplier (Cariello, 2008).

The Italian wholesale market started to operate as a Pool in April 2004 and became an Exchange in 2005 with the liberalization of the demand-side bidding (Gianfreda and Grossi, 2009).

The Italian Power Exchange (IPEX) spot market is an auction market, where transactions take place the day ahead of the day in which electricity is physically produced and consumed. The IPEX spot market (MGP) operates with hourly load periods (i.e., J = 24). The Unique National Price (PUN) represents the price for end customers, namely the MCP. It is computed as the average of the zonal prices³ weighted by zonal consumption (GME, 2009).

³There are virtual zones with foreign markets which are connected with the Italian electricity network, such as Austria, Corsica, France, Greece, Slovenia and Switzerland, and physical national macro-regions as Northern-Italy, Central-Northern Italy, Central-Southern Italy, Southern Italy and the islands (Sicily and Sardinia).

The main production source of the Italian EM (year 2011 update) is the fossil-fuel, which covers more than 50% of the total demand, followed by the renewable energies (approx. 20%), such as hydro power, geo-thermal power and photovoltaic power plants. A singular feature of Italian EM is that a wide part of the national energy demand is imported from abroad (approx. 13%). Since the green allowances have an important role in the Italian market and the legislation which regulates them is still an evolving process, in terms of supply stack, it is rather complicated to establish the real marginal production cost for each production utility. Tendentially, the energy imported from abroad is the most expensive, while hydro power is the cheapest.

Influencing variables Although the Italian EM shows a different structure and a relatively low degree of liberalization, there is no reason to think that the Australian spot prices supply-side influencing variables do not influence spike occurrences within the PUN time series at all. Therefore, load and temperature may constitute two factors influencing the spike occurrences. Additionally, we consider the unsold load, which represents the quantity of unsold energy present in the market⁴, as another supply-side influencing factor.

At the end, as the majority of the Italian energy production comes from the fossil-fuel sources, extreme values of the spot prices could be influenced by the fluctuations of the BRENT Crude Oil Index for Europe (Calento et al., 2006), which can be classified as a demand-side influencing variable.

4.3.1 Data

The data for the estimation process consist of a series of 52,584 hourly observations of the PUN and load values, starting from the 1st January 2006 to and including 31st December 2010. Similarly to what we did for the NSW market data in §4.2, the whole year 2011 is kept for evaluating the forecasts.

Fig. 4.5 shows the Italian spot prices TS available for our analysis. Although the PUN time series exhibits the typical liberalized EM stylized prop-

⁴In a given load period of a given day, the unsold energy is the quantity of energy which is bought by the GME, but not sold to the customers.
Electricity prices (Italian EM)



Figure 4.5: The Unique National Price (PUN) time series, from the 1st January 2006 to 31st December 2011.

erties, described in §1.2, it looks rather different than the NSW series shown in Fig. 4.1. Mainly, the PUN time series does not have any peak, in which PUN value is as extreme as the spikes values observed in the Australian market, even if it shows extreme and time-varying volatility. Besides, the time series shows a marked trend presence, which changes the mean-value of the spot prices across the time.

We can say that the Italian spot prices TS is in some ways more similar to the British TS analyzed in §2.5 than the Australian one, since the 'normalregime' fluctuations of the spot price are more nervous, as well as the spot prices mean clearly exhibits a trend component.

A first inspection on graphs and related ACFs shows that the PUN time series for each load period does not exhibit a well defined long run behavior. Still, the ACF and PACF functions remark a strong weekly periodic component and a persistent autocorrelation function (see Fig. 4.6). A clear annual dynamic is also shown by the ACFs of the time series associated to the load periods from 9:00 to 21:00. The ADF test rejects the null hypothesis at 5% significance level, except for PUN time series of the 23rd and 24th load periods.

Exogenous variables The exogenous variables T_{max} , T_{min} and Load are constructed by using the same method adopted for the Australian EM (see §4.2.1), while the additional variables chosen are obtained as follows:



Figure 4.6: The Italian EM spot prices ACF and PACF (load periods 6, 12, 18).

- Unsold t is constructed adopting the same procedure used for Load t.
- Brent_t is constructed by detrending the European BRENT Crude Oil closing-prices, using a 2-week moving median and a monthly rolling volatility. Such deseasonalizing procedure for $Brent_t$ is chosen because we believe that the MCP needs no more than one week for having a potential (extreme) response to an oil price fluctuation around the normal regime.

Referring to Eq. 3.12, $\mathbf{z}_t = [1, \text{Load}_t, T_{\max,t}, T_{\min,t}, \text{Unsold}_t, \text{Brent}_t].$

Sources Spot prices and loads time series were provided by ACSM S.p.A., a small company involved in the hydro energy production process. The other EM data were provided by the GME. The daily temperature data were provided by Ilmeteo S.r.l., while the time series of the Europe BRENT Crude Oil daily closing prices were obtained with the support of the Bloomberg Terminal computer system. More precisely, based on the method used for NEM analysis (i.e. choosing the temperatures coming from the biggest city), the temperature data collection concerns the city of Rome, which is the biggest Italian city indeed. Since the Italian weather conditions may radically change depending on the specific area considered, the choice of Rome could be nonsense. However, as we focus on the daily absolute deviations of the minimum and maximum temperatures, which should be almost the same all over the Country, our choice is still reasonable.

4.3.2 Spike detection

Since a price spike identification method which uses a quantile-based approach has been traced in §2.4, for the Italian EM we adopt such framework for detecting spikes. However, since it is a common knowledge⁵ that no extreme spot price values have been observed within the Italian EM during a public holiday, nor during the weekends, we exclude such days from our detrending, deseasonalizing, estimating and forecasting procedures. Specifically, the Italian public holidays omitted are: 1st and 6th January, 25th April, 1st May, 2nd June, 15th August, 1st November and 8th, 25th and 26th December.

4.3.3 Estimating ACH(1,1) parameters

Similarly to what we did for the Australian EM spot prices of NSW, we chose to analyze only the load periods which exhibit a more accentuated spiky behavior. Our choice is fairly justified by the quantile-based approach for spike detection. In fact, since our method classifies from 5% to 10% of the spot price as 'spikes', if we analyzed the time series which do not show enough extreme values, we would occur in a high probability to classify as 'spike' a price which is only slightly above the normal-regime.

As we can see in Fig. 4.7, the more interesting load periods are for $j = 8, 9, 13, \ldots, 23$. These time series show more than the others an interesting number of extreme spot price values, as we can see looking at the boxplots positive outvalues.

Still, coherently to what we did for the Australian analysis in §4.2, the results of the ACH models applied to the whole set of load periods are summarized in Table A.4.

⁵Such knowledge was empirically confirmed by the company ACSM S.p.A.



Figure 4.7: The Italian $p'_{j,t}$ (see Eq.2.5) boxplots, from the 1st of January, 2006 to the 31st of December, 2011, $\forall j = 1, ..., J$.

So, we are now able to estimate $\tilde{\theta}_j$, $j = 8, 9, 13, \ldots, 23$, where j represents the jth load period.

Estimation results

Table 4.4 shows the estimates for the ACH model, for each of the 13 load periods considered.

Unfortunately, the maximization process reports a high number of NAs in estimating the standard errors. It happens because the MaxLik routine may fail in computing some elements of the Hessian matrix, especially as the number of parameters is particularly large. This is the reason why they are omitted from Table 4.4. It follows that any analysis about significance of the coefficients can not be done.

Even so, at least we are able to qualitatively analyze the direction, in which the probability of a spike occurrence moves, for each of the exogenous variables chosen.

Positive fluctuations of the constructed variables Load, T_{max} , T_{min} and Brent at time t bring the hazard of the next day h_{t+1} to be tendentially higher, while Unsold induces an opposite behavior of the hazard rate. Such directions were expected, as they are the same reported in the ACH analysis of NSW prices. Concerning the Brent variable, it is normal that a value above the average of the European BRENT Crude Oil Index leads to a higher risk of a spike occurrence, since the fossil-fuel power stations cover the main part of Italian electricity production.

Forecast results

Table 4.5 reports the results of model estimations and forecasts, obtained by computing the hazard rate h_t for each load period j = 8, 9, 13..., 23 (see §3.20). Fig. 4.8 shows an example of the spike analysis adopting the estimated model for the 9th load period, which compares the observed spikes with the estimated spikes obtained with the values of $h_{9,t}$.

Taking into account that spike occurrence probability is set by the spike detection procedure to around 10% (the on-peak hours method for spikes identification specified in §2.4 was chosen for the interesting load periods), the

j		ACH(1,1) Parameter (Variable)										
	α_1	β_1	γ_0	γ_1	γ_2	γ_3	γ_4	γ_5	v	b	с	λ
			(const.)	(Load)	(T_{\max})	(T_{\min})	(Unsold)	(Brent)				
8	0.036	0.572	1.040	0.811	-0.239	-0.221	2.630	-0.107	2.352	-0.952	-1.436	2.796
9	0.839	0.117	8.094	-0.787	-0.112	-0.186	2.943	0.135	7.658	0.706	0.190	1.631
13	0.114	0.836	2.700	-1.816	-0.354	2.395	-0.453	1.381	2.432	-1.342	-1.784	0.306
14	0.335	0.650	0.772	-1.479	-0.472	0.360	-0.274	-0.684	1.550	-1.701	-0.310	0.058
15	0.033	0.280	3.524	-1.452	-0.016	0.854	0.072	-0.475	2.675	0.560	-0.179	3.235
16	0.053	0.830	0.803	-1.014	-0.574	0.107	0.774	0.759	1.046	0.967	0.829	2.919
17	0.282	0.683	2.766	-1.252	-0.413	1.399	-0.177	-0.401	1.877	1.568	-0.302	3.589
18	0.141	0.625	2.730	-2.036	-0.730	-0.315	-0.125	-0.444	2.592	-1.028	-0.255	3.557
19	0.512	0.455	0.304	-0.841	0.347	0.238	2.307	0.780	4.232	-0.646	-0.991	0.168
20	0.125	0.872	3.225	-1.671	-0.309	-0.004	1.977	-2.344	4.039	-1.593	0.381	0.403
21	0.583	0.416	0.577	1.848	-0.441	-0.650	1.241	-1.335	2.490	-0.906	-0.111	0.920
22	0.312	0.684	1.112	1.042	-0.222	-0.008	2.412	-1.730	1.493	-1.696	-0.850	0.022
23	0.740	0.180	-0.473	-1.990	0.106	0.176	1.570	-1.097	2.555	-1.706	0.301	0.589

Table 4.4: Estimates for the ACH(1,1) model, with the Italian EM data. The table shows, for each j considered, the values of the parameter vector $\tilde{\theta}_j$.

Table 4.5: ACH(1,1) spike occurrence analysis for Italian data. The estimated spikes are split between estimation (2006-2010) and forecast (2011). *Spike detected* indicates the number of spikes detected by the ACH(1,1) model and the number of spikes observed within the TS, while *False alarm* indicates the number of spikes detected which are not observed for real over the total number of estimated spikes.

	Estima	tion	Forec	ast
j	Spike detected	False alarm	Spike detected	False alarm
8	1/120	0/1	0/20	0/0
9	33/126	87/120	11/10	13/24
13	67/121	338/405	16/35	53/69
14	67/119	260/327	19/31	30/49
15	91/128	296/387	17/32	85/102
16	85/128	339/424	19/25	131/150
17	29/128	88/117	0/39	0/0
18	49/113	265/314	15/32	55/70
19	30/107	193/223	14/19	45/59
20	41/112	268/309	8/23	51/59
21	0/124	0/0	0/24	0/0
22	3/126	2/5	0/21	0/0
23	52/126	372/424	17/39	61/78
Tot.	548/1578	2508/3056	136/350	524/660
	(34.7%)	(82.1%)	(38.9%)	(79.4%)

first thing we can say is that the aggregate results of the ACH modelling show a high number of false alarms. Basically, the probability of a false alarm in forecast is about 4 times bigger than the probability that the model captures a real spike. Having a total number of 3,289 hourly observations considered in year 2011, the model classifies 20% of them as a spike occurrence. We believe that it is a rather low-quality result.

4.4 Comparing results: a benchmark Logit model

The results shown in Table 4.2 establish that the rate of spike event occurrences partially depends upon factors relating to load and temperature effects



Figure 4.8: The ACH(1,1) spike analysis of the Italian 9th load period. The observed extreme price events are indicated by the grey bar in the top panel. The bottom panel shows the ACH estimated probabilities, with the horizontal grey bar indicating the threshold of 0.5.

(coefficients of Load, T_{max} and T_{min}), as well as the history of the process (coefficients α_1 and β_1). It means that the ACH model might produce a better spike occurrence forecasting than those produced by any memoryless model which uses the same set of exogenous variables.

Therefore, we need to compare the ACH results with a benchmark model, which forecasts the probability of a spike event by means of the exogenous variables alone. The (memoryless) Logit model

$$p_{t+1} = \frac{1}{1 + \exp\left(-\boldsymbol{\delta}' \mathbf{z}_t\right)} \tag{4.4}$$

provides a straightforward basis for comparative forecast evaluations, where p_{t+1} is the one-step-ahead forecast probability of a spike occurring at time t + 1 and \mathbf{z}_t is a vector of exogenous variables known at time t, similarly to Eq. 3.20 (Christensen et al., 2011).

Numerical aspects Concerning the Logit model in Eq. 4.4, R glm routine, which allows us to estimate a Logit model, is used for an efficient evaluation

	Estima	tion	Forec	ast
Model	Spike detected	False alarm	Spike detected	False alarm
ACH	356/1858	742/1098	18/76	131/149
	(19.2%)	(67.6%)	(23.7%)	(87.9%)
Logit	249/1858	155/404	5/76	7/12
	(13.4%)	(38.4%)	(6.6%)	(58.3%)

Table 4.6: Comparison with Australian data between spike detection capability of the overall ACH(1,1) and Logit models. The estimated spikes are split between estimation (2003-2009) and forecast (2010).

of the parameter δ by means of the log-likelihood maximization, namely $\hat{\delta}$, and the standard errors.

4.4.1 New South Wales EM analysis

Regarding the NEM region of NSW the results from the estimation of the Logit model (see Eq. 4.4) are roughly consistent with those of the ACH model, with the strong significance of the coefficients of Load and T_{max} for $j = 25, \ldots, 34$, and Load and T_{min} for $j = 35, \ldots, 38$. The interpretation of the parameter $\hat{\delta}_j$ leads to the same direction of the ACH estimates interpretation (see §4.2.2 and Table A.1).

Table 4.6 compares the spike occurrences which are detected both in estimation and in forecast by the ACH and Logit models, for the NSW data. As we can see, the aggregate result of ACH modelling shows a better spike detection capability than the Logit model, especially in forecasting, where ACH detects 18 spikes instead of 5, over a number of 76 observed. However, the number of false alarms is higher for the ACH model (873), compared to the Logit model (162), suggesting a higher precision of the Logit model in isolating spike events without generating false alarms. Still, we must also take into account that the ACH false alarms are concentrated in the 30th and 33rd load periods, as we can see in Table 4.3 and partially in Fig. 4.4. These observations suggest both benefits and drawbacks of including the process memory as an hazard rate influencing factor.

Looking at Fig. 4.9, we can see that the Logit model tendentially detects spikes with a higher price value than the ACH model. Specifically, median



Figure 4.9: Price values boxplots of the actual spikes detected by both ACH and Logit models, over the whole TS (i.e., 2003-2010). Values of prices are shown in logarithmic scale.

and 3rd quartile are respectively 181 AUD/MWh and 445.5 AUD/MWh for ACH against 276.2 AUD/MWh and 744.9 AUD/MWh for Logit. This is possibly due to the fact that Logit should catch better than the ACH model the isolated most extreme demand situations upon the electricity network, attributable only to weather and load fluctuations, which induce electricity price to be extremely high. Differently, the ACH model seems to be able to capture also the spikes which do not have an extremely high price value.

At last, comparing the results between the ACH model and the benchmark model, we can say that the history of the process matters, even if the Logit model produces a lower number of false alarms. In particular, we believe that the history of the process allows the Australian ACH model to capture the network features (such as technical aspects) which can not be explained by the chosen supply-side influencing exogenous variables.

4.4.2 Italian EM analysis

The results from the estimation of the Logit model highlight that all the exogenous variables chosen have a rather important role to explain the probability of a spike occurrence. As we can see in Table A.2, it is clear that the exogenous variables have a certain influence in determining when a spike occurs, since

	Estima	Forec	ast	
Model	Spike detected	False alarm	Spike detected	False alarm
ACH	548/1578	2508/3056	136/350	524/660
	(34.7%)	(82.1%)	(38.9%)	(79.4%)
Logit	91/1578	74/165	15/350	1/16
	(5.7%)	(44.8%)	(4.3%)	(6.3%)

Table 4.7: Comparison with Italian data between spike detection capability of the overall ACH(1,1) and Logit models. The estimated spikes are split between estimation (2006-2010) and forecast (2011).

the largest part of the coefficients is significant. Still, we can notice that the role of the variables is the same observed within the NSW market, with a odd role of T_{\min} . Concerning the additional variables considered on the Italian case study Unsold and Brent, they also have a good influence for the largest part of the estimated Logit models. Therefore, although the ACH estimation process did not allow us to analyze the significance of the computed coefficients, the result of the Logit estimation allows us to confirm that the variable chosen are significant.

Table 4.7 compares the result in spikes occurrences prediction between the ACH model and the Logit models. As we can see, similarly to what we found in the Australian EM analysis, the number of actual spike detected is higher for the ACH model than the Logit. By the way, Logit exhibits again an incredibly lower number of false alarms, comparing to the ACH. Basically, ACH forecast predictions show that the probability of a false alarm is four times bigger than the probability of a real spike detection, while such probability is less than 10% for the Logit predictions.

Since the method for spike detection outlined in §2.4, which identifies the relatively extreme price values by adopting a quantile-based approach, was chosen for the Italian EM case study, a deeper qualitative analysis upon which kind of spikes are detected by the ACH and the Logit models, as we did in Fig. 4.9 for the NSW models, is nonsense.

In conclusion, despite the number of effective spikes detected in both models (about 35% by the ACH and 5% by the Logit), which shows that Logit is rather poor in forecasting, the main issue concerns the number of false alarms.

Table 4.8: Comparison between Italian and Australian forecasts (year 2011 for Italy and 2010 for NSW). *Obs.* indicates the total number of observations considered for the forecast evaluation, *Tot* indicates the sum of the real spikes detected and the false alarms. The results over the total number of observations are shown in parenthesis.

So	ırce			Forecast		
$\mathbf{E}\mathbf{M}$	Model	Obs.	Spikes	Real spikes detected	False alarms	Tot.
Italy	ACH	3,289	350	136	524	660
			(10.64%)	(4.13%)	(15.93%)	(20.07%)
	Logit	$3,\!289$	350	15	1	16
			(10.64%)	(0.46%)	(0.03%)	(0.49%)
NSW	ACH	$5,\!110$	76	18	131	149
			(1.49%)	(0.35%)	(2.56%)	(2.92%)
	Logit	$5,\!110$	76	5	7	12
			(1.49%)	(0.10%)	(0.14%)	(0.23%)

Since only one false alarm is observed, the Logit model shows a good effectiveness in forecasting spikes. On the other hand, 524 false alarms are probably too many in order to conclude that the ACH framework is better than the Logit.

4.4.3 Comparison between Italian and Australian forecasts

In order to have a wider view on the results of the developed and estimated models for the Australian and the Italian EMs, in Table 4.8 we propose a comparison based on the spike detection accuracy in forecast between the previously estimated models.

As we can see, the presence of the memory component generally brings to have a larger number of false alarms, which is more evident for the Italian EM. For both of the EMs analyzed, the Logit framework shows a higher accuracy in forecasting real spikes. However, its loss in real spike detection is high compared to the ACH framework results. We believe that on the Australian market the ACH gains something, while on the Italian market it shows a too high propensity to generate false alarms, compared to the total number of available observations. In conclusion, we believe that the memory component may be relevant for modelling price spikes occurrences, depending on the EM considered.

4.4.4 Applying the Logit to the whole NSW spot prices TS

In §4.2 and §4.3 the ACH model has been applied to every single load period. The most interesting results have been reached by applying the ACH framework to the Australian NSW data. The comparisons between the results of ACH and Logit models have been outlined in §4.4.1 and §4.4.2.

Since we introduce this chapter with the ACH results obtained without individually estimating every single load period parameter $\tilde{\theta}_j$, we propose a further comparison between ACH and Logit models by applying them to the whole NSW spot prices TS.

Following the approach adopted by Christensen et al. (2011), we use the same set of exogenous variables which was chosen in §4.2.1 for the NSW analysis – Load_t , $T_{\text{max},t}$ and $T_{\min,t}$. Two spike detection analyses are proposed, based on two different estimation and forecast periods.

- 1. In the first analysis, the dataset is the same of the NEM analysis based on the load period approach (see §4.2). It means that 122,736 half hourly price observations are used for the model estimation and 17,520 for the forecasting evaluation.
- However, as (Christensen et al., 2011) adopt different estimation and forecast periods (i.e. from 1st March 2001 to 30th June 2007 for estimation and from 1st July 2007 to 30th September 2007 for forecasting, using a three-months forecast horizon), the second analysis considers the estimation period from 1st January 2003 to 30th June 2007 and from 1st July 2007 to 30th September 2007 for evaluating the forecasts (see §4.1, which also reports the forecast results summary of this analysis in Table 4.1 Attempt 4).

The choice of analyze two different sets of periods (called for the next part as 'first dataset' and 'second dataset') allows us to have a wider view on the real spike detection capability of the ACH model applied to the entire spot prices time series.

ACH and Logit Estimates

The ACH(1,1) estimates of both datasets are following listed.

- 1. Concerning the first dataset, the ACH estimates are $\tilde{\theta} = [\tilde{\alpha}_1, \tilde{\beta}_1, \tilde{\gamma}_0, \tilde{\gamma}_1, \tilde{\gamma}_2, \tilde{\gamma}_3, \tilde{v}, \tilde{b}, \tilde{c}, \tilde{\lambda}] \simeq [0.053, 0.947, 4.186, -2.228, -1.367, -0.162, 1.582, -0.569, -0.295, 1.564].$
- 2. Concerning the second dataset, the ACH estimates are $\hat{\theta} = [\tilde{\alpha_1}, \tilde{\beta_1}, \tilde{\gamma_0}, \tilde{\gamma_1}, \tilde{\gamma_2}, \tilde{\gamma_3}, \tilde{v}, \tilde{b}, \tilde{c}, \tilde{\lambda}] \simeq [0.074, 0.926, 1.646, -1.889, -1.254, -0.068, 1.473, -0.711, -0.069, 1.364].$

Coherently with the previous analysis of the NEM ACH estimation results (see §4.2.2), the coefficients of the three supply-side influencing variables show that the effect of a positive value of Load_t , $T_{\max,t}$ and $T_{\min,t}$ is to increase the hazard rate h_{t+1} . Since the load period based approach is not adopted anymore, t indicates the tth load period present within the whole spot prices TS.

The Logit estimates are $\hat{\boldsymbol{\delta}} = [\hat{\delta_0}, \hat{\delta_1}, \hat{\delta_2}, \hat{\delta_3}] \simeq [-5.440, 1.028, 0.175, 0.114]$ for the first dataset and $\hat{\boldsymbol{\delta}} \simeq [-5.024, 1.002, 0.101, 0.172]$ for the second one. All the Logit parameters are strongly significant (all the p-values of t-statistic are lower than 0.001). Again, the values of $\hat{\delta_1}, \hat{\delta_2}$ and $\hat{\delta_3}$ confirm that the effect of a positive value of Load_t, $T_{\max,t}$ and $T_{\min,t}$ is to increase the probability of a spike event occurring at time t + 1.

Estimation results Table 4.9 shows the comparison between ACH and Logit modelling results, considering both the first and the second dataset.

The ACH model still exhibits a higher propensity to detect spikes than the Logit model. Concerning the first analysis, the ACH can detect 368 spikes over the 2,777 observed within the estimation time series, despite the Logit which detects only 144 spikes within the same time interval. In the second dataset, the ACH detects 160 spikes over the 1,989 observed within the estimation period, while the Logit detects only 48 spikes.

Table 4.9: Spike detection capability of the ACH(1,1) and Logit on the whole NSW spot prices time series. The estimated spikes are split between estimation (2003-2009) and forecast (2010). Forecast h. indicates the forecast horizon, *Estimation* indicates the estimation period and *F. alarm* indicates the number of false alarms.

			Estimati	Estimation		st
Estimation	Forecast hor.	Model	Spike detected	F. alarm	Spike detected	F. alarm
2003-2009	1 year	ACH	368/2777	67/429	0/93	0/0
			(13.3%)	(15.6%)	(0.0%)	-
2003-2009	1 year	Logit	144/2777	70/214	0/93	62/62
			(5.2%)	(32.7%)	(0.0%)	(100.0%)
2003-2007	3 months	ACH	160/1989	49/209	15/299	22/37
			(8.0%)	(23.4%)	(5.0%)	(59.5%)
2003-2007	3 months	Logit	48/1989	7/55	0/299	0/0
			(2.4%)	(12.7%)	(0.0%)	-

Forecast results (first dataset) Differently than the previous ACH modeling results, the ACH framework exhibits a lower number of false alarms than the Logit model. However, the forecast performances of both the ACH and the Logit models are surprisingly poor. In fact, none of the spikes observed in year 2010 – used for the forecast evaluation – is detected by the two estimated models. It follows that both of the developed models do not have any capability in forecasting spike occurrences.

If we look at the Fig. 4.10, the ACH model apparently fits the spike occurrences within a very short time interval. Such hazard values are totally unexpected comparing to the previously obtained hazard functions. Unless ACH specification is totally inappropriate for the dataset considered, one of the possible reasons why the ACH framework works so badly in modelling spike occurrances with the whole time series is that the anomalous high number of spike events in years 2006 and 2007 (around the 78,000th observation in Fig. 4.10) force the estimation process to be particularly unstable, leading a low-quality calibration and forecasting. A relatively high number of false alarms is observed within the ACH model, where about 60% of the spikes predicted are not observed for real.

Forecast results (second dataset) The estimated ACH(1,1) forecasts 15 spikes over the 299 observed, with 22 false alarms. Hence, the ACH works better with the second dataset than with the first one. The Logit model is not able to forecast any of the spikes observed in the third quarter of year 2007 (see Fig. 4.11). Yet, since the Logit does not predict any spike, it indicates that the exogenous variables are slightly different from the original adopted by Christensen et al. (2011) (in the research paper, the Logit was able to predict 23 spikes with any false alarm), coherently with the analysis outlined in §4.1.

Conclusion

Although Christensen et al. (2011) reports that the ACH model forecasts about 48% of the spikes with a relatively low number of false alarms (i.e. approx. 19% of the predicted spikes), in our analyses the lack of forecasting is clear and induces us to believe that the load period based approach works better if the model described in §3 is considered, as it can be seen in Tables 4.3 and A.3.

However, since the use of conditional expected duration and hazard rate specifications specified by Christensen et al. (2011) generally gives better results (see Table 4.1), our choice of adopting the generalized framework could be contested.

At the end, in general we cannot consider the ACH as a satisfying framework to fit and forecast the spike occurrences in the EMs, since its forecasting capability depends too much upon the time period analyzed and it produces a sizeable number of false alarms.

4.5 Yet another ACH(1,1) estimation

Since a method to detect price spikes was outlined in §2.5 and applied on the APX-PUK data, we propose a further ACH application on the British spot prices time series. Differently than the Australian and Italian cases, studied in §4.2 and §4.3, estimation and analysis are made for each of the 48 available load periods.

The main reason why we are going to analyze the British EM is that it



Figure 4.10: The ACH(1,1) and Logit spike analysis of the whole NSW spot prices time series from 1st January 2003 to 31st December 2010 (forecast horizon of one year). The top panel shows the spot prices time series of the New South Wales. The middle panel shows the ACH estimated hazards, with the horizontal grey bar indicating the threshold of 0.5. The bottom panel shows the Logit estimated probabilities.



Figure 4.11: The ACH(1,1) and Logit spike analysis of the whole NSW spot prices time series, from 1st January 2003 to 30 September 2007 (three-months forecast horizon). The top panel shows the spot prices time series of the New South Wales. The middle panel shows the ACH estimated hazards, with the horizontal grey bar indicating the threshold of 0.5. The bottom panel shows the Logit estimated probabilities.

has been particularly relevant for the latter born European EMs. In fatc, the British electricity privatization, conduced by the government under Margaret Tatcher, has been widely observed as a possible model for EM liberalization and deregulation reforms in a number of countries, especially in Europe. The British EM born in 1990, as a consequence of the British Electricity Act of 1989, which set out dramatic structural changes to the electricity supply industry that came into effect on 31st March 1990 (Green and Newbery, 1992).

As it was mentioned in §2.5, the British EM is an half hourly market where prices are determined in advance, for each level of demand expected during the following day.

The data for estimation process consist of a TS of 83,328 half hourly observations of the spot prices and margin values, starting from the 1st April 2005 up to and including the 31st December 2009. Year 2010 is kept for the forecast validation. The whole spot prices TS is shown in Fig. 2.1, while Fig. 2.2 shows the ACF and PACF functions of the load periods 12, 24 and 36. A first analysis on the main dynamic characteristics of the series (such as seasonality, persistence, spikes, etc.) has been outlined in §2.5.

Differently than Italian and Australian cases, only the exogenous variable Margin is considered. Margin is a variable constructed from the market expectation of the day next load fluctuations, detrended and deseasonalized by means of the smoothing splines and the dummy variables techniques (see §2.4.2 for further details).

The aggregate results of the estimation procedures are shown in Table 4.10. As we can see, the results are particularly poor. In fact, both in the estimation and forecast parts, the estimated models seem to be able to respectively capture 4% and 5.6% of the spike occurrences, with a high number of false alarms.

Table 4.10: Spike detection capability of the overall ACH(1,1) on the British EM data. The estimated spikes are split between estimation (2007-2009) and forecast (2010).

Estimation			Forec	ast
Model	Spike detected	False alarm	Spike detected	False alarm
ACH	237/5905	634/871	70/1241	124/194
	(4.0%)	(72.8%)	(5.6%)	(63.9%)

Chapter 5

Concluding Remarks

5.1 Short summary

In Chapter 1 we introduced the energy market: why it was born and how it works in general. The main common features shared by the world's energy spot prices were then listed $(\S1.2)$.

In Chapter 2 we defined the term *price spike* by investigating the electricity spot prices knowledge. After that, by means of a short literature study, we described the world of methods for detecting price spikes ($\S2.3.2$, $\S2.3.3$), starting from the existing spot price modelling and forecasting awareness ($\S2.3.1$). In the last part of Chapter 2 we proposed a specific method for spike detection ($\S2.4$) and we applied it on the British Energy Market spot prices ($\S2.5$).

In Chapter 3 we explained the ACH framework and the parameter estimation procedure adopted (3.3, §3.4), starting from the description of how the connected standard ACD model works (§3.2). Chapter 4 was introduced with a comparison between our ACH results and the results obtained by Christensen et al. (2011) (§4.1). Later, we applied the ACH model to the Australian and Italian energy markets data (§4.2, §4.3), evaluating the results achieved for each analyzed market. A further analysis of the ACH results had been proposed by the comparison with a memoryless benchmark model (§4.4). In the last part of Chapter 4 an additional estimation attempt, concerning the application of the ACH model on the British EM data, had been outlined (§4.5).

5.2 Significance of the result

First of all, we have to list some remarkable aspects and limitations of our research.

- This document is inspired by the research of Christensen et al. (2011).

However, the ACH modelling procedure developed and adopted was slightly different than that used by the authors. In fact, our ACH framework adopts a generalized specification of the conditional duration, as we can see in Eq. 3.15 and the hazard rate specification originally proposed by Hamilton and Jordá (2002). Additionally, a different perspective was proposed for the largest part of the developed models (the only exceptions are outlined in §4.1 and §4.4.4): Christensen et al. (2011) studied the entire spot prices TS, while we mainly focused on each (interesting) load period, individually taken.

- Since any EM need to be studied for a better approach to its own challenges and characterizing features, the set of exogenous variables, which helps to explain when a spike occurs, could be different for each market subject to analysis.
- Sometimes, numerical optimization may produce unsatisfying results, not necessarily connected to an incorrect specification of the model. As our ACH specification has a large number of parameters, a better estimation of such parameters could be probably reached with more attempts for maximizing the log-likelihood function (e.g., by using different routines for numerical optimization).
- In this document a unique model was proposed for each EM, which was the same for every load period. Anyway, since some coefficients were not significant, the approach that we used constitutes a limitation for spike occurrences modelling. For instance, we could consider a different model (with a different set of explicative variables), for each load period analyzed.
- Concerning the British and Italian EMs, the developed method for spikes identification is an arbitrary method, which apparently brings to having good results, but suffers from some limitations, as we observed in §2.4.4.

First of all, we can say that the ACH modelling does not generally allow to predict spike events with a good accuracy and efficiency. Although Christensen et al. (2011) achieve incredible results in forecasting, our analysis highlights that ACH estimation and forecasting results strongly depend both on the time period and on the specific EM considered. Besides, the number of false alarms is rather high – if compared to the spike prediction capability of the model – to believe that the ACH framework could be considered as a benchmark for the EMs. Therefore, both the EM analyzed (which shows a particular type of spiky behavior as we can quickly realize looking at Fig. 4.2.1) and the time period chosen by the Australian authors, with a very unique situation in year 2007, constitute a special case.

The electricity markets modeled with ACH and Logit show that the exogenous variables chosen, both in supply and in demand sides, have a certain influence in determining spike occurrences. Since the history of the process actually permits to detect more spike events with a number of false alarms lower than in the other EMs, the NSW EM highlights the relevance of the historical component for detecting extreme price events. However, such evidence is not remarkable within the Italian and British EMs, where a rather high number of false alarms does not allow us to have the same conclusion, even if the history permits to detect more spike events as well.

For instance, within the Italian EM analysis (see Table 4.8), knowing that the price occurrence probability is set to 10%, 15.9% of the spot prices within the TS used for forecasting are predicted as spikes by the model and they are not real spikes at all. Such results are even worse within the British EM.

The poor results of the Italian EM analysis are probably due to two facts:

- Italian EM does not have such a high level of liberalization maturity as the Australian and British EMs. As a consequence, its spike occurrences are determined by a wider amount of factors than those considered in our research;
- 2. since the price values are probably too close to the vertically-integrated monopoly prices, Italian spot price fluctuations do not exhibit a remarkable spiky behavior, if compared to the Australian or British prices. Thus, it is rather complicated both to identify a spike and to model spike occurrences.

Both of the previously described aspects may lead to a unstable estimation of the ACH model coefficients.Still, we have to consider that the 2008 World economic crisis drastically changed the energy consumption in Italy (as reported from the GME), probably bringing the market prices to a new balance. Such situation leads to face another challenge for any model calibration process.

Since APX-PUK is one of the most liberalized and deregulated EM all over the World, we do not know the reasons why the results of the ACH modelling on APX-PUK data are so poor, comparing to the NEM results. intuitively, comparing the differences between the spot prices time series of NEM and APX-PUK, the ACH model is not able to fit and forecast spike occurrences within the British spot prices. However, a larger set of exogenous variables should be considered before arriving at any conclusion.

Italian and British EMs likely needs a deeper technical study, both for defining and identifying spike occurrences and for recognizing an appropriate set of factors influencing spike events, unless the ACH framework is totally inappropriate for modelling spike occurrences of such EMs.

5.3 Future work

Firstly, a connection between every single load period model and its neighboring models is auspicious, both for numerically determining the maximum likelihood and for a possible significance of the interrelationships. As a matter of fact, it is largely expected that an observed spike, within a certain load period of a given day, may be relevant in causing other spike events in the neighboring load periods. For instance, a multi-step estimation method could lead to more interesting forecast results, since sometimes a spike spreads within more than one single load period of a given day.

Secondly, an alternative approach for detecting spikes can be used. In such case, a different identification of the spikes occurrences within the time series could bring to rather variant estimation values and consequent results.

Another observation concerns the ACH log-likelihood. We can use a different maximization criterion, based on a weighted relevance of false alarms produced by the model. Alternatively, a deeper study on the estimated hazard rate h_t could be conducted, possibly addressing either to a barrier (set to 0.5 in our research) change or to a single sudden shift of h_t to be taken into account to identify a spike.



Figure 5.1: Analysis of the hazard function h_t barrier (NSW data, ACH model estimated using the entire spot price TS, from 01.01.2003 to 30.06.2007). The black line shows the percentage of the real spikes detected by the model. The grey line shows the percentage of false alarms.

For instance, Fig. 5.1 shows a comparison between the percentage of false alarms and real spikes detected by the model, varying the barrier value of h_t , referring to the second analysis outlined in §4.4.4. As we can see, the percentage of false alarms is constantly higher than the percentage of real spikes detected, leading us to believe that the ACH lack of false alarms is generally persistent.

Moreover, a different set of explicative variables which explains the reasons why a spike occurs in a better way, can be found by a careful analysis of the considered EM, as well as a different order of the parameters m and q in Eq. 3.13.

Eventually, since the hardware sources have an important role in the numerical maximization of the log-likelihood, especially when a model has a wide number of parameters, we believe that the use of the cloud computing could help to make a higher number of attempts, which could naturally lead to better results.

Appendix A

Definitions, Models and Formulas

A.1 Definitions

Definition A.1 (Point Process). A (univariate) point process (PP) is a sequence $\mathbf{\Phi} = (T_n)_{n\geq 1}$ of positive random numbers T_n , which may also take the value $+\infty$. We may interpret T_n as the time at which a certain (random) event occurs the *n*th time and assume that (everywhere on Ω):

$$T_n < T_{n+1}$$
, if $T_n < \infty$,
 $T_n = T_{n+1}$, if $T_n = \infty$.

(Last and Brandt, 1995).

Definition A.2 (Marked Point Process). Assume that $(\mathbf{X}, \mathcal{X})$ is a measurable space, define $\mathbf{X}_{\infty} := \mathbf{X} \cup \{x_{\infty}\}$ and let \mathcal{X}_{∞} be the σ -field of subsets of \mathbf{X}_{∞} , which is generated by \mathcal{X} and $\{x_{\infty}\}$. We define a marked point process (MPP) as a sequence $\mathbf{\Phi} = ((T_n, X_n))_{n \geq 1}$ of pairs of random elements T_n of $(0, \infty]$ and X_n of \mathbf{X}_{∞} that meets the point process conditions mentioned in A.1 (Last and Brandt, 1995).

Definition A.3 (Survival Function). Let X be a random variable with cumulative distribution function $F_X(t)$ on the interval $[0, +\infty)$ and probability function $f_X(t)$. Its survival function is defined as

$$R(t) \equiv \mathbb{P}[X > t] = \int_t^\infty f_X(u) \, du = 1 - F_X(t).$$

Definition A.4 (Box-Cox Transformation). The Box-Cox transformation (without shift parameter) is defined as a continuously varying function, with respect to the parameter v, in a piece-wise function form that makes it continuous at the point of singularity (v = 0). Given a data vectors $y = [y_1, \ldots, y_n]$, with $y_i > 0$, the transformation is

$$y_i = \begin{cases} \frac{y_i^v - 1}{v[(GM(y)]^{v-1})}, & \text{if } v \neq 0\\ GM(y) \ln y_i, & \text{if } v = 0 \end{cases}$$

,

where GM(y) is the geometric mean of the observations y_1, \ldots, y_n .

A.2 Models

A.2.1 Seasonal ARIMA model

A time series $\{X_t\}$ is a SARMA $(p, d, q) \times (P, D, Q)_s$ process with period s if it satisfies a difference equation of the form

$$\phi(B)\Phi(B^{s})(1-B)^{d}(1-B^{s})^{D}X_{t} = \theta(B)\Theta(B^{s})Z_{t}, \ Z_{t} \sim \mathcal{N}(0,\sigma^{2}),$$

where p, d, q, P, D and Q are nonnegative integers; $\phi(z) = 1 - \sum_{i=1}^{p} \phi_i z^i$, $\Phi(z) = 1 - \sum_{i=1}^{p} \Phi_i z^i$, $\theta(z) = 1 + \sum_{j=1}^{q} \theta_j z^j$ and $\Theta(z) = 1 + \sum_{j=1}^{q} \Theta_j z^j$; B is the backward shift operator (i.e., $B^j X_t = X_{t-j}$) and Z_t is the error term. The parameters ϕ_1, \ldots, ϕ_p are the AR coefficients, Φ_1, \ldots, Φ_p are the seasonal AR coefficients, $\theta_1, \ldots, \theta_q$ are the MA coefficients and $\Theta_1, \ldots, \Theta_q$ are the seasonal MA coefficients. d and D are the degrees of differencing required to achieve stationarity (Lai et al., 2001).

A.2.2 GARCH model

A time series $\{Y_t\}$ is a (strong) generalized autoregressive conditional heteroscedasticity GARCH(m, q) process with ARCH order m and GARCH order q if it satisfies the equation

$$Y_t = \sigma_t X_t, \ X_t \sim IID(0,1),$$

where σ_t is the function

$$\sigma_t^2 = \omega_0 + \sum_{i=1}^m \alpha_i Y_{t-i}^2 + \sum_{i=1}^q \beta_i \sigma_{t-i}^2,$$

with $\omega > 0$ and $\alpha_i, \beta_j \ge 0, \forall i = 1, \dots, m; j = 1, \dots, q$ (Lisi, 2010).

A.3 Formulas

ARMA representation of the standard ACD(m,q) model

Given the ACD(m, q) model §3.2.1, such that

$$\psi_{N(t)} = \omega + \sum_{j=1}^{m} \alpha_j x_{N(t)-j} + \sum_{k=1}^{q} \beta_k \psi_{N(t)-k},$$

a very useful feature is that it can be formulated as an ARMA(r, q) model for durations x_i .

$$x_{N(t)} = \omega + \sum_{j=1}^{r} (\alpha_j + \beta_j) x_{N(t)-j} - \sum_{j=1}^{q} \beta_j w_{N(t)-j} + w_{N(t)}; \quad x_{N(t)} = \psi_{N(t)} + w_{N(t)},$$

where $r = \max\{m, q\}$ and w_i is a martingale difference.

A.4 Logit estimates

Table A.1: Estimates for the Logit model, with the NSW data. The table shows, for each j, the values of the parameter vector $\hat{\delta}_j$ and the standard errors; significant codes are '* * *' (0.001), '**' (0.01), '*' (0.1).

j	Logit	Paramet	ter (Varia	ble)
	δ_0	δ_1	δ_2	δ_3
	(interc.)	(Load)	(T_{\max})	(T_{\min})
25	-4.930^{***}	0.830***	0.259***	0.142**
	(0.259)	(0.111)	(0.033)	(0.055)
26	-4.424^{***}	0.767***	0.235***	0.033
	(0.221)	(0.101)	(0.030)	(0.058)
27	-4.344^{***}	0.830***	0.245***	0.007
	(0.215)	(0.098)	(0.060)	(0.030)
28	-4.235^{***}	0.746***	0.259***	-0.002
	(0.206)	(0.094)	(0.029)	(0.060)
29	-4.146^{***}	0.789***	0.246***	-0.017
	(0.200)	(0.093)	(0.029)	(0.060)
30	-4.284^{***}	0.804***	0.274^{***}	0.012
	(0.208)	(0.093)	(0.030)	(0.060)
31	-4.390^{***}	0.893***	0.271***	0.029

	(0.215)	(0.095)	(0.030)	(0.060)
32	-4.478^{***}	0.958***	0.266***	0.048
	(0.221)	(0.097)	(0.031)	(0.057)
33	-4.597^{***}	0.880***	0.254^{***}	0.064
	(0.231)	(0.102)	(0.032)	(0.055)
34	-4.981^{***}	0.875***	0.277***	0.120^{*}
	(0.263)	(0.114)	(0.034)	(0.054)
35	-4.659^{***}	0.746***	0.191***	0.199***
	(0.244)	(0.125)	(0.034)	(0.050)
36	-3.156^{***}	0.649***	0.016	0.175***
	(0.154)	(0.038)	(0.035)	(0.038)
37	-3.276^{***}	0.406***	0.008	0.268***
	(0.161)	(0.106)	(0.039)	(0.0410)
38	-4.096^{***}	0.523***	0.041	0.315***
	(0.215)	(0.131)	(0.047)	(0.051)

Table A.2: Estimates for the Logit model, with the Italian EM data. The table shows, for each j, the values of the parameter vector $\hat{\delta}_j$ and the standard errors; significant codes are '* * *' (0.001), '**' (0.01), '*' (0.1).

j	Logit Parameter (Variable)							
	δ_0	δ_1	δ_2	δ_3	δ_4	δ_5		
	(interc.)	(Load)	(T_{\max})	(T_{\min})	(Unsold)	(Brent)		
8	-2.364^{***}	-0.001	0.022	0.014	-0.557^{***}	-0.053		
	(0.185)	(0.104)	(0.025)	(0.032)	(0.090)	(0.090)		
9	-2.410^{***}	0.042	0.040	0.032	-0.754^{***}	-0.039		
	(0.186)	(0.115)	(0.025)	(0.032)	(0.100)	(0.090)		
13	-2.961^{***}	0.629***	0.105***	-0.001	-0.185^{*}	0.327***		
	(0.225)	(0.124)	(0.026)	(0.039)	(0.101)	(0.098)		
14	-3.198^{***}	0.596^{***}	0.138***	0.026	-0.220^{*}	0.248^{*}		
	(0.239)	(0.123)	(0.027)	(0.040)	(0.105)	(0.100)		

15	-3.490^{***}	0.574^{***}	0.202***	0.045	-0.197^{*}	0.267^{***}
	(0.255)	(0.119)	(0.027)	(0.042)	(0.107)	(0.101)
16	-3.160^{***}	0.518***	0.177^{***}	-0.011	-0.192^{*}	0.132^{*}
	(0.240)	(0.116)	(0.026)	(0.044)	(0.105)	(0.098)
17	-3.048^{***}	0.513***	0.143^{***}	0.023	-0.378^{***}	0.224^{*}
	(0.228)	(0.120)	(0.025)	(0.038)	(0.102)	(0.096)
18	-3.260^{***}	0.412**	0.144^{***}	0.066^{*}	-0.422^{***}	-0.057
	(0.231)	(0.126)	(0.025)	(0.037)	(0.104)	(0.097)
19	-2.527^{***}	0.182	0.0282	-0.025	-0.761^{***}	0.070
	(0.200)	(0.138)	(0.026)	(0.036)	(0.108)	(0.099)
20	-2.663^{***}	0.231^{*}	0.068**	0.004	-0.415^{***}	0.394***
	(0.202)	(0.126)	(0.025)	(0.035)	(0.101)	(0.096)
21	-3.552^{***}	-0.305^{***}	0.162^{***}	0.182***	-0.341^{***}	0.294**
	(0.237)	(0.084)	(0.025)	(0.032)	(0.088)	(0.091)
22	-2.828^{***}	-0.161^{*}	0.066**	0.094**	-0.275^{***}	0.420***
	(0.205)	(0.084)	(0.024)	(0.031)	(0.087)	(0.091)
23	-2.913^{***}	0.549***	0.062^{*}	0.068^{*}	-0.425^{***}	0.343***
	(0.214)	(0.112)	(0.026)	(0.033)	(0.087)	(0.095)

A.5 Overall results: Australian and Italian cases

Since in §4.2 and §4.3 we focused on the ACH estimation upon the load periods which exhibit – from our viewpont – a more 'interesting' spiky behavior, we report here the results of ACH estimation on overall the load periods for both NEM and Italian EM data.

	Estima	tion	Forecast		
Load periods	Spike detected	False alarm	Spike detected	False alarm	
25-38	356/1858	742/1098	18/76	131/149	
	(19.2%)	(67.6%)	(23.7%)	(87.9%)	
1-24, 39-48	185/874	571/756	3/17	121/124	
	(21.2%)	(75.5%)	(17.6%)	(97.6%)	
Tot.	541/2732	1313/1854	21/93	252/273	
	(19.8%)	(70.8%)	(22.6%)	(92.3%)	

Table A.3: Summary of the ACH(1,1) spike detection capability for each of the 48 load periods (Australian data). The estimated spikes are split between estimation (2003-2009) and forecast (2010).

Table A.4: Summary of the ACH(1,1) spike detection capability for each of the 24 load periods (Italian data). The estimated spikes are split between estimation (2006-2010) and forecast (2011).

	Estimation		Forecast	
Load periods	Spike detected	False alarm	Spike detected	False alarm
8, 9, 13 - 23	548/1578	2508/3056	136/350	524/660
	(34.7%)	(82.1%)	(38.9%)	(79.4%)
$1\!-\!7,10\!-\!12,24$	211/715	1388/1599	40/152	147/187
	(29.5%)	(86.8%)	(26.3%)	(78.6%)
Tot.	759/2293	3896/4655	176/502	671/847
	(33.1%)	(83.7%)	(35.1%)	(79.2%)

Appendix B

R code

B.1 ACH likelihood and estimation functions

```
### LIKELIHOOD FUNCTIONS : LOG LIKELIHOOD FUNCTIONS FOR ACH ###
3
  ***********
4
6
  # function: u_series - computes historical durations
7
  # input: prices (vec num), spikes (vec boolean), add_first (boolean)
8
  # output: x (vec num) - historical durations vector u_{N(t)}
9
  10
11
  u_series <- function(prices, spikes, add_first=TRUE) {</pre>
12
   t<-which(spikes==TRUE)
13
    # input check
    if (length(prices)!=length(spikes)) stop("u_series: input lengths
14
       differ")
15
    if (!is.logical(spikes)) stop("u_series: vector 'spikes': wrong input
       format")
    if (length(t)==0) stop("u_series (warning): no spikes presence")
16
17
18
    spikes<-as.logical(as.vector(spikes))</pre>
19
    x<-rep(NA, length(t)-1)</pre>
    for (i in 2:length(t))
20
     x[i-1]<-t[i]-t[i-1]
21
22
    if(add_first==TRUE)
23
     return(c(mean(x),x))
24
    else
25
     return(x)
26 }
27
28
29 ********
30 # function: psi_model - computes psi function values
31 # input: prices (vec num), spikes (vec boolean), m (int), q (int), Alpha (vec num), Beta (vec num), v (num), b (num), c (num), lambda (num)
32 # output: psi_v (vec num) - psi function values psi_1, ..., psi_N_t; t=N
      _t=1,...length(which(spikes==TRUE))
33
  34
35 psi_model <- function (prices, spikes, m, q, Alpha, Beta, v, b, c, lambda)
      {
36
    AACD_abs_par<-1e-4
    psi<-NULL
37
38
    # input check
    if (v<=0) stop("psi_model: 'v' must be positive")</pre>
39
40
    x<-u_series(prices, spikes, add_first=TRUE)</pre>
41
    N \leq -length(x)
42
    # initialization
    init_psi<-function(x) {</pre>
43
     psi<-rep(sum(Alpha)*x[1]/(1-sum(Beta)),length(x))</pre>
44
45
     return(psi)
46
    }
47
    shock<-function(x, psi, r, index, b, c, lambda, gpsi) {</pre>
```

```
48
       temp<-x[(index-1):(index-r)]/(gpsi * psi[(index-1):(index-r)]) - b #</pre>
           b: shift parameter, c: rotation parameter
 49
       #(abs(temp) - c*temp)^lambda
50
       return((sqrt(AACD_abs_par^2+temp^2) - c*temp)^lambda)
 51
     }
52
     if (is.null(psi))
53
       psi<-(init_psi(x))^v</pre>
54
     else
       psi<-psi^v
 55
 56
57
     for(i in (max(m,q)+1):N) {
58
       temp<-t(Alpha) %*% (psi[(i-1):(i-q)] * shock(x,psi,r=q,index=i,b=b,c</pre>
           =c,lambda=lambda,gpsi=1)) + t(Beta) %*% psi[(i-1):(i-m)] # gpsi
           :=1 <=> EAACD
 59
       psi[i]<-temp
60
     7
61
     psi[1] <-mean(psi[-1])</pre>
62
     return(psi^(1/v))
 63 }
 64
65 *****
66 # function: lambda_function - computes lambda function values
67
   # input: x (vec num), delta (num)
 68 # output: out (vec num) - lambda function values
 69
   *****
 70
 71 lambda_function<-function(x, delta=0.1) {
 72
     x<-as.vector(x)
73
     lambda<-x
 74
     # input check
 75
     if (!is.numeric(x)) stop("lambda_function: 'x' must be numeric")
 76
     lambda[which(x \le 1)] \le -1.0001
 77
     lambda[which((x>1)&&(x<(1+delta)))]<-1.0001+((2*delta*(x[which((x>1)&&
         (x \le (1 + delta))) = -1)^2)/(delta^2 + (x[which((x>1)\&\&(x \le (1 + delta))))))
         ]-1)^2))
 78
     lambda[which(x>=1+delta)] < -0.0001+x[which(x>=1+delta)]
 79
     return(lambda)
80 }
 81
 82
   83 # function: h_model - autoregressive conditional hazard model
84 # input: prices (vec num), spikes (vec boolean), m (int), q (int), Alpha
        (vec num), Beta (vec num), Gamma (vec num), v (num), Z (matr num),
       delta (num)
 85
   # output: h (vec num) - conditional hazard values
86 ******
87
88 h_model <- function (prices, spikes, m, q, Alpha, Beta, Gamma, v, b, c,
       lambda, Z, delta=0.1) {
 89
     T<-length(prices)
90
     psi out <- rep(NA.T)
91
     Alpha <- as.vector(Alpha)
92
     Beta <- as.vector(Beta)
93
     Gamma<-as.vector(Gamma)
 94
     v<-as.numeric(v)
95
     Z < -as.matrix(Z)
96
     # input check
97
     if ( (any(Alpha<0)) || (any(Beta<0)) || (v<=0) || (length(Alpha)!=m) || (
         length(Beta)!=q) ) stop("h_model: wrong inputs")
     if ( (any(is.na(Z))) || (length(Z[1,])!=length(Gamma)) ) stop("h_model:
98
          wrong 'Z' or 'Gamma' format")
     if ( length(Z[,1])!=T ) stop("'h_model (warning): wrong number of
        observations in 'Z'")
99
100
     x<-u_series(prices,spikes)</pre>
101
     psi<-psi_model(prices, spikes, m, q, Alpha, Beta, v, b, c, lambda)</pre>
102
     psi<-rev(psi)
103
     s_time<-rev(which(spikes==TRUE))[1:length(psi)]</pre>
104
     s_time<-c(length(prices)+1,s_time)</pre>
105
     for (i in 1:length(psi))
106
       psi_out[(s_time[i+1]+1):s_time[i]]<-psi[i]</pre>
```

```
107
     start<-s time[length(s time)]+1</pre>
     h < -1/lambda_function( psi_out[start:T] + (as.vector(Z%*%(Gamma))[start])
108
         :T]) )
109
     return(list(h=h, start=start))
110 }
111
112
   113 # function: ln_likelihood_EXP_ACH - ln likelihood of ACH model with EXP
       errors
114 # input: prices (vec num), spikes (vec boolean), garch.order (int), arch
       .order (int), Alpha (vec num), Beta (vec num), Gamma (vec num), v (
       num), Z (matr num)
115 # output: LnLik (num) - ln likelihood value
117
118 In_likelihood_EXP_ACH<-function(prices, spikes, garch.order, arch.order,
        Alpha, Beta, Gamma, v, b, c, lambda, Z) {
119
     # defining U and c st U %*% theta - c <= 0, theta:=[Alpha, Beta, Gamma</pre>
         , v, b, c, lambda]
120
     U<-diag(arch.order+garch.order+length(Z[1,])+1)
121
     c<-rep(0, arch.order+garch.order+length(Z[1,])+1)
122
     \texttt{U} < -\texttt{U} \left[ -((\texttt{arch.order+garch.order+1}):(\texttt{arch.order+garch.order+length}(\texttt{Z} + \texttt{Z})) \right] 
         [1,]))),]
123
     c<-c[-((arch.order+garch.order+1):(arch.order+garch.order+length(Z
         [1,])))]
124
     if(garch.order>0) {
125
       U<-rbind( U,c(rep(0,arch.order),rep(1,garch.order),rep(0,length(Z
           [1,])),0),c(rep(0,arch.order),rep(-1,garch.order),rep(0,length(Z
           [1,])),0) )
126
       c<-c(c,0,-1)
127
128
     U<-rbind( U,c(rep(-1,arch.order),rep(-1,garch.order),rep(0,length(Z
         [1,])),0) ) # stationarity constr
129
     c<-c(c,-1) # stationarity constr</pre>
130
     T<-length(prices)
131
     v<-rep(NA, T)
132
     LnLik<-0
133
     y[which(spikes==TRUE)]<-1
     y[which(spikes==FALSE)]<-0
134
135
      # input and constraints check
     if (any(is.na(y))) stop("ln_likelihood_EXP_ACH: 'y' not computed")
136
137
138
     if ( (any((U%*%as.vector(c(Alpha, Beta, Gamma, v))-c) <=0)) ) {
       LnLik<-rep(NA,length(h_model(prices, spikes, arch.order, garch.order</pre>
139
           , rep(0,length(Alpha)), rep(0,length(Beta)), rep(0,length(Z[1,])
           ), v=1, b=0, c=0, lambda=1, Z)$h))
140
     7
141
     else {
       h_lik<-h_model(prices, spikes, m=arch.order, q=garch.order, Alpha,
142
          Beta, Gamma, v, b, c, lambda, Z)
143
       v lik<-v[h lik$start:T]
144
       if (length(h_lik$h)!=length(y_lik)) stop("ln_likelihood_EXP_ACH: 'h'
            and 'y' lengths differ")
145
       LnLik <-(y_lik*log(h_lik$h))+((1-y_lik)*log(1-h_lik$h))
146
     }
147
     return(list(value=sum(LnLik), lnl=LnLik))
148 }
149
151
   # function: Optim_LnL_EXP_ACH - numerical optimization
152 # input: prices (vec num), spikes (vec boolean), garch.order (int), arch
       .order (int), Alpha (vec num), Beta (vec num), Gamma (vec num), v (
       num), Z (matr num), init (vec num), stationary.constr (boolean),
       maxit (num)
153 # output: opt (obj) - optim function output
155
156 Optim_LnL_EXP_ACH<-function(prices, spikes, arch.order, garch.order, Z,
       init=NA, grad=FALSE, maxit=100000, method=c("BFGS", "BFGSR", "SANN", "
       NM")) {
```

```
157
      require(maxLik)
158
      require(sandwich)
159
       if (is.na(init))
160
        init<-c(rep(0.4, arch.order+garch.order), rep(0, length(Z[1,]))</pre>
              ,1,0,0,1)
161
      if ( (grad==TRUE)&&(arch.order==1)&&(garch.order==1) ) {
      、 (grac
res<-NA
}
162
163
164
      else {
165
         logLik_function<-function(theta) {</pre>
166
           if (arch.order>0)
167
             Alpha<-theta[1:arch.order]
168
           if (garch.order>0)
169
             Beta<-theta[(arch.order+1):(garch.order+arch.order)]
170
           if (length(Z[1,]>0))
171
             Gamma <- theta [(arch.order+garch.order+1):(garch.order+arch.order+
                  length(Z[1,]))]
172
           v<-theta[length(theta)-3]
173
           b<-theta[length(theta)-2]
174
           c<-theta[length(theta)-1]
175
           lambda <- theta [length (theta)]
           out<-ln_likelihood_EXP_ACH(prices, spikes, garch.order, arch.order
, Alpha, Beta, Gamma, v, b, c, lambda, Z)
176
177
           return(out$lnl)
178
         }
179
         res<-maxLik(logLik=logLik_function, grad=NULL, hess=NULL, start=init
   , method=method, print.level=0, iterlim=maxit)</pre>
180
      }
181
       return(res)
182 }
```

```
R code
```
Bibliography

- T.M. Christensen, A.S. Hurn, and K.A. Lindsay. Forecasting spikes in electricity prices. *International Journal of Forecasting*, 2011.
- R. Weron. Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach. Wiley, New York City, U.S., 2006.
- D. Edwards. Energy Trading and Investing: Trading, Risk Management and Structuring Deals in the Energy Market. McGraw-Hill, New York City, U.S., 2009.
- William W. Hogan. Regulation and Electricity Markets. In Western Power Trading Forum, Cambridge, Massachussets, 2009. Harvard University.
- Javier Contreras, Oscar Candiles, José Ignacio de la Fuente, and Tomás Gómez. Auction Design in Day-Ahead Electricity Markets (Republished). IEEE Transactions on Power Systems, 16(3):409–417, 2001.
- Afzal S. Siddiqui. Price-Elastic Demand in Deregulated Electricity Markets. In Institute for Operations Research and the Management Sciences (INFORMS). Environmental Energy Technologies Division Ernest Orlando Lawrence, Berkeley National Laboratory, California, November 2003.
- Hélyette Geman and Andrea Roncoroni. Understanding the fine structure of electricity prices. *Journal of Business*, 79:1225–1261, 2006.
- Jon Exley, Shyam Mehta, and Andrew Smith. Mean Reversion (Draft 21st June 2004). In *Finance and Investments Conference*, Brussels, Belgium, 2004.
- Shijie Deng. Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes, February 2000.
- Carlos Blanco and David Soronow. Mean Reverting Processes Energy Price Processes Used For Derivatives Pricing & Risk Management. *Commodities* Now, pages 68–72, 2001.

- Severin Borenstein. The trouble with electricity markets (and some solutions), January 2001.
- Claudia Klüppelberg, Thilo Meyer-Brandis, and Andrea Schmidt. Electricity spot price modelling with a view towards extreme spike risk. *Quantitative Finance*, 10(9):963–974, 2010.
- Takashi Kanamura and Kazuhiko Ohashi. A structural model for electricity prices with spikes: Measurement of spike risk and optimal policies for hydropower plant operation. *Energy Economics*, 29(5):1010–1032, September 2007.
- Francesco Lisi and Fany Nan. Components estimation for electricity prices: procedures and comparisons. Padova, Italy, 2012.
- Joanna Janczura and Rafal Weron. Modeling electricity spot prices: Regime switching models with price-capped spike distributions. MPRA Paper 23296, University Library of Munich, Germany, June 2010.
- Gero Schindlmayr. A regime-switching model for electricity spot prices. Scientific paper, EnBW Trading GmbH, May 2005.
- Sanjeev Kumar Aggarwal, Lalit Mohan Saini, and Ashwani Kumar. Electricity price forecasting in deregulated markets: A review and evaluation. International Journal of Electrical Power & Energy Systems, 31(1):13–22, January 2009. ISSN 01420615. doi: 10.1016/j.ijepes.2008.09.003.
- Michael Bierbrauer, Stefan Trueck, and Rafal Weron. Modeling electricity prices with regime switching models. Econometrics 0502005, EconWPA, February 2005.
- Alicia Troncoso Lora, Jesús Riquelme Santos, José Cristóbal Riquelme Santos, Antonio Gómez Expósito, and José Luís Martínez Ramos. A comparison of two techniques for next-day electricity price forecasting. In *Proceedings of* the Third International Conference on Intelligent Data Engineering and Automated Learning, IDEAL '02, pages 384–390, London, UK, 2002. Springer-Verlag. ISBN 3-540-44025-9.

- M. Zhou, Z. Yan, Y. X. Ni, G. Li, and Y. Nie. Electricity price forecasting with confidence-interval estimation through an extended ARIMA approach. In *Generation, Transmission and Distribution, IEE Proceedings-*, volume 153, pages 187–195. IET, 2006.
- Stefan Trueck, Rafal Weron, and Rodney Wolff. Outlier treatment and robust approaches for modeling electricity spot prices. MPRA Paper 4711, University Library of Munich, Germany, August 2007.
- A. Azzalini and B. Scarpa. Analisi dei dati e data mining. UNITEXT / Collana di Statistica e Probabilità Applicata. Springer, 2004. ISBN 9788847002722.
- Ruey S. Tsay. Analysis of Financial Time Series. Wiley Series in Probability and Statistics. John Wiley & Sons, Hoboken, New Jersey, 3rd ed. edition, 2010. ISBN 9780470414354.
- Francesco Lisi. Modelli e Metodi per Serie Storiche Finanziarie. Universitá degli Studi di Padova, Italy, 2010.
- Robert F. Engle and J.R. Russell. Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data. *Econometrica*, 66: 1127–1162, 1998.
- Maria Pacurar. Autoregressive Conditional Duration (ACD) Models in Finance: A Survey of the Theoretical and Empirical Literature. SSRN eLibrary, 2006. doi: 10.2139/ssrn.933120.
- Guillaume Weisang. Acd models: Models for data irregularly spaced in time. Bentley College, Massachusetts, April 2008.
- James D. Hamilton and Oscar Jordá. A Model for the Federal Funds Rate Target. Journal of Political Economy, 110(5):1135–1167, October 2002. University of Chicago Press.
- M. Fernandes and J. Grammig. A family of autorefressive conditional duration models. *Journal of Econometrics*, 130:1–23, 2006.
- Yiu-Kuen Tse and Yang Tao. Estimation of High-Frequency Volatility: Autoregressive Conditional Duration Models versus Realized Volatility Methods. School of Economics, Singapore Management University, 2010.

- Marco R. Steenbergen. Maximum Likelihood Programming in R. University of North Carolina, Chapel Hill, 2006.
- Francesco Cariello. The Italian Electricity Market. Technical report, Autoritá per l'energia elettrica e il gas, October 2008.
- Angelica Gianfreda and Luigi Grossi. Zonal Price Analysis of the Italian Wholesale Electricity Market. In Energy Market, 2009. EEM 2009. 6th International Conference on the European, pages 1–6. University of Verona, Italy, May 2009.
- GME. Vademecum della borsa elettrica italiana. Gestore Mercati Energetici, January 2009.
- Rita Calento, Raffaella Centurelli, and Marco Tomasi. Analisi econometrica del Prezzo di borsa dell'Energia Elettrica. Technical report, Scuola Mattei Master Medea - ENI, May 2006.
- Richard J. Green and David M. Newbery. Competition in the British Electricity Spot Market. *Journal of Political Economy*, 100(5):929–953, October 1992. University of Chicago Press.
- Günter Last and Andreas Brandt. Marked Point Processes on the Real Line: The Dynamical Approach. Probability and Its Applications. Springer, Berlin, Germany, August 1995.
- E.K.M. Lai, Y.W. Cheng, and M. McAleer. Predicting Monthly Catch for Some Western Australian Coastal Finfish Species with Seasonal ARIMA -GARCH Models. In *MODSIM01*, pages 1487–1492. University of Western Australia, December 2001.