NTNU
Norwegian University of
Science and Technology

Time Series Analysis of Electricity Prices

A comparative study of power markets

Erik Løkken Walter

Master of Science in Physics and Mathematics
Submission date: July 2011
Supervisor: Ingve Simonsen, IFY

Norwegian University of Science and Technology
Department of Physics

NTNU
Faculty of Natural Sciences
and Technology
Department of Physics

Master's thesis

Time Series Analysis of Electricity Prices

- A comparative study of power markets

Erik Løkken Walter
Supervisor: Ingve Simonsen

July 2011

Abstract

During the past few decades, the power sectors of several countries have been substantially reorganized, and liberalized markets for trading of electricity have been established. In this report, evidence from six electricity markets are studied in order to identify characteristics of electricity prices. The market structures, statistical quantities, as well as long-term dependence, are investigated. Detrended fluctuation analysis and the average wavelet coefficient method are employed in order to estimate the Hurst exponent, which quantifies the presence of long-termed dependence. Since it is concluded that the price series are periodic on several time scales, all characteristics are investigated for both the original and deseasonalised versions of the time series. In particular, it is confirmed that the electricity prices are volatile, but that a considerable amount of the volatility is caused by the daily and weekly periodicities. Furthermore, the characteristic return distributions, volatility clustering and price spikes are analysed.

Preface

This thesis is written as the final part of the Master's Degree Programme in Applied Physics and Mathematics at the Norwegian University of Science and Technology. Most of the programming are coded in the Python language, but an interface to Fortran95 is also employed in order to utilize the wavelet transform. Since the thesis include analysis of several time series, there are a lot of figures to be presented. I hope that the report still is easy to follow. Furthermore, it should be mentioned that getting access to data from electricity markets is not easy. This is the reason why some of the investigated data sets are shorter than desired.

My fellow students and good friends Nikolai Hydle Rivedal and Vebjørn Tveiterås deserve a big thanks for carefully reading through the report, detecting unpleasant misprints and suggesting improvements. I am also thankful to my mathematician friends Kristoffer Hellton and Ole Thomas Helgesen for valuable discussions.

Finally, I would like to express my gratitude to Professor Ingve Simonsen, for a determined guidance throughout the work with this thesis. He always meets his students (which is an impressive amount of people!) with a friendly smile. I also appreciate that he introduced me to the wavelet transform, which certainly could be regarded as a mathematical multi-tool.

Erik Løkken Walter,
July 2011

Contents

List of Figures

Chapter 1

Introduction

I react pragmatically. Where the market works, I'm for that. Where the government is necessary, I'm for that. (...) I'm in favor of whatever works in the particular case.

John Kenneth Galbraith

There is no doubt that the idea of free markets has had an enormous impact on the world society. Ever since the pioneering thoughts of the political economist Adam Smith, liberalized markets have been the ideal for trading of commodities. However, some commodities and services, such as railway, telecommunication and electricity, require large and expensive facilities. It has therefore been regarded that the easiest way of organising such markets is to let a governmental institution be responsible for all parts of the service. Still, during the last few decades, attempts have been made to liberalize markets like these. As will be seen for the electricity markets, the liberalization process is not trivial. For that reason, in the spirit of the economist and author John Kenneth Galbraith quoted above, it is important to study evidence from the markets that have been liberalized, in order to clarify what is working and what is not.

As electricity has become one of the most important commodities of the western society, well-functioning power markets are of essential importance. In particular is the reliability of the electricity supply crucial. As will be seen, there are different solutions to how the supply capacity is ensured. Do these solutions influence the characteristics of the electricity prices? Extended understanding of the price dynamics is necessary in the search for the best way of organizing the markets. Identifying the characteristics of electricity prices could help revealing the condition of the market. Furthermore, knowledge about the price process is necessary for modelling purposes, which is important for any market participant.

Aim and Structure of Thesis

The main objective of this thesis is to investigate characteristics of time series of electricity system prices. Evidence from six electricity markets will be studied in detail, and some of the features are compared to those of other financial markets in order to get an understanding of how electricity markets differ from these. Since, as will be seen, the electricity prices exhibit periodic patterns on several time scales, it is necessary to moderate the effect of these periodicities before the electricity prices are compared to time series other types of markets. In order to explain the characteristics, it is also necessary to study the structure, in particular the price setting mechanisms, of the different markets.

Chapter 2 gives an overview of the basic theory of liberalized electricity markets, as well as case studies of the structure of six different markets. Furthermore, the mathematics of time series analysis and the wavelet transform are introduced. In Chapter 3, the two employed methods for estimating the long-range dependence of time series are given. The six analysed data sets are described in Chapter 4, before the results are presented and discussed in Chapter 5. There are two appendices attached to this report. In the first, basic statistical theory that are employed in this thesis is included. The second appendix gives an introduction to how the wavelet transform can be utilized in order to smooth a time series.

Chapter 2

Theory

In this chapter, an overview of common electricity market structures is given. The price setting mechanisms in six electricity markets are particularly investigated. Then, a few important aspects of time series analysis are discussed, before the wavelet transform, which is a fundamental part of one of the methods for estimating the Hurst exponent, is introduced in short. Finally, the characteristics later to be investigated in detail are defined at the end of the chapter.

2.1 Deregulation of Electricity Markets

Up until the 1980's, electricity was produced, transported and sold by state-owned companies. These vertically integrated companies operated in a monopoly market and they often had the full responsibility of supplying industry, households and other consumers with electricity. However, the success of liberalization of other vertically integrated economies such as telecom and railway markets, led to the belief that electricity with advantage could be traded in a deregulated market. This belief was supported by the common political ideology of free markets: it was argued that introducing competition in the market would stimulate technological innovation, allocate resources in a better way and improve the effectiveness in the supply of services [1, 2]. In addition to this ideological motivation, technological progress in generation and transmission of electricity made the change in market structure possible.

Most of the historical details of electricity market liberalization in different countries and regions will not be studied here, but it should be mentioned that the military dictatorship of Chile, ironically enough, was the pioneer in liberalizing electricity. A reform of 1982 paved the way for separate generation and distribution companies, and extensive privatization began in 1986.

The British electricity sector was reorganized in 1990, followed by the establishment of *Nord Pool*, the Nordic market, in Norway in 1992. Motivated by the experience from well developed markets, more and more electricity markets are now being liberalized worldwide [2].

The deregulation of electricity markets is still an object of heated discussions, and the experiences from other markets do not give any obvious answers: some markets have successfully been operated for up to two decades, others have failed with calamitous consequences. The breakdown of the Californian market in the beginning of this century, which led to the bankruptcy of Enron, as well as frequent electricity blackouts in Europe and North America in 2003 are used as warnings against liberalization. However, Weron [2] points out that market liberalization can not be blamed solely for these failures.

One of the main objections against the liberalization of electricity markets is the lack of incentives for market participants to invest in new generation and transmission capacity [1]. Expensive and time consuming projects are often avoided in the deregulated market, since the profit maximizing and least risky option often is to build power plants with short construction times (such as gas-fuelled plants). The risk that private electricity generation companies have to take when building new power plants is often high due to the uncertainty in the expected payback. The risk averse nature of participants in a free market may therefore lead to capacity shortage. In order to stimulate market participants to increase their generation (and transmission) capacity, the authorities may intervene in different ways. Possible interventions are discussed in the following subsections.

2.1.1 Capacity Payments

The basic idea behind capacity payments is to award companies that contribute to the *reliability* of the power system. It is a daily payment rewarded to each generator, based on the level of their availability for power generation. Even if capacity payments in theory should stimulate the construction of new capacity, it is in practice seen that the payments are weak incentives for building capacity. Capacity payments are still present in several markets, but various solutions exist. However, they are all based on the basic idea of rewarding those who contribute to generation capacity.

2.1.2 Capacity Markets

In *capacity markets*, it is common that *distribution companies* are required by law to fulfil some capacity obligations. The size of the capacity obligations may differ from market to market, but it is common that it is calculated as the company's expected monthly peak load plus a safety margin. The company can meet the capacity requirement either through their own generators (if they have any), through bilateral agreements, or by purchasing a right to buy electricity from other generators. The latter option is termed *recall rights* and are the product that is traded in the capacity markets. Since the generation companies with idle production capacity are able to sell recall rights in the capacity market, they get a motivation for investing in installation of new capacity. The capacity markets are independent markets but are at the same time closely connected to the electricity markets.

It is of uttermost importance for the electricity market that the capacity is correctly valued. If for instance a bubble were present in the capacity market, i.e. capacity was valued higher than the real value, it would encourage construction of a lot of new capacity. This would again lead to overcapacity, pinched electricity prices and lower income than expected for the utility companies. If, on the other hand, the value of capacity were underestimated, it would lead to a shortage of installed capacity, high electricity prices and possibly blackouts [3].

To summarize, electricity markets with capacity payments or parallel capacity markets may drive the amount of installed capacity to both shortage and overcapacity. This implies that a balance must be possible, but it strongly depends on the (constructed) way that the capacity (or the capacity payments) are valued.

2.1.3 Energy Only Markets

In *energy only* markets, the electricity price is the only compensation generators receive, and it should therefore compensate for both fixed and varying costs. These markets rely on the belief that a shortage of capacity will encourage investments in new power plants, since there is a demand for new plants. However, since the demand for electricity varies a great deal throughout a year, the demand in periods with low electricity consumption may not be high enough for new plants to enter the market; it may simply be too risky. The result may be that companies decline to build new capacity since it is too risky compared to the expected income. On the other hand, if the demand for new generation capacity is high, companies may find it profitable to build new plants. In this way, energy only markets may balance at a level of minimum capacity. In other markets, the balance of supply and demand is desired, since it leads to low prices for the consumers and "survival of the fittest" for the suppliers. In electricity markets, however, the supply capacity must be high enough to ensure the supply in times of high demand. Consequently, the supply capacity will have to be far greater than necessary in times of low demand.

As argued above, the incentives for suppliers to invest in generation capacity in order to meet high demand periods may not be strong enough. This may result in occasional shortage of supply which, as will be shown, again results in price spikes. If consumers are not willing to accept such possibly extreme price spikes, it is necessary for the authorities to introduce *price caps* on the electricity price. This intervention again leads to the necessity of other regulatory incentives for companies to invest in increased capacity.

The reason why several structures for ensuring generation capacity exist is that the perfect solution to the issue is not known. This acknowledgement emphasize the importance of studying evidence from the already existing market structures.

2.2 Construction of Electricity Markets

In order to understand the dynamics of price fluctuations in electricity markets, it is obviously necessary to have knowledge about and understand the price setting mechanisms of the market in question. These mechanisms include bidding structure, how the price is set based on the bids, as well as the time horizon for the bids. This section will give an overview of different market solutions with emphasis on the price setting mechanisms of two fundamental market kinds: *power pools* and *power exchanges*. However, before these are introduced, it should be emphasized what makes the trading of electricity so different from trading of many other commodities.

2.2.1 Electricity as a Commodity

Among other trade goods, electricity stands out as one of the commodities that require the most sophisticated trading facilities. The main reason for this is the physical constraint of electricity storage: *electric energy can not be stored to a large extent with today's technology*[1]. The consequence of this limitation on the market organization is substantial. It means that power generation and demand must be strictly balanced at all times. This results in the need for a transmission system operator (TSO), whose task is to balance the supply and demand. The details of how the generation-demand

[1]The only exception from this rule of thumb is the conversion of electric energy into hydro potential energy by means of hydro pumps. However, pumping facilities are not common in most countries.

equilibrium is assured will not be investigated here; it is enough to mention that the balancing procedure is handled in separate markets, often referred to as *balancing* markets. The focus of attention in this work will be on the *day-ahead markets* for electricity, whose structure and mechanisms will be investigated later on.

In addition to the non-storability feature, there is a constraint on the *transmission* of electricity. Since the transmission is bounded to the transmission grid, there are no prospects of a global electricity market with the current technology. Furthermore, as will be discussed later on, grid congestions may restrict the flow of electricity between distant market participants, setting further constraints on the free market.

2.2.2 Sub-markets within Electricity Markets

In many deregulated electricity markets, electricity is traded in different kinds of sub-markets. These have different purposes and may be described as physical or financial markets. They are categorized by means of the time horizon for the agreements traded.

Day-Ahead and Intra-day Markets

The most common market for physical delivery of electricity is *day-ahead markets*. These are markets where contracts with physical delivery the subsequent day are traded. Day-ahead markets are the main sub-markets in most electricity markets, and form the basis for the related financial sub-markets. Some market operators also operate *Intra-day markets*. These are sub-markets for delivery of electricity at the same day, which enable power generators to react to unforeseen levels of demand. Here, it should also be mentioned that in many regions, most electricity is still traded through bilateral agreements, so-called over-the-counter (OTC) transactions. An example of this is long-term agreements where the price is set through tender rounds or negotiations.

Balancing Markets

These are markets where the balance between supply and demand are adjusted by a regulator. The details of how this equilibrium is assured are different from market to market and will not be studied here.

Forward and Futures Markets

The *forward* and *futures* markets are specifically created to allow market participants to hedge against the risk of losses. The contracts offered in these markets usually have longer time horizons than a day, and the products are often traded extensively before the agreed delivery date. This indicate that some of the participants at these markets are pure traders who speculate in the forward and futures contracts.

In the following, the emphasis will be put on the day-ahead markets, since these are the most important markets for physical delivery of electricity. Also, the pricing of the financial derivatives are usually dependent on the day-ahead markets.

2.2.3 Power Pools

Power pools, also termed *economic pools*, are markets where generators compete for dispatching electricity. The pools are commonly established by governmental authorities who want to introduce competition between the power generators. Every generator evaluates what price they are willing to sell electricity for, and place bids based on that. All the bids are collected by the market administrator, and an aggregated *supply curve*[2] of the market is created based on the bids. The supply curve is then compared with an estimated demand, and the *market clearing price* (MCP), also termed *system price*, is established as the intersection of these curves, see Fig. 2.1 (a). Power pools are criticised for having a low level of transparency: the price determination is a complex optimization mechanism because the technical constraints possibly make the bids complicated [2]. In areas with power pools, it is usually not allowed to trade electricity outside the pool.

2.2.4 Power Exchanges

Power exchanges (PX) are often established as *private* initiatives to facilitate trading of electricity between generators, distributors and customers [2]. As opposed to power pools, power exchanges are based on two-sided auctions, which means that both supply *and* demand curves are calculated based on bids. The intersection of these curves determines the system price for the effective life of the bids (usually an hour), see Fig. 2.1 (b). For the day-ahead markets, it is common that the auctions are conducted once a day, where the participants place bids for each hour of the succeeding day, and an hourly

[2]A curve where the supply is given as a function of the price.

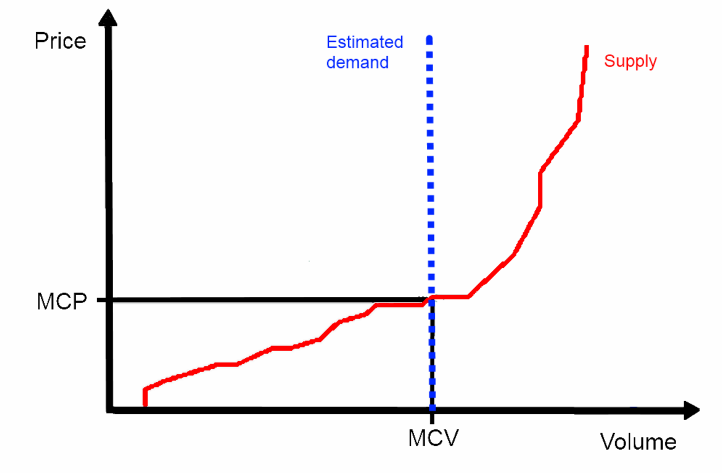

(a)

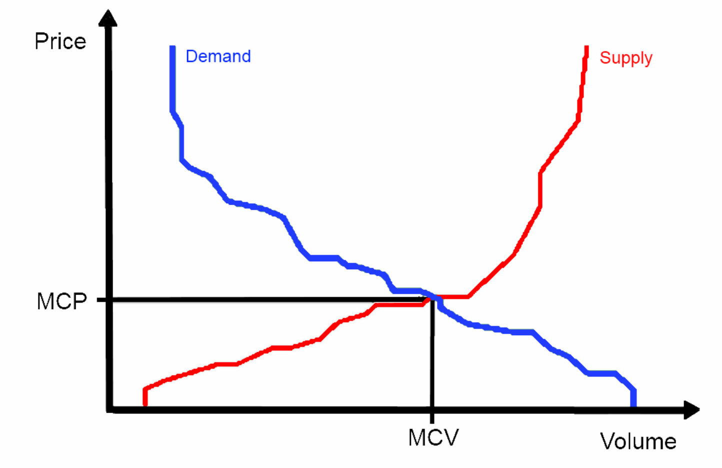

(b)

Figure 2.1: The price determining mechanism for: (a) power pools and (b) power exchanges. It is seen that in a power pool, the market clearing volume (MCV) (and demand) is estimated, while in a power exchange, the aggregated demand curve is calculated based on the bids. The market clearing price (MCP), or system price, is given as the intersection between the supply and demand curves. The figure is based on a figure taken from the web pages of Nord Pool Spot [4].

market clearing price is set based on the hourly supply and demand curves. There exist two common auction types to determine the price the customers have to pay: *uniform price* and *pay-as-bid* auctions.

Uniform price auction

In a uniform price auction, buyers with bids above or equal to the system price get to buy electricity to that price. Likewise, suppliers with offers below or equal to the system price get to dispatch electricity and are paid the system price. This means that customers who are desperate for electricity may place very high bids, but still don't have to pay any more than the system price. Likewise, a supplier may offer electricity to a price equal to his marginal cost (which is most probably his *reservation price*[3]), but is paid the system price, which may be considerably higher. In this way, the generators are rewarded when they have a marginal cost lower than other market participants. In order to prevent too high bids, many markets have a price cap, which sets a restriction to the maximal bidding level.

Pay-as-bid auction

In a pay-as-bid auction, suppliers are paid exactly the price at which they offered electricity as long as there exist buyers willing to pay that price. This type of auction is also referred to as *discriminatory* auction [2].

Both in uniform price and pay-as-bid auctions it is possible for one market participant to place multiple offers and/or bids with different prices and electricity quanta. An example of a bid/offer order from a generation company is given in Fig. 2.2. It is an order valid for a given hour of the day. From the figure, it is seen that the fictitious generation company is willing to sell V_1^{A} MWh at price level p_1^{A}. In addition, the company is willing to sell V_2^{A} MWh at price level p_2^{A}. Such a sell order could either originate from that the company has two generators with different production costs, or that one generator has a non-linear marginal production cost function. Both examples give the result that it is cheap to produce up to V_1^{A} MWh, but more expensive to produce above this level.

In Fig. 2.2, it is also seen that the fictitious generation company is willing to buy V_1^{B} MWh at price p_2^{B} NOK/MWh. This is commonly appearing in reality and is explained by the fact that many generation companies have special delivery agreements with the industry that they have to fulfil. If the electricity price is especially low, the generation company may choose to buy

[3]The lowest price at which a seller is willing to sell electricity [5].

electricity from other generators in order to fulfil these agreements. If the company mainly had hydro power plants, which are known as very flexible regarding production speed, the profit optimizing option for the company could be to buy electricity from less flexible power plants, such as nuclear power plants, who would have to produce energy anyway.

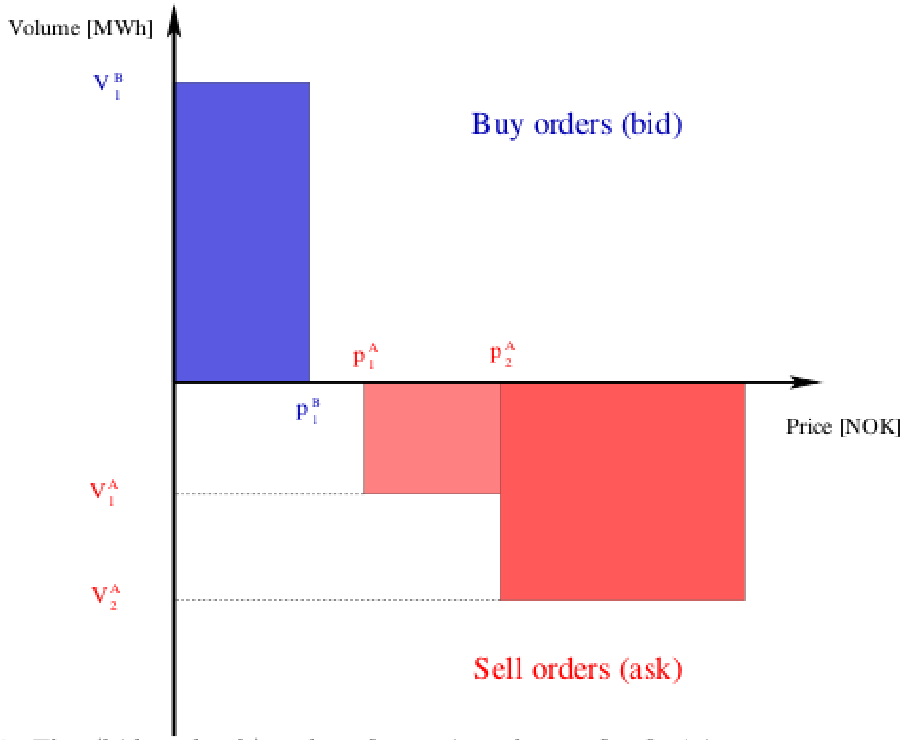

Figure 2.2: Typical bid order from a fictitious generation company. Figure taken from Simonsen, Weron and Mo [6].

2.2.5 Bottlenecks and Local Prices

Transmission of electricity in a deregulated market requires a well constructed electrical grid. However, in a market where some generators do not want to deliver electricity below a certain level of the system price, bottlenecks (grid congestions) may arise. Such congestions may reduce the availability of the cheapest electricity for some areas in the market, and these areas are forced to buy more expensive electricity from other generators.[4] It is important to

[4]As an example, it is worthwhile to mention the county of Sør-Trøndelag in Norway, which has been infamous for incoming grid congestions and shortage on self-production of electricity. During wintertime, this has often led to a higher electricity price in this area than in other parts of Norway.

be aware that local prices may differ significantly from the system price. In this work, however, only the system prices are considered, since these are the ones that reflect the nature of the market mechanisms, while specific area prices mainly attribute to limitations in the electric network.

2.3 Case Study of Electricity Markets

In the previous section, some common mechanisms and solutions in the construction of a deregulated electricity market were introduced. Now it is time to take a look at how some of the markets world wide are organized, starting with the Nordic market: Nord Pool.

2.3.1 Nord Pool

As the oldest and one of the most mature power exchanges in the world, Nord Pool[5] has been leading the way as an example for other electricity markets [2, 4]. It was established in Norway in 1992, but has later expanded to Sweden (1996), Finland (1998) and Denmark (2000). Nord Pool consists of both markets for physical delivery of electric power (*Elspot* and *Elbas*), and a financial market (*Eltermin*) where various derivatives are traded. A balancing market is operated by the transmission system operator in order to secure delivery. In this work, the focus of attention will be the large day-ahead Elspot market, which has about three times as many participants as the intra-day Elbas market. Participation in the market is voluntary. Still, Nord Pool Spot AS claimed that they had a market share of about 70% in the Nordic region in 2008 [4]. A large fraction of the power generated in Nord Pool is hydro power. This makes bids such as the ones seen in Fig. 2.2, where generation companies want to buy power in case of low prices, common. Since many generation companies have delivery agreements with the industry, they may find it profitable to buy cheap power from other generators and save their own easily adjustable hydro power.

The price setting at Elspot is a *two-sided, uniform price auction*, where the system price is given as the intersection between the aggregated supply and demand curves, as described in Section 2.2.4. Every day at noon, market participants submit their bids and offers for every hour of the succeeding day to the market administrator. There are three different ways of bidding at Elspot. *Hourly bidding* is the intuitive way of bidding, where pairs of price and volume for each hour are submitted. *Block bidding* is a special case

[5] Nord Pool has a misleading name, since it in fact is a power exchange and *not* a power pool.

of hourly bidding where the bidding price and volume are fixed for several subsequent hours. In *flexible hourly bidding*, a coupled price and volume sale offer is valid for the whole day until it is redeemed when the spot price exceeds the price indicated by the bid [2]. Nord Pool does not have a capacity market or any capacity payments, and is characterised as an energy-only market.

2.3.2 The Iberian Market: OMEL/MIBEL

Spain was the first country in continental Europe to liberalize the electricity market. Compania Operadora del Mercado Español de Electricidad (OMEL) was assigned the task of organising and managing the new market, also called OMEL, in 1997. As Portugal was included in 2007, the market changed its name to MIBEL, but OMEL is still the market operator of the MIBEL day-ahead market [2]. An additional intra-day market intends to function as an adjustment market, where only buyers who participated in the corresponding daily market auction are allowed to participate. In this work, the focus of attention will be the Spanish part of the day-ahead market. When analysing historical time series from this market, the market expansion must be kept in mind.

MIBEL is described as a *power exchange*, but as Weron [2] argues, it is actually a hybrid solution since the use of *capacity payments* usually is associated with power pools. The MIBEL daily market is characterised as a *two-sided, uniform price auction* where bids for every hour of the succeeding day are submitted before 10 a.m. every day. OMEL then matches the supply and demand and calculates the system price, which is adopted as the price of electricity for that hour. Bidders have the possibility to specify some complex bidding conditions in their bids, which makes the matching of bids more complex for the market administrator. For a description of the different complex bid structures possible in MIBEL, a reference to the market administrator OMELs website is made [7].

MIBEL is a voluntary market, but as Weron [2] argues, the capacity payments, which are only employed at the market, make trading outside the market less attractive.

2.3.3 Germany: EEX

Today's German power exchange, *The European Energy Exchange* (EEX)[8], was formed in 2002 as a fusion between two separately existing German power exchanges. EEX, or its spin-off companies, now operate both a day-ahead and intra-day power spot market, as well as day-ahead markets for natural gas and emission rights. Furthermore, EEX operates financial derivatives markets for these products. As earlier, it is the day-ahead market for electricity that is the focus here. This market is a power exchange that is operated by the European Power Exchange (EPEX)[6], which is a spin-off from EEX and the french Powernext SA. The German day-ahead market for electricity is a *two-sided, uniform price auction* where hourly bids and offers are submitted within noon every day [10]. Participation in the market is *voluntary.*

In terms of consumption, Germany represents the largest power market in Europe. However, because of the extensiveness of bilateral (over-the-counter, OTC) transaction contracts, the volume traded on EEX was lower than in Nord Pool and the Italian IPEX in 2008 [11]. Still, the trading on EEX corresponded to about one quarter of Germany's total demand in 2008, and the amount is rising [10]. The German market is described as an oligopoly since the four largest production companies have a market share of almost 85%. Moreover, the two largest production companies have a market share of approximately 55%, according to Weight and Hirschhausen [12]. In their study from 2006, they find indications that the German market still is not sufficiently competitive to overcome market abuse. Hence the prices over time are higher than expected from a perfect competition model.

2.3.4 Poland: TGE/PolPX

Towarowa Gielda Energii (TGE), in English referred to as the Polish Power Exchange (PolPX), was established in the end of 1999 and the day-ahead market was launched in July 2000 [11]. TGE suffers from a relatively small trading volume. Several causes for this are pointed out, among them inappropriate structure and high charges for participation [2]. Furthermore, long-term delivery agreements that were made before the liberalization of the electricity market started are obstacles in the market development [2]. TGE now operates an intra-day market as well as financial derivatives and CO_2 emission right markets. The physical day-ahead market is a *two-sided, uniform price auction* that finishes one day before the delivery day.[7]

2.3.5 American East: PJM Interconnection

The PJM (Pennsylvania-New Jersey-Maryland) Interconnection[14] is the world's largest competitive market for electricity wholesale. A day-ahead

[6]EPEX operates power exchanges in Germany, Austria, Switzerland and France. On their webpages, they claim that these countries account for more than a third of the European energy consumption [9].

[7]See TGE's webpages [13] for details on the bidding schedule.

market, a realtime market, two generating capacity markets (daily and long-term), a financial market allowing to hedge against price differences between locations, and an ancillary services market are operated [2]. The day-ahead market is a *two-sided, uniform price* auction. However, different uniform prices are calculated for every *node* of the transmission grid (locational marginal pricing, LMP). This means that if there is a grid congestion at the transmission lines connecting one node to the others, the node may have a price that differs from the price at the other nodes. This is reminiscent of the area prices in Nord Pool.

2.3.6 California: The Miserable Collapse of CalPX

Beginning in 1998, California was the first US state to liberalize their electricity sector. A comprehensive process full of compromises resulted in the creation of an independent system operator (CAISO) and a power exchange, CalPX [2]. The exchange conducted *two-sided, uniform price* auctions in a day-ahead market as well as an hour-ahead market. CalPX was also responsible for real-time balancing and congestion management. Also being an energy only market, the main structures were equal to the structure of Nord Pool. In year 2000 the electricity prices began rising rapidly. From the second half of 1999 until the second half of 2000, the prices increased by a factor of 50 [2]! CalPX stopped operating on Jan 31, 2001 and went bankrupt soon after. What went wrong?

Weron [2] points out a fundamental flaw in the market design of CalPX. Even if CalPX was meant to be a voluntary market, the major utilities were obligated to trade electricity through CalPX exclusively. This exposed them to an enormous risk. Weron explains [2]:

> *On the one hand, their retail revenues were fixed at the regulated rates; the utilities did not receive any additional compensation in the event wholesale prices exceeded the regulated rates. On the other, they were barred from hedging by purchasing power in advance of the day-ahead market. This restriction made the market vulnerable to manipulation. For a disaster to strike, all that was needed was a period of tight supply.*

A coincidence of several factors[8] caused the prices to rise during 2000, leading to the fatal destiny of CalPX.

[8] Recognised and explained by Joskow [15].

2.4 Basic Concepts of Time Series Analysis

A time series is a sequence of observations. The observations can generally be taken through any dimension, but are usually taken in terms of equally spaced time intervals. A well-known example of a time series is the Dow Jones Industrial Average index, given in Fig. 2.3, showing how the daily closing value of the index has evolved during the period 1928-2011. In order to describe such a time series, some important aspects of time series analysis are introduced in this section. A real-valued time series $\{Z_t : t = 0, \pm 1, \pm 2, ...\}$ will, in this thesis, commonly be denoted as $\{Z_t\}$.

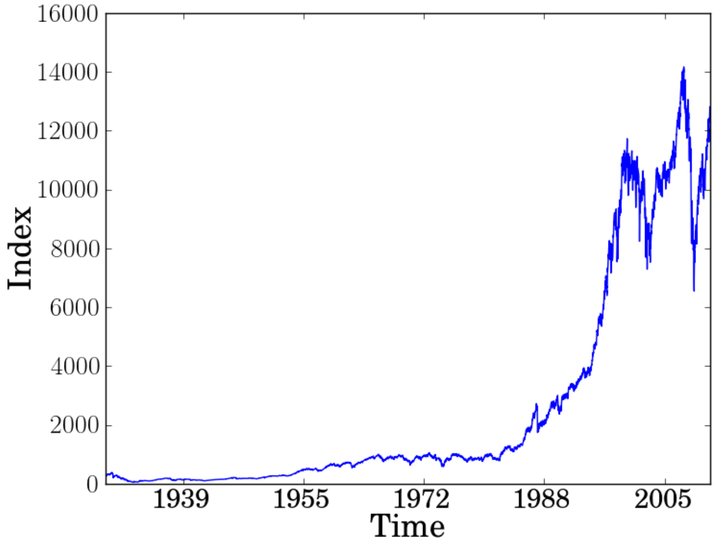

Figure 2.3: Example of a time series, showing the daily closing value of the Dow Jones Industrial Average index during 1928-2011 [16].

2.4.1 Stationarity

Strictly speaking, an n'th order distribution stationary time series $\{Z_t\}$ is a time series for which

$$F_{Z_{t_1},...,Z_{t_n}}(x_1, ..., x_n) = F_{Z_{t_1+k},...,Z_{t_n+k}}(x_1, ..., x_n),$$

for any $(t_1, ..., t_n)$ and integer k, where F is the n-dimensional cumulative distribution function defined as the probability

$$F_{Z_{t_1},...,Z_{t_n}}(x_1, ..., x_n) = P\{Z_{t_1} \leq x_1, ..., Z_{t_n} \leq x_n\}.$$

This strict definition of stationarity is however too strict for practical applications. Thus, a weaker sense of stationarity is defined as the existence and time invariance of the joint moments.

Definition 1. *A process is* n*'th order weakly stationary if all its joint moments up to order n exist and are time invariant* [17].

In this work, *stationarity* will be used as a term for second order weak stationarity. Hence a stationary time series is understood as a time series with constant mean and covariance.

If a time series has a covariance that varies with time, transformations such as the log-transform may be tried out in order to get a stationary series. If the series has a mean that varies with time, a stationary time series may be obtained by applying the *difference operator*

$$Z_t^* = \nabla Z_t = Z_t - Z_{t-1}. \tag{2.1}$$

2.4.2 The Autocorrelation Function

For a stationary time series, the autocorrelation function (ACF) is defined as

$$\begin{aligned} \rho_k &= \frac{\mathrm{Cov}(Z_t, Z_{t+k})}{\sqrt{\mathrm{Var}(Z_t)}\sqrt{\mathrm{Var}(Z_{t+k})}} && (2.2) \\ &= \frac{\mathrm{Cov}(Z_t, Z_{t+k})}{\mathrm{Var}(Z_t)} && (2.3) \end{aligned}$$

where the covariances and variance are defined as

$$\begin{aligned} \mathrm{Cov}(Z_t, Z_{t+k}) &= E\Big[(Z_t - E\Big[Z_t\Big])(Z_{t+k} - E\Big[Z_{t+k}\Big])\Big], && (2.4) \\ \mathrm{Var}(Z_t) &= E\Big[(Z_t - E\Big[Z_t\Big])^2\Big], && (2.5) \end{aligned}$$

where $\mathrm{E}\Big[\ \Big]$ is the expectation value operator. For a sample of n observations, the ACF can be estimated by

$$\hat{\rho}_k = \frac{\sum_{t=1}^{n-k}(Z_t - \hat{\mu})(Z_{t+k} - \hat{\mu})}{\sum_{t=1}^{n}(Z_t - \hat{\mu})^2}, \tag{2.6}$$

where $\{Z_t\}$ is the sequence of observations and $\hat{\mu} = E\Big[Z_t\Big]$ its sample mean. The ACF is often given graphically, where the spike at lag k indicates the correlation between points separated by $k-1$ points, cf. for example Fig. 5.5.

In this way, the ACF expresses how an observation tends to be correlated to earlier observations. Many programming languages have built-in functions for estimation of the ACF of a time series.

The confidence bands of an ACF-plot are given by

$$\Delta = \pm \frac{z_{1-\alpha/2}}{\sqrt{N}}, \tag{2.7}$$

where z is the percent point function of the standard normal distribution, α is the significance level, and N is the sample size. The confidence bands signify the interval in which $(1-\alpha)$ % of the auto-correlation values are expected to be if there in fact is no correlation present. Therefore, a value $|\hat{\rho}_k| > |\Delta|$ indicates that the time series has an auto-correlation at time lag k.

2.5 The Wavelet Transform

The wavelet transform is a mathematical tool that has been applied in numerous fields of science and engineering, ranging from compression of fingerprints and other digital images [18] to sound synthesis, and recovery of signals from noisy data [19]. The fundamental idea is that it decomposes data into a new basis, hereby called the *wavelet basis*. The data are decomposed, at a given scale, into a smooth and detailed component, and this procedure is applied recursively to coarser and coarser scales (starting from the finest scale). In the following, the wavelet transform will be roughly introduced, partly based on the approaches of Jawerth and Swelden [20] and Press, Teukolsky, Vetterling and Flannery [21].

2.5.1 The Continuous Wavelet Transform

The continuous wavelet transform (CWT) of a square-integrable function $f(t)$, is defined as [20]

$$W[f](a,b) = \int_{-\infty}^{\infty} \psi^{\dagger}_{a;b}(t) f(t) \mathrm{d}t = \langle \psi_{a;b}, f \rangle \,, \tag{2.8}$$

where

$$\psi_{a;b}(t) = \frac{1}{\sqrt{a}} \psi \left(\frac{t-b}{a} \right), \qquad a, b \in \Re \text{ and } a > 0 \tag{2.9}$$

define the wavelet basis. In equation (2.9), different a's and b's correspond to dilations and translations, respectively, of the so-called *mother wavelet* $\psi(t)$, which itself is a square-integrable function.

It can be shown [22] that the wavelet transform (2.8) is invertible if

$$C_\psi = \int_{-\infty}^{\infty} \frac{\mid \hat{\psi} \mid^2}{\omega} \mathrm{d}\omega < \infty, \tag{2.10}$$

where $\hat{\psi}(\omega)$ denotes the Fourier transform of $\psi(t)$. When this requirement is met, then the inverse wavelet transform is given by

$$f(t) = \frac{1}{C_\psi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \psi_{a;b}(t) W[f](a,b) \frac{\mathrm{d}a\mathrm{d}b}{a^2}. \tag{2.11}$$

Requirements on the mother wavelet

To meet requirement (2.10), one must have that $\hat{\psi}(0) = 0$, which leads to the requirement

$$\hat{\psi}(0) = \int_{-\infty}^{\infty} \psi(t)\mathrm{d}t = 0. \tag{2.12}$$

This means that $\psi(t)$ must oscillate around zero, i.e. $\psi(t)$ is a wave or simply a *wavelet*.[9] A wavelet ψ must hence meet two critera: it should be a square-integrable function, and it should fulfill requirement (2.12). There are many possible wavelets to choose from, all of which have different characteristics when it comes to the trade-off between localization in space (or time) and frequency. Two of these wavelets, namely the simple Haar wavelet, and a DAUB4 (introduced in the next sections), are given in Fig. 2.4. The class of which the latter is a member of, was discovered by Ingrid Daubechies in 1988, and is subject of attention in the following.

2.5.2 The Discrete Wavelet Transform (DWT)

The wavelet transform may be discretized by letting a and b take on discrete values restrictively. In particular, it can be assumed that $a = 2^i$, for integer

[9] This, together with the limited extension in the time domain, is the reason why ψ is called a *wavelet*, meaning "small wave" (French: *ondelette*) [20].

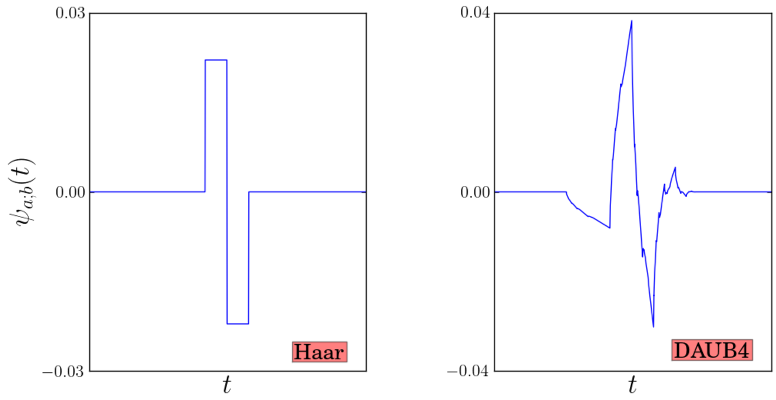

Figure 2.4: The Haar and DAUB4 wavelets. A wavelet basis can be made by scaling and translations of these mother wavelets.

i, and that b are restricted to be multiples of a. For a discrete signal $s(k)$, equation (2.8) may then be written as [23]

$$W[f](2^i, 2^i n) = \frac{1}{\sqrt{2^i}} \sum_k \psi^\dagger(\frac{k}{2^i} - n)s(k), \tag{2.13}$$

where i, n and k are integers. The inverse wavelet transform may be discretized in the same manner. In the next few subsections, it will be focused on a rough introduction to an *implementation* of the discrete wavelet transform (DWT). For a detailed description of the DWT, refer to Refs. [20, 23].

Daubechies Wavelet Filter Coefficients

In the discrete case, wavelets are specified by a set of numbers, called *wavelet filter coefficients* [21]. For Daubechies wavelets, the number of wavelet filter coefficients is referred to as *the order* of the wavelet. For simplicity, the focus in this section will be on the 4^{th} order Daubechies wavelet, DAUB4. The DAUB4 will have four wavelet filter coefficients, named c_0, c_1, c_2, c_3. These coefficients may form a transformation matrix in the following way:

$$A = \begin{bmatrix} c_0 & c_1 & c_2 & c_3 & & & & & & & & \\ c_3 & -c_2 & c_1 & -c_0 & & & & & & & & \\ & & c_0 & c_1 & c_2 & c_3 & & & & & & \\ & & c_3 & -c_2 & c_1 & -c_0 & & & & & & \\ \vdots & \vdots & & & & & \ddots & & & & & \\ & & & & & & & c_0 & c_1 & c_2 & c_3 \\ & & & & & & & c_3 & -c_2 & c_1 & -c_0 \\ c_2 & c_3 & & & & & & & & c_0 & c_1 \\ c_1 & -c_0 & & & & & & & & c_3 & -c_2 \end{bmatrix}, \tag{2.14}$$

where blank entries indicate zeroes. By letting this matrix operate on a vector to its right, a transformation of the vector will be obtained. The odd rows generate a weighted moving average of four points in the original vector, and is called the *smoothing* filter. The even rows, on the other hand, perform a "weighted differencing" and generate what is referred to as the data's *detail* information.

In order to make use of the transformation (2.14), it is necessary to be able to reconstruct the original data through an inverse transformation matrix. By requiring (2.14) to be orthogonal, the inverse transformation is given by the transpose of (2.14):

$$A^T = \begin{bmatrix} c_0 & c_3 & & & \cdots & & & & & c_2 & c_1 \\ c_1 & -c_2 & & & \cdots & & & & & c_3 & -c_0 \\ c_2 & c_1 & c_0 & c_3 & & & & & & & \\ c_3 & -c_0 & c_1 & -c_2 & & & & & & & \\ & & & & \ddots & & & & & & \\ & & & & & c_2 & c_1 & c_0 & c_3 & & \\ & & & & & c_3 & -c_0 & c_1 & -c_2 & & \\ & & & & & & & c_2 & c_1 & c_0 & c_3 \\ & & & & & & & c_3 & -c_0 & c_3 & -c_2 \end{bmatrix}. \tag{2.15}$$

It can be shown [21] that if the DAUB4 filter coefficients equal

$$c_0 = (1+\sqrt{3})/4\sqrt{2}, \quad c_1 = (3+\sqrt{3})/4\sqrt{2},$$
$$c_2 = (3-\sqrt{3})/4\sqrt{2}, \quad c_3 = (1-\sqrt{3})/4\sqrt{2},$$

then (2.14) is orthogonal, and (2.15) is hence the inverse transformation.

The Pyramidal Algorithm

Having introduced the transformation (2.14), the DWT (still using the DAUB4 wavelet) may be found through the following procedure, often called a *pyramidal algorithm* [21]:

1. Apply an $N \times N$ transformation matrix like (2.14) to a vector with N components, where $N = 2^n$ and n is a positive integer. This will generate a vector where odd rows are "smooth" data and even rows are "detailed" data.

2. Permute the data vector so that the first half contains the smooth data and the second half contains the detailed data.

3. Apply an $\frac{N}{2} \times \frac{N}{2}$ transformation matrix like (2.14) to the smooth vector of size $\frac{N}{2}$.

4. Permute the resulting data vector into smooth data and detailed data.

5. Continue until there is only two smooth datapoints left.

The following diagram, taken from Ref. [21], where d's signify detailed information and s's signify smooth information, should make the procedure clear in the case of $N = 16$:

$$
\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \\ y_7 \\ y_8 \\ y_9 \\ y_{10} \\ y_{11} \\ y_{12} \\ y_{13} \\ y_{14} \\ y_{15} \\ y_{16} \end{bmatrix}
\xrightarrow{\substack{\textit{apply} \\ (2.14)}}
\begin{bmatrix} s_1 \\ d_1 \\ s_2 \\ d_2 \\ s_3 \\ d_3 \\ s_4 \\ d_4 \\ s_5 \\ d_5 \\ s_6 \\ d_6 \\ s_7 \\ d_7 \\ s_8 \\ d_8 \end{bmatrix}
\xrightarrow{\textit{permute}}
\left[\begin{array}{c} s_1 \\ s_2 \\ s_3 \\ s_4 \\ s_5 \\ s_6 \\ s_7 \\ s_8 \\ \hline d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \\ d_7 \\ d_8 \end{array}\right]
\xrightarrow{\substack{\textit{apply} \\ (2.14)}}
\left[\begin{array}{c} S_1 \\ D_1 \\ S_2 \\ D_2 \\ S_3 \\ D_3 \\ S_4 \\ D_4 \\ \hline d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \\ d_7 \\ d_8 \end{array}\right]
\xrightarrow{\textit{permute}}
\left[\begin{array}{c} S_1 \\ S_2 \\ S_3 \\ S_4 \\ \hline D_1 \\ D_2 \\ D_3 \\ D_4 \\ \hline d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \\ d_7 \\ d_8 \end{array}\right]
\xrightarrow{\textit{etc.}}
\left[\begin{array}{c} \mathcal{S}_1 \\ \mathcal{S}_2 \\ \hline \mathcal{D}_3 \\ \mathcal{D}_4 \\ \hline D_1 \\ D_2 \\ D_3 \\ D_4 \\ \hline d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \\ d_7 \\ d_8 \end{array}\right]
\tag{2.16}
$$

To transform back to the original vector, the above algorithm is performed backwards with the inverse matrix (2.15) as the operating matrix.

2.6 Measures of Electricity Market Characteristics

Since the main objective of this thesis is to compare the statistical characteristics of different electricity markets, a few well-known measures, as well as some stylized facts, of electricity system price time series, will now be introduced and defined. As a point of reference and an example of a time series of hourly system prices, the analysed time series from the Nord Pool day-ahead market is given in Fig. 2.5.

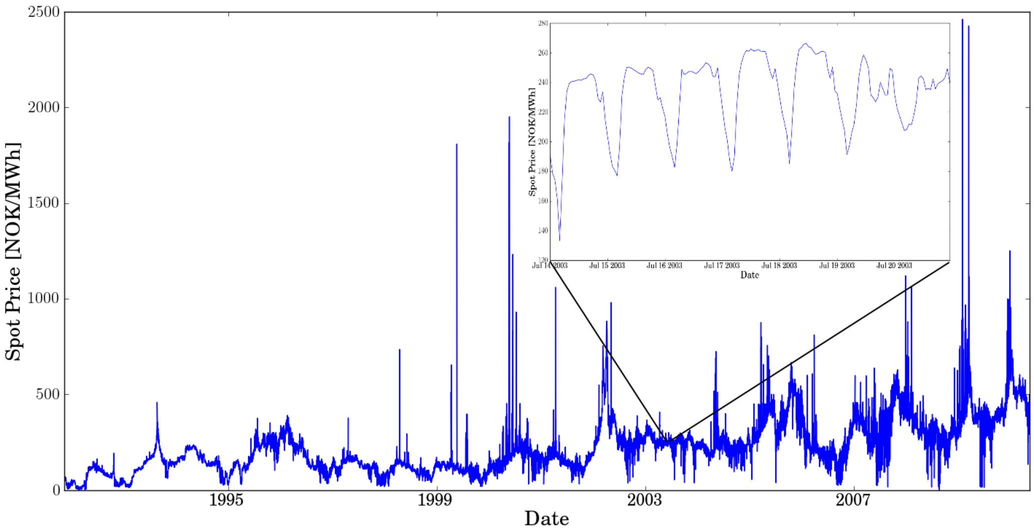

Figure 2.5: Time series of the hourly Nord Pool system price from May 4, 1992 until May 4, 2011. A plot of an arbitrarily chosen week is added in the top right corner, revealing a daily periodic pattern.

2.6.1 Spikes

As is visualized in Fig. 2.5, spikes are commonly appearing in time series of electricity prices. Spikes are by definition *short intervals in time* at which the price *surpasses a threshold value* [24]. However, the time interval and threshold value may be defined in different ways, both static and time varying. As Trueck, Weron and Wolff [24] point out in their literature study, it seems hard to reach a general agreement on what the time interval and threshold value should be. Due to this ambiguous definition of spikes, several ways of identifying spikes have been proposed.

The literature study of Trueck et al. [24] reveals the following approaches. The simplest is to set a fixed price threshold, and classify all prices that exceed the threshold price as spikes [25]. However, in time series with a non-constant mean, this definition may lead to a misleading classification since some periods of time might have persistent, natural high prices. To overcome this, some authors suggest instead to set a threshold, fixed or time-varying, for the log-price *change*[10] [26, 27]. Yet another possibility is to identify spikes as prices above a *varying* threshold value, given by the moving mean price plus a number of moving standard deviations [28]. This way of identifying spikes is robust to non-stationary time series with a significant price trend. However, as will later be seen, clustering of the volatility is common in the spot price time series. In periods with high volatility, the moving standard deviation will also be high, the spike threshold may grow unreasonably high and hence reject some spikes as fluctuations within the "normal". Of course, the discussion then falls back to the definition of a spike. The employment of wavelets to separate spikes has also been proposed [2, 29].

It is obvious that the different methods to identify the spikes will yield different results, but it is a highly subjective manner which definition that suits different tasks the best. In this thesis, it is desirable to identify spikes only in order to compare the level of "spikedness" in the different markets. For this purpose, a spike will be defined as any value above the moving weekly mean price plus three standard deviations. The moving mean price is calculated from an hourly price time series P_t as

$$\overline{P}_{t,\text{week}} = \frac{1}{168} \sum_{i=1}^{168} P_{t-i}, \tag{2.17}$$

and the standard deviation σ_P is calculated over the entire series.

Definition 2. *A spike is defined (in the context of this thesis) as any value above the time varying threshold value* $\Theta_t = \overline{P}_{t,\text{week}} + 3\sigma_P$.

2.6.2 Returns

As claimed in the previous section, electricity spot prices may change drastically within short periods of time. An extensively used measure of price changes in financial markets is the return. The return is simply the relative

[10]Which equals the logarithmic *return*, cf. Section 2.6.2.

change in price,

$$R_{\Delta t}(t) = \frac{p(t+\Delta t) - p(t)}{p(t)} = \frac{\Delta p(t)}{p(t)}, \tag{2.18}$$

within a certain time period, Δt. An alternative quantity is the *logarithmic return*,

$$r_{\Delta t}(t) = \ln\left(\frac{p(t+\Delta t)}{p(t)}\right) = \ln\left(1 + \frac{\Delta p(t)}{p(t)}\right). \tag{2.19}$$

It can easily be shown, through the Taylor series expansion of the natural logarithm[11] for $-1 < \Delta p(t)/p(t) \leq 1$, that these two quantities are equal to the lowest order:

$$\begin{aligned} r_{\Delta t}(t) &= \ln\left(1 + \frac{\Delta p(t)}{p(t)}\right) \\ &\approx \left(1 + \frac{\Delta p(t)}{p(t)} - 1\right), \qquad\qquad \text{for } \ |\Delta p(t)/p(t)| \ll 1 \\ &= R_{\Delta t}(t). \end{aligned}$$

Since $|\Delta p(t)/p(t)| \ll 1$ usually holds for stock markets, the return (2.18) and logarithmic return (2.19) are often not distinguished in the literature. However, the presence of price spikes in electricity markets may result in $\Delta p(t)/p(t) > 1$, or even $\Delta p(t)/p(t) \gg 1$. Consequently, the return and logarithmic return should be distinguished when analysing power markets. In this report it is chosen to analyse the logarithmic returns, since these are less sensitive to large increments in the system price.

Originally, the return quantity was introduced as a way to measure the gain or loss of an investment [30], but the technological restrictions on storing electricity makes it hard to buy an amount of electricity today and sell it tomorrow. Therefore it is questionable whether the return is a meaningful quantity to study for the power markets. Still, the logarithmic return of electricity markets will be studied here, in order to compare this quantity with results from other financial markets. Furthermore, the return is a necessary quantity to establish in order to investigate other interesting characteristics such as the volatility. Precautions should, however, still be made. The concept of return as a measure of unforeseen changes in the price only makes sense under the assumption that the data in question have no seasonality.

[11] $\ln x = \sum_{n=1}^{\infty} (-1)^{n+1} \frac{(x-1)^n}{n}$, for $0 < x \leq 2$.

The exception from this rule of thumb is when Δt fulfil $\Delta t = nT$, where T is the period of the seasonality and n an integer. As electricity system prices often exhibit periodic behaviour on several time scales, the seasonal patterns should ideally be filtered out before calculating the return. This filtering may be a challenging task, since it is hard to define exactly how much of a price level that is caused by the seasonal pattern, and how much that is caused by other phenomena. In order to get around any daily periodic pattern, return periods of $\Delta t = n \cdot 24$ hours could be considered. A weekly cycle could be phased out in a similar manner, but in order to calculate the daily returns it may be desirable to try to remove the weekly periodicity instead, cf. Section 5.1.1. Finally, when calculating the returns over a few days, or at most a few weeks, any annual periodicity may be neglected [6].

2.6.3 Volatility

The (logarithmic) volatility is defined as the standard deviation of the (logarithmic) return [30]. In this work, the term volatility will refer to the logarithmic volatility unless the opposite is clearly stated. A stylized fact of electricity markets is that the electricity prices are very volatile compared to other commodity or financial markets. Some common volatility values are listed in table 2.1.

Table 2.1: Typical values for daily volatilities found in different kinds of markets [2, 6].

Market	Volatility [%]
Stock indices	$1 - 1.5$
Individual stocks	< 4
Bonds	< 0.5
Crude oil	$2 - 3$
Natural gas	$3 - 5$
Short-term interest rates	0.03
Electricity markets	Up to 50!

Chapter 3

Hurst Analysis

In 1951, the British hydrologist H.E. Hurst posted his study [31] of an 850 year long record[1] of the Nile River's overflows. He had observed that a year with large river overflow tended to be followed by still larger overflows, while years with small overflow was likely to be followed by other years with small overflows. The time series of the overflow levels tended to *persist* its value, before it suddenly changed. It seemed that cycles, as well as some kind of long-termed memory, were present in the data. Since standard analysis did not uncover long-term correlation, Hurst developed a new measure called the *rescaled range*. The scaling properties of the rescaled range gives the *Hurst exponent*, which may reveal the presence of long-term correlations in a time series [32].

3.1 The Rescaled Range and the Hurst Exponent

For a discrete time series $\{Z_i : i = 1, 2, ..., N\}$ of length N, from now on denoted in short-form as $\{Z_i\}$, the rescaled range can be calculated by first dividing $\{Z_i\}$ into d subsequences of length n, such that $N = dn$. For each subsequence $\{X_j^m : j = 1, 2, ..., n\}$ with $m = 1, ..., d$, perform the following algorithm [2, 33]:

1. Calculate the mean μ_m and standard deviation S_m of $\{X_j^m\}$.
2. Create the mean-adjusted accumulated series $\{Y_j^m\}$, such that $Y_j^m = \sum_{k=1}^{j}(X_k^m - \mu_m)$ for $j = 1, 2, ..., n$.

[1] The Egyptians kept record of the overflows for a long time!

3. Find the range $R_m(n) = \max(Y_1^m, Y_2^m, ..., Y_n^m) - \min(Y_1^m, Y_2^m, ..., Y_n^m)$.

4. The rescaled range for $\{X_j^m\}$ is then $(R/S)_m(n) = R_m(n)/S_m(n)$.

Finally, the characteristic rescaled range for subsequences of length n is found as the mean value of the rescaled ranges of each subseries, i.e.

$$\overline{(R/S)}(n) = \frac{1}{d}\sum_{m=1}^{d}(R/S)_m(n). \tag{3.1}$$

As will be seen in the next paragraph, the Hurst exponent may be be found by calculating the characteristic rescaled range for subsequences of different lengths n, and finding a scaling relation for these rescaled ranges.

The Hurst Exponent

The Hurst exponent, H, is implicitly estimated through the asymptotic behaviour of the rescaled range [33]:

$$\overline{(R/S)}(n) \sim Cn^H, \tag{3.2}$$

where C is a constant. H is strongly connected to the autocorrelation of a time series, and is investigated in order to identify some of the correlation characteristics of the series. A process with $0 < H < 0.5$ will have negative auto-correlation, it is anti-correlated, meaning that the values are likely to alternate around its constant mean. If $0.5 < H < 1.0$, the process has positive autocorrelation, meaning that the series has a higher probability of sustaining its either positive or negative value. The increments of an ordinary random walk process, which are uncorrelated and commonly termed white noise, will have a Hurst exponent of $H = 0.5$.

For time series such as the electricity system prices, it is of interest to investigate the Hurst exponent of the increments, i.e. the differenced series (cf. (2.1)). This would indicate the nature of the movements of the system prices. If $H < 0.5$ one would expect that the system prices usually do not move far away from the mean over large time periods. In such cases, it is said that the prices are anti-persistent or mean-reverting. For $H > 1/2$, on the other hand, the increments could have the same sign for a longer period of time, and the system price may hence move away from the mean for a longer period of time. In this case, the price is said to be persistent. Persistent time series are known to be smoother than anti-persistent ones.

The entrance of the rescaled range analysis led to the development of models for long memory processes. Mandelbrot and van Ness [34] introduced the *fractional Gaussian noise*, which by definition is a stationary process with an ACF given by [35]

$$\rho_k = \frac{1}{2}\Big(|k+1|^{2H} - 2|k|^{2H} + |k-1|^{2H}\Big), \tag{3.3}$$

where k is the time lag, see Section 2.4.2. Fractional Gaussian noise can be considered as the increments of *fractional Brownian motion*, which also is characterized by the parameter H [35]. This means that both fractional Gaussian noise *and* fractional Brownian motion are characterized by H. Hence, the statement that a time series is characterized by some H may either imply that the series itself or that the differenced series has an ACF that can be approximated by (3.3). From now on, when it is said that the Hurst exponent of a time series is estimated, it is implicitly meant that the time series itself has an ACF that can be approximated by (3.3), i.e. the time series can be compared to fractional Gaussian noise. As will be seen, this may be the case for the increments of electricity system prices.

To estimate the Hurst exponent of a time series is in general a difficult task. The simple estimation method based on the scaling of the rescaled range, termed R/S analysis, has proven to produce increasingly biased estimates as H drops below 0.5 [36]. Moreover, R/S analysis is not able to distinguish presence of long-term correlation from non-stationarity. Since earlier evidence has suggested that $H < 0.5$ for the increments of electricity prices [2, 37], the R/S analysis is omitted here. However, several other estimation methods are suggested in the literature. Empirical studies have also been carried out in order to agree upon which methods that produce the best estimates [38, 39]. Unfortunately, it seems that there is no or little consensus, and that each method has different advantages and disadvantages depending on the time series under investigation.

In this work, two independent methods termed Detrended Fluctuation Analysis and the Average Wavelet Coefficient method, are employed. Furthermore, in order to introduce the wavelet based method, another method based on the Fourier transform is briefly introduced.

3.2 Detrended Fluctuation Analysis

The Detrended Fluctuation Analysis (DFA), introduced by Peng et al. in 1994 [40], is a method that, as opposed to the R/S analysis mentioned above, can handle a non-stationary mean. Given a time series $\{Z_t : t = 1, 2, ..., N\}$, the DFA begins by dividing $\{Z_t\}$ into d subseries $\{X_i^m : i = 1, 2, ..., n\}$ of length n (with $m = 1, 2, ..., d$). Then, for each subseries [2]:

1. Define the accumulated subseries as $Y_i^m = \sum_{j=1}^{i} X_j^m$ for $i = 1, 2, ..., n$.
2. Perform a least squares fit of the line $\widetilde{Y}^m(i) = a_m i + b_m$ to $\{Y_i^m : i = 1, 2, ..., n\}$.
3. Calculate the root mean square (RMS) of $\left\{\left(Y_i^m - \widetilde{Y}^m(i)\right) : i = 1, 2, ..., n\right\}$, i.e.

$$\mathrm{RMS}(n, m) = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(Y_i^m - \widetilde{Y}^m(i)\right)^2}.$$

Finally, calculate the mean RMS for all subseries of length n,

$$\overline{\mathrm{RMS}}(n) = \frac{1}{d}\sum_{m=1}^{d} \mathrm{RMS}(n, m). \tag{3.4}$$

As the rescaled range, $\overline{\mathrm{RMS}}(n)$ will obey a power law scaling $\overline{\mathrm{RMS}}(n) \sim cn^H$ (if the original series $\{Z_t : t = 1, 2, ..., N\}$ can be approximated as fractional Gaussian noise). The slope of a linear least squares fit to $\log(\overline{\mathrm{RMS}}(n))$ versus $\log(n)$ would thus give an estimate $\hat{H}_{\mathrm{DFA}}$ of the Hurst exponent.

To the knowledge of the author, it is not yet developed any distribution theory for $\hat{H}_{\mathrm{DFA}}$. However, Weron [41] performed a Monte Carlo study which resulted in confidence intervals for $\hat{H}_{\mathrm{DFA}}$ in the case $H = 0.5$ (for different sample sizes N). Using his results, a test against a null hypothesis of having $H = 0.5$ can be performed. From this, a conclusion of whether $\hat{H}_{\mathrm{DFA}}$ is significantly different from $H = 0.5$ or not can be drawn. Unfortunately, confidence intervals for H can not be obtained using the results of this study.

As already mentioned, a method based on the Fourier transform will now be introduced as a logical step before introducing the wavelet method. It should be noted that these two methods assume that the input is fractional Brownian *motion*, i.e. employing these methods to a time series should yield equivalent results for H as employing DFA to the increments of the same series.

3.3 Power Spectral Density Methods

The Hurst exponent can also be estimated by first estimating the power spectral density of the time series and then fit a power-law to the spectrum. If the series is a fractional Brownian motion, the power spectrum will scale approximately as $|f|^{-1-2H}$ for small f [35].

3.3.1 The Periodogram

For a discrete time series $\{x_t : t = 1, ..., n\}$, the power spectrum may be estimated by the *periodogram* [2]

$$I_n(\omega_k) = \frac{1}{n}\Big|\sum_{t=1}^{n}\left(x_t e^{-i(t-1)\omega_k}\right)\Big|, \tag{3.5}$$

where $\omega_k = 2\pi k/n$ are the Fourier frequencies, for integers $k = 1, ..., \mathrm{int}(n/2)$. Here, $\mathrm{int}(n/2)$ denotes the largest integer smaller than or equal to $n/2$. If n is a power of two, i.e. $n = 2^N$ for integer N, then the Fast Fourier Transform algorithm can be employed. The line that appears in a plot of $\log(I_n(\omega_k))$ versus $\log(\omega_k)$ will then have a slope $\beta = -1 - 2H$ which may be estimated by a least squares fit, implicitly estimating H.

In a log-log scatter plot of $I_n(\omega_k)$ versus ω_k, most points will be at the right hand side of the plot[2]. This causes the large frequencies to be overweighted in a regular least squares fit. Since the power law decay of the periodogram is most pronounced for small frequencies, as stated above, the overweighting is not desired. To get around this, one may resample the periodogram such that all points are approximately equidistant in the log-log plot. On the other hand, due to the limited length of $\{x_t\}$, there will be increasingly large estimation errors associated with the $I_n(\omega_k)$ at low frequencies. Therefore, one could say that it is a subjective manner which points should be counted in the regression and which should not.

The periodogram is not investigated further in this work. However, there exist possible modifications to the method that will improve the estimation of H [2]. Instead of investigating these methods, a less famous approach of estimating the power spectral density will now be introduced. This wavelet based approach is chosen because the employment of the wavelet transform bring along some desired properties, as will be seen.

3.3.2 The Wavelet Variance

The power spectral density of a discrete signal may also be estimated by means of the discrete wavelet transform (DWT), which was introduced in Section 2.5. There are several reasons for employing the DWT for this purpose. First, the DWT is known to *decorrelate* a wide variety of time series that frequently occurs in physical sciences. The application of common statistical theory can be more appropriate on the uncorrelated wavelet coefficients

[2]The reason for this is that a decade at the right hand side, for instance $[10^{-1}, 10^0]$, spans a wider interval than a decade at the left hand side.

than on the original time series[3]. Second, the fast DWT algorithm is faster than the FFT algorithm [35]. Third, non-stationary effects that can affect the statistical measures in time space, may have less influence on the measures in the wavelet space. Last, but not least, a wavelet based approach can be developed to be robust against any extreme-valued artefact. In the case of electricity prices, the presence of extreme valued spikes will influence the entire Fourier spectrum[4]. The spikes are usually less distinct in the daily average prices than in the hourly prices, but their influence on the Fourier spectrum can still be considerable. In the wavelet space, however, a spike will only contribute to coefficients nearby in time and the spike can thus be ignored, *if* that is desirable.

By means of the DWT, the power spectral density may be estimated by estimating the *wavelet variance*, namely the variance of the wavelet coefficients corresponding to each dyadic time scale. Since the expectation values of the wavelet coefficients by definition are zero, the following simple estimator for the wavelet variance at scale j is obtained:

$$\hat{\nu}_j^2 = \widehat{\mathrm{Var}}_b\{W[X_t](a,b)\} \equiv \langle |W[z(t)](a,b)| \rangle_b = \frac{1}{N/2^j} \sum_b \Big| W[X_t](a,b) \Big|^2. \tag{3.6}$$

Here, $N/2^j \equiv N_j$ is the number of wavelet coefficients at scale $a = 2^j \Delta t$, where Δt is the time difference between two subsequent points in the analyzed time series $\{X_t\}$. Similarly as the power spectral density scales as $|f|^{-(2H+1)}$ for a fractional Brownian motion $\{X_t\}$, the wavelet variance will scale as $\hat{\nu}_j^2 \sim a^{2H+1}$.

Instead of estimating H by means of the wavelet *variance*, the corresponding wavelet *standard deviation* could be employed,

$$\hat{\nu}_j = \sqrt{\hat{\nu}_j^2} \sim a^{H+1/2}. \tag{3.7}$$

An estimate of H is found through a linear least squares fit of $\log(\hat{\nu}_j)$ versus $\log(a)$. In the next section, confidence intervals for $\hat{\nu}_j$ are developed.

3.3.3 Confidence Intervals for the Wavelet Standard Deviation

A linear least squares fit of the log-log plot of the wavelet standard deviation versus the time scale will produce a fair estimate of the Hurst exponent.

[3]For a demonstration of this, think of how the sample variance of a time series with long positive correlations could fail to estimate the true variance of the process.

[4]Think of the resulting Fourier transform of a Dirac delta function. It will give contributions to the entire frequency domain.

However, a better estimate can be found by making certain assumptions on the statistics of the wavelet coefficients.

Since the number of wavelet coefficients are decreasing with the scale ($N_j = N/2^j$), the wavelet standard deviations at large time scales are calculated from only a few number of coefficients. This brings on an increasing uncertainty associated with the standard deviation estimates (3.7) at large time scales. By calculating a confidence interval for each estimate of the wavelet standard deviation, the width of these intervals can be used as basis for a *weighted* least squares fit of $\log(\hat{\nu}_j)$ versus $\log(a)$. The confidence interval for $\hat{\nu}_j$ can be calculated as explained in the following.

First, assume that wavelet coefficients corresponding to the same time scale make up a random sample from a Gaussian distribution with zero mean and standard deviation $\nu_j = \mathrm{STD}_b\{W[X_t]_{a,b}\}$. As Percival and Walden [35] argues, the assumption of having a random sample, i.e. no auto-correlation, is in fact a reasonable approximation because of the decorrelation property of the DWT. The assumption that the wavelet coefficients are Gaussian distributed may not always be true. However, this assumption is made only in order to be able to say *something* about the confidence intervals and hence attribute weights in the weighted least squares fit.

Second, under the assumptions stated above, the sum of squares of the wavelet coefficients at scale j will correspond to a sum of squares of independent and identically, normally distributed random variables. This means that the sum in (3.6) divided by its expectation value is chi-square distributed with N_j degrees of freedom, i.e.

$$\frac{\hat{\nu}_j^2}{\nu_j^2} N_j \sim \chi^2_{N_j}. \tag{3.8}$$

Consequently, the square root of this is chi distributed,

$$\frac{\hat{\nu}_j}{\nu_j}\sqrt{N_j} \sim \chi_{N_j}. \tag{3.9}$$

Having stated this, a confidence interval for the wavelet standard deviation can be found. Letting $Q_{N_j}(p)$ signify the $p{\cdot}100$ % point of the χ_{N_j} probability density function, i.e.

$$\mathrm{P}\Big\{\chi_{N_j} \leq Q_{N_j}(p)\Big\} = p, \tag{3.10}$$

the $(1-p){\cdot}100$ % confidence interval for the chi distributed variable in (3.9) will be given by

$$\mathrm{P}\Big\{Q_{N_j}\Big(\frac{p}{2}\Big) \leq \frac{\hat{\nu}_j}{\nu_j}\sqrt{N_j} \leq Q_{N_j}\Big(1-\frac{p}{2}\Big)\Big\} = 1-p. \tag{3.11}$$

By manipulation of this expression, a $(1-p)\times 100\%$ confidence interval for the wavelet standard deviation ν_j at scale j can be found as

$$\left[\frac{\hat{\nu}_j\sqrt{N_j}}{Q_{N_j}\left(1-\frac{p}{2}\right)}, \frac{\hat{\nu}_j\sqrt{N_j}}{Q_{N_j}\left(\frac{p}{2}\right)}\right]. \tag{3.12}$$

3.3.4 The Roughness Parameter

By generalizing (3.6) and (3.7), a generalized Hurst exponent, referred to as the *Hölder* exponent or the *roughness parameter* [42], denoted $H(q)$, may be introduced through the scaling of the qth moment [43]:

$$\left\langle \left| W[X_t](a,b) \right|^q \right\rangle_b^{1/q} \sim a^{H(q)+1/2}. \tag{3.13}$$

For so-called *mono-affine* processes, $H(q)$ is not dependent on q. However, for some processes, termed *multi-affine* processes, $H(q)$ can vary with q. In these cases, the Hurst exponent is taken to be $H(q=1)$. There have been indications during this work that some of the electricity system price series may be multi-affine. Therefore, the Hurst exponents of all the system price increment series are estimated through the scaling of the first moment, i.e. $q=1$. The wavelet based method based on $q=1$ is termed the Average Wavelet Coefficient (AWC) method [37].

The disadvantage of choosing $q=1$ is that the distribution of $\left\langle \left| W[X_t](a,b) \right| \right\rangle_b$ does not have a closed-form expression[5]. The distribution could be estimated empirically, but this is outside the scope of this thesis. For the purpose of determining the weights for performing a weighted least squares fit, as described above, it will here be assumed that $\left\langle \left| W[X_t](a,b) \right| \right\rangle_b$ (the wavelet absolute mean) follows the same distribution as $\left\langle \left| W[X_t](a,b) \right|^2 \right\rangle_b^{1/2}$ (the wavelet standard deviation). Hence, the confidence interval from (3.12) is adopted, but $\hat{\nu}_j = \left\langle \left| W[X_t](a,b) \right|^2 \right\rangle_b^{1/2}$ is replaced by $\hat{\nu}_j{}^* = \left\langle \left| W[X_t](a,b) \right| \right\rangle_b$.

[5]If X_t is normally distributed, then $|X_t|$ is said to be *half-normally* distributed. To the author's knowledge, there exists no closed-form expression of the distribution of a sum of half-normally distributed random variables.

Chapter 4

Description of the Data

4.1 Nord Pool

The analysed time series from Nord Pool is provided by Nord Pool Spot AS [4] and consists of the hourly Elspot market clearing price from May 4, 1992[1] until May 4, 2011. This adds up to a total number of 166 560 data points, making it the longest time series analysed in this work. The series is already introduced in Section 2.6, Fig. 2.5. The arbitrarily chosen week 29 of 2003 is added to the plot in order to graphically show the common daily pattern in the electricity prices.

4.2 OMEL

The data from MIBEL/OMEL consists of hourly system prices in the Spanish branch of MIBEL from Jan 2, 1998 until May 28, 2011, resulting in a time series of 117 480 data points shown in Fig. 4.1. Since the time series is from the Spanish branch of the market, it will be denoted OMEL in this report. The series can be downloaded from OMEL's web pages [7]. Having a first look at the series, it seems that it is vertically "fatter", meaning that it fluctuates more on small time scales than the Nord Pool series.

The OMEL time series includes a total of 367 zero valued data points. In order to be able to analyse the logarithm of the series, the zero points are replaced with values of 0.1 EUR/MWh.

[1] Which was the very first Elspot auction day.

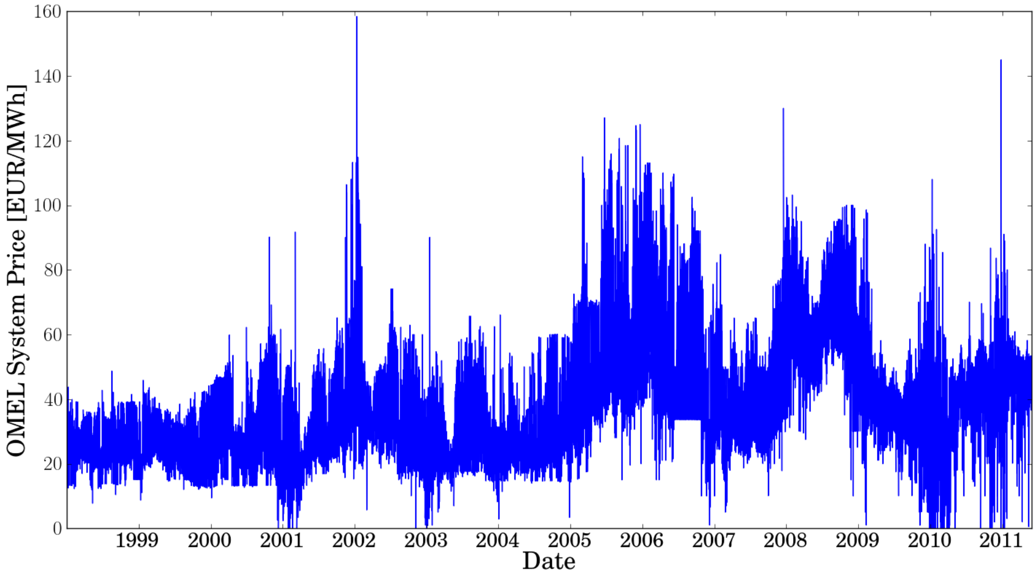

Figure 4.1: Time series of the hourly OMEL market clearing prices from Jan 2, 1998 until May 28, 2011.

4.3 EEX

The analysed data from the German EEX market consists of 29 664 data points and covers the time period from Jan 1, 2007 until May 20, 2010. The time series of hourly system prices is shown in Fig. 4.2. An interesting feature of this series is seen clearly here: the EEX system price has several data points below zero. In fact, as many as 93 negative valued data points are registered in this time period. This means that, during these short periods of time, consumers are paid to use electricity. Such a feature is hardly seen in any other financial or commodity markets without having any strings attached. In order to find a possible explanation for this mechanism, one has to remember the special characteristics of electricity as a commodity, described in Section 2.2.1: electricity *can not be stored* to a significant extent. In addition, the German power market has a considerable high amount of nuclear power plants[2], which are known to have a low flexibility regarding production volume adjustments. When the demand is low, this may lead to excess of power. Once the electricity is produced, it needs to be distributed

[2]However, the German government announced in May 2011 that they will shut down all German nuclear power plants within year 2022 [44].

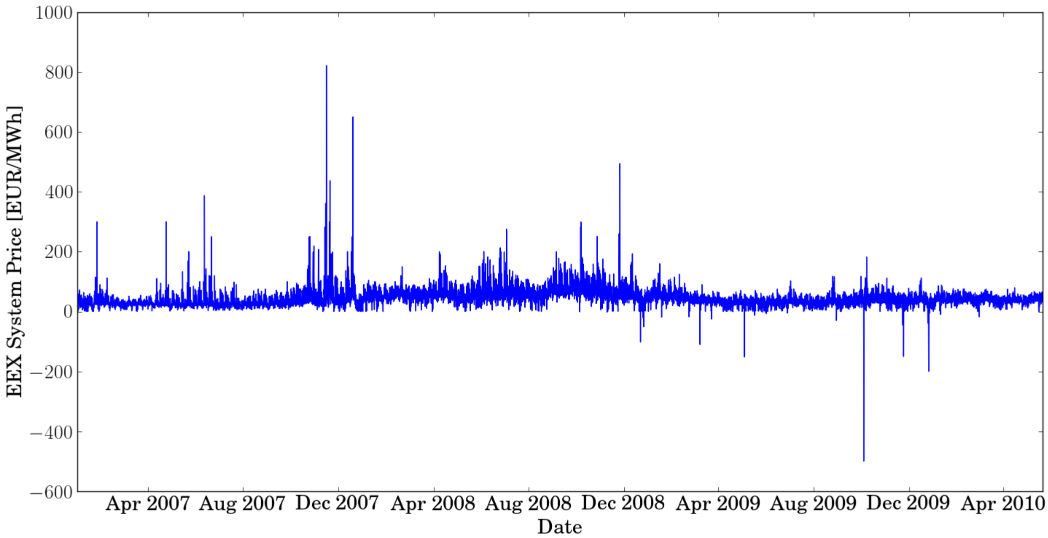

Figure 4.2: Time series of the hourly EEX market clearing prices from Jan 1, 2007 until May 20, 2010. A total of 93 hours had negative valued system prices.

and consumed instantaneously. In extreme cases, consumers are then paid to receive electricity. Fig. 4.3 shows histograms of which hours of the day, as well as months of the year, that the 93 negative values occurred. As expected, it is seen that they occurred during night and mid-day, the low-consumption time periods. The events seem to spread throughout the whole year, which also is not that surprising. As a mid-European country, Germany's electricity consumption is not as season dependent as for instance in the Nordic countries (heating during winter) or California (air-conditioning during summer).

In order to be able to employ the logarithmic transform to the series, the 93 negative valued are replaced by 0.1 EUR/MWh. The same goes for the 50 zero valued data points. Some of the results of the analysis to be performed in the next chapter are marginally sensitive to the choice of the replacement value. This goes especially for the kurtosis and skewness[3] of the logarithmic return distributions. The replacement value is a highly subjective manner, but it is reasonable to choose a value that is close to zero but still large enough to not result in too large return values. For replacement values of 0.01 EUR/MWh and 1.0 EUR/MWh, the volatility of the series remains the

[3]Defined in Appendix A.

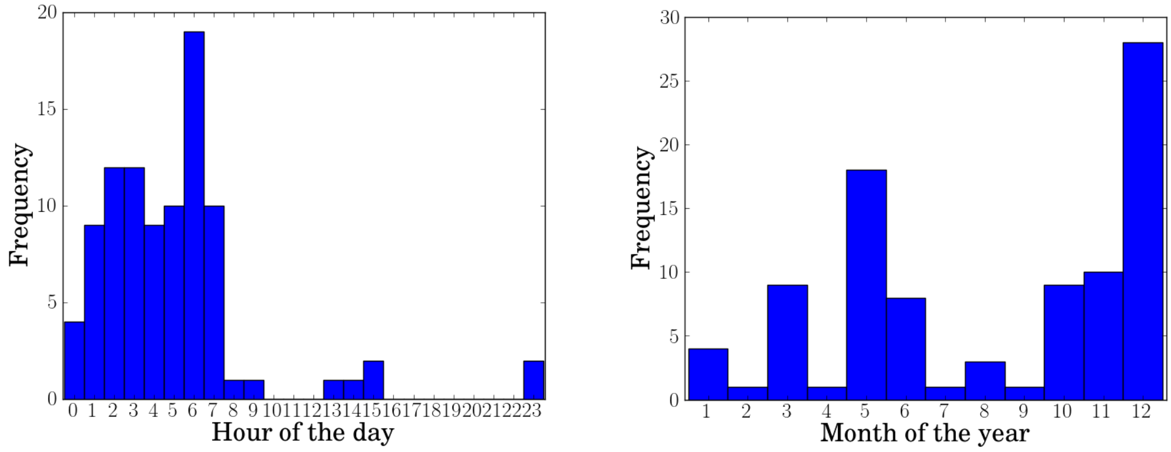

(a) (b)

Figure 4.3: Histograms showing in which (a) hours of the day and (b) months of the year the EEX system prices have been negative valued.

same.

4.4 TGE

The analysed data from the Polish power exchange is a time series, shown in Fig. 4.4, of the hourly market clearing price going from Jan 1, 2007 until Dec 4, 2010, giving a total of 34 416 data points. There are three zero valued data points which are replaced by 0.1 PLN/MWh.

4.5 CalPX

Fig. 4.5 shows the analysed time series from the Californian market, consisting of 24 888 hourly market clearing prices going from Apr 1, 1998 until the market was shut down on Jan 31, 2001. It is apparent that the market changed during this period. The last part of the series consists of severe fluctuations and extreme price spikes. The all-time high value occurred on Jan 21, 2001, and measured 2499.6 \$/MWh, approximately 38 times the mean value of the series.

There is a total number of 152 zero valued points in the time series, all of which are replaced by 0.1 \$/MWh.

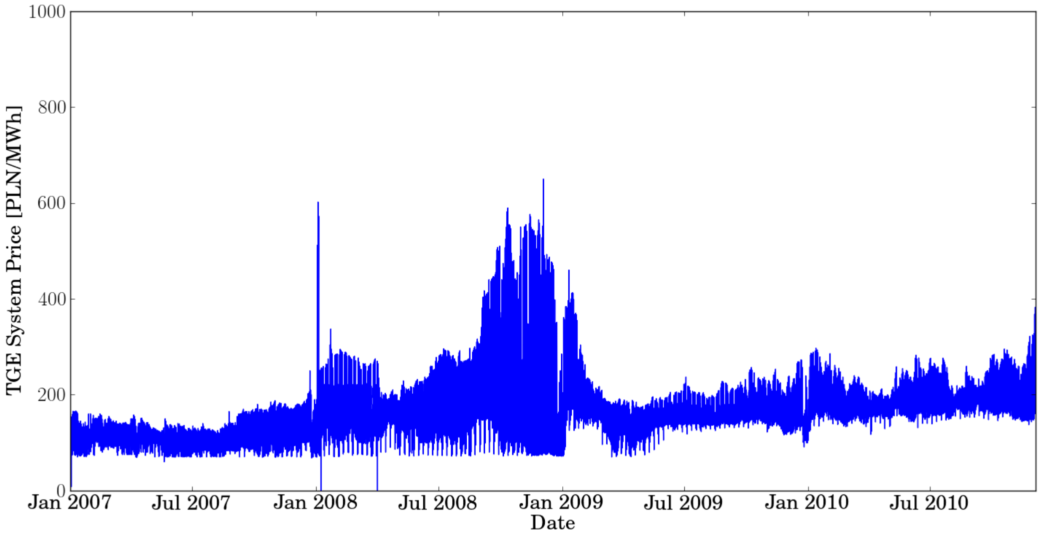

Figure 4.4: Time series of the hourly Polish TGE market clearing prices from Jan 1, 2007 until Dec 4, 2010.

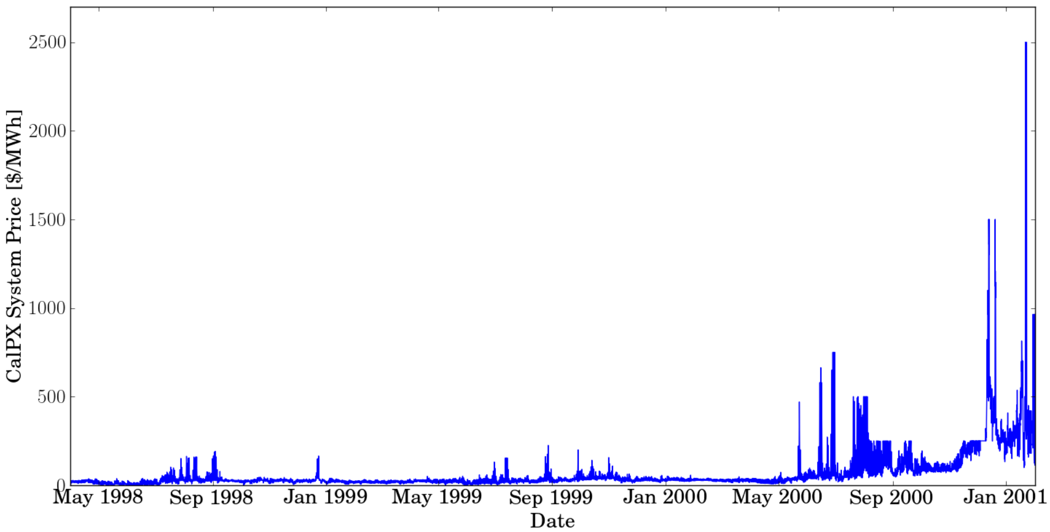

Figure 4.5: Time series of the hourly CalPX market clearing prices from Apr 1, 1998 until the market shut down on Jan 31, 2001.

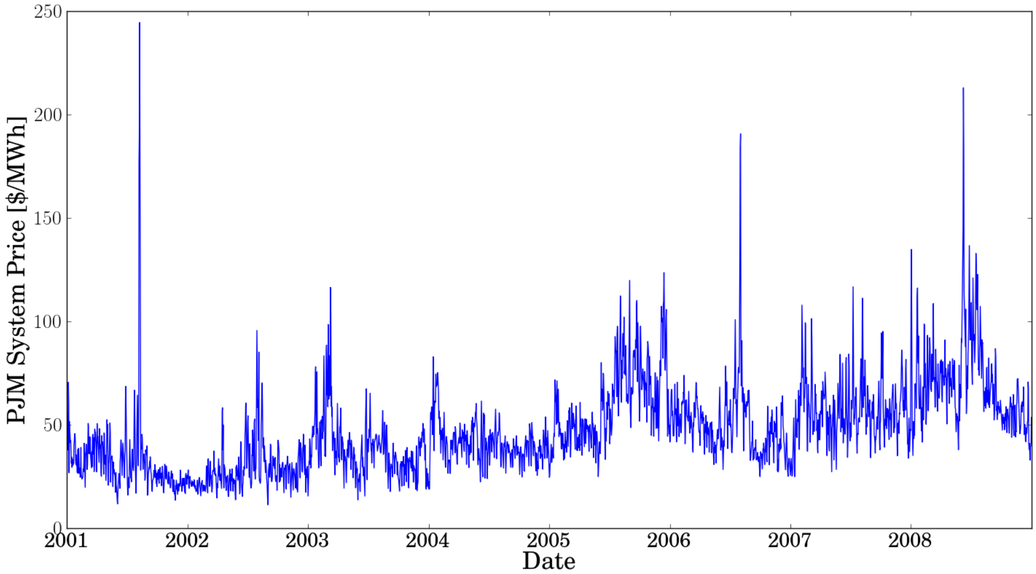

Figure 4.6: Time series of the daily PJM Interconnection market clearing prices from Jan 1, 2001 until Jan 1, 2009.

4.6 PJM Interconnection

The time series from the PJM Interconnection, given in Fig. 4.6, consists of *daily* mean prices from the day-ahead market. It spans the time period from Jan 1, 2001 until Jan 1, 2009, giving a total number of 2 922 data points.

Chapter 5

Results and Discussion

5.1 Seasonality

As already indicated, time series of electricity system prices are often periodic, both on the daily, weekly and annual time scale. Fig. 5.1 shows the power spectral density (PSD), introduced in Section 3.3.1, of the daily mean system price time series. A distinct weekly cycle is observed in all the series. Furthermore, all series except the CalPX series indicate two peculiar peaks at frequencies corresponding to cycles of 7/2 and 7/3 days. Both of these frequencies are smaller than the Nyquist frequency[1], and can therefore not be explained by aliasing effects. Instead, a likely explanation is that these two cycles are caused by the difference in demand between working days and weekend days. This effect will not be seen in only one frequency component because the weekend (two days) is shorter than the working day period (usually five days). The odd weekend-workday periodicity will therefore have contributions from several frequencies and it seems evident that $f = 2/7$ days^{-1} and $f = 3/7$ days^{-1} are the two most important contributors to these fluctuations. National holidays and other frequently occurring events will also influence the PSD. However, since there are only a few national holidays outside the weekends during the year, their contribution to the PSD is here assumed to be negligible.

The daily periodicity of electricity prices can not be observed from Fig. 5.1 because the data employed to produce these figures are of daily resolution. However, when investigating the PSDs of the *hourly* system prices, the presence of a daily periodicity is also apparent. Fortunately, when analysing system prices at time scales equal to or larger than a day, the daily periodicity is averaged out.

[1]For daily data, the Nyquist frequency is $f_N = 1/2$ days^{-1}.

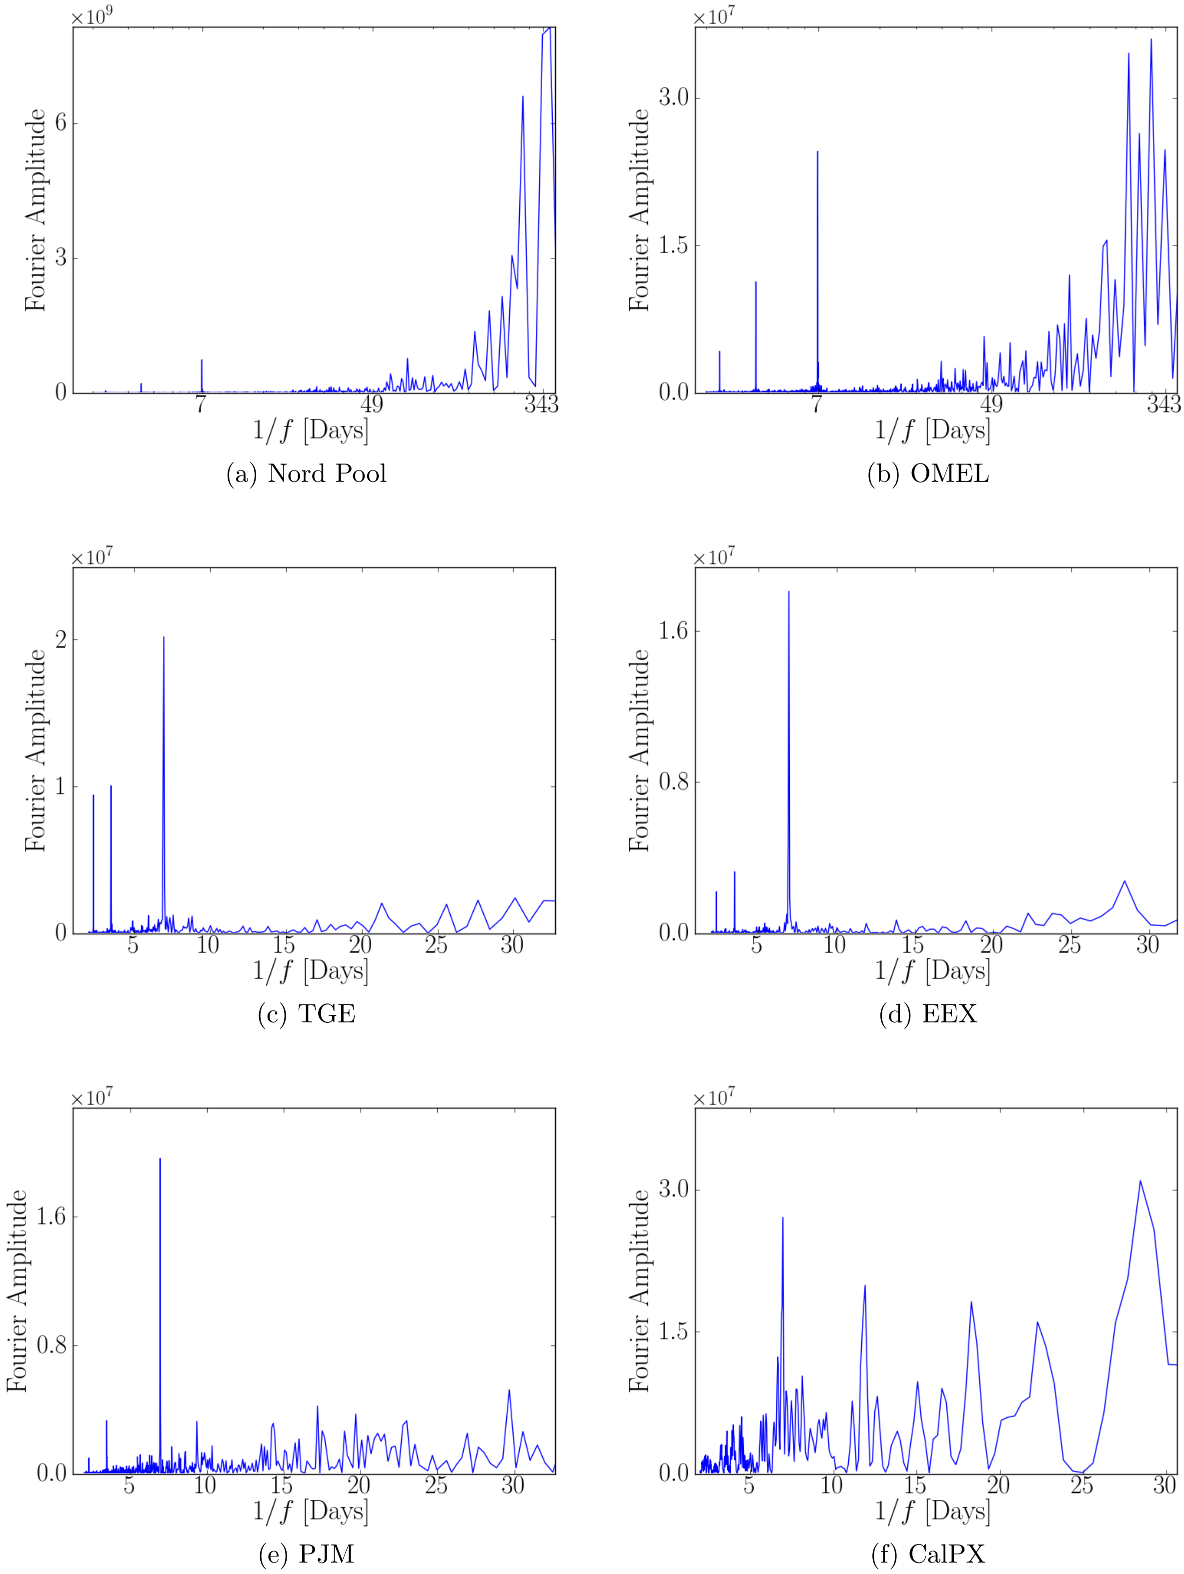

Figure 5.1: The power spectral densities of the system prices of the different electricity markets.

In the PSDs for Nord Pool and OMEL, frequencies up to the annual periodicity are included. The limited lengths of the datasets make it hard to identify one peak as the main contributor. Furthermore, as described in Section 3.3.1, the PSD is expected to increase exponentially for smaller frequencies. Still, Fig. 5.1 (a) and (b) show strong indices of a yearly cycle in Nord Pool as well as cycles with annual, 3/4 year and approximately half year periods in OMEL. The other series were not long enough to investigate annual cycles. However, as will be discussed, it seems probable that annual cycles are present in some of the markets because of the seasonal weather fluctuations.

The daily and weekly periodicity is mainly caused by human consumption patterns. Electricity is most demanded at the peak hours during the morning and afternoon. As discussed above, the demand pattern is also different on working days than on weekends and holidays. On the annual scale, there is an additional climate factor, especially the temperature, that leads to the seasonal fluctuations in demand[2]. For markets dominated by hydro-power, such as Nord Pool, there are also seasonal fluctuations in the amount of available power.

Looking past the periodic contributions

When the returns and volatilities of the system prices are measured, it is of interest to reduce the contributions from seasonal factors. The reason for this is that the seasonal fluctuations could give severe contributions to for instance the level of volatility. This could give a wrong picture of the market, since these fluctuations are more or less predictable. When comparing electricity markets to other types of commodity or finance markets, it is only fair to compare the fluctuations that can not be attributed to periodic behaviour. When calculating the *daily* return, the daily periodicity has no contribution. Furthermore, the annual periodicity can be neglected, as described in Section 2.6.2. For a precise measure of the daily returns, it is therefore the weekly periodicity that is the most important to reduce. Recall that this periodicity is caused by two factors. The first factor, namely that the demand for electricity is differently *distributed* throughout working days than throughout the weekends, is eliminated by employing the daily mean prices to calculate the daily return. The other factor, which is that the overall demand are less in the weekends, causes weekly periodicity in the daily mean prices as well.

[2]In areas with a warm climate, it is common that the demand, and hence the electricity prices, are higher during summertime, because of extensive use of air conditioning. In cold areas, such as the Nordic countries, the demand for electricity is highest during winter due to the extensive use of indoor electricity heating.

A method to remove the weekly periodicity in the daily mean prices, based on a moving median filter combined with trend removal, will now be introduced in order to compare the return levels and volatilities to the original price data.

5.1.1 Removing the Weekly Periodicity

This method is based on an assumption that the daily electricity system prices can be decomposed into three components: a time-varying trend T_t; a periodic component P_t; and a stochastic component Y_t. It is also assumed that these components combine in an additive way[3], i.e. $x_t = T_t + P_t + Y_t$ [2]. As seen in the previous section, the periodic component P_t contains several periodicities. The aim of this method is to remove the weekly component of the daily mean prices. This is accomplished in the following way.

First, a *smoothed* version of the time series of daily mean prices is subtracted from the time series itself. This will result in a time series where the time-dependent mean is removed. Since the aim of the method is to extract the weekly periodicity, it is important that the weekly periodicity itself is not included in the smoothed time series, i.e. the smoothing should be performed over time scales larger than a week. The smoothing can be performed in several ways, for example by employing a moving average. Another approach, using wavelets, is chosen here. This method for smoothing a time series is described in Appendix B. Independent of the method chosen, the subtraction of the smoothed time series should result in a time series with approximately zero mean, showing only the fluctuations at the weekly time scale and below.

Second, the resulting detrended time series is sorted into seven columns so that each column represents a day of the week and every row represents a week. Then, a median filter of length[4] $n = 13$ is applied to each column in the following way. Each element is replaced with the median value of the interval spanning from the preceding six values to the subsequent six values. This can be illustrated with an example: the value at the first column and seventh row, i.e. Monday at week 7, is replaced with the median value of the 13 Mondays from week 1 until week 13. This is performed for all days until the seventh last week of the dataset. It should be stressed that all median values are calculated before any values are replaced. Cutting off the first and last six weeks of the dataset results in a time series containing typical

[3]A different approach would be to assume that the components combine multiplicatively [2].

[4]This choice of n represents the assumption that the weekly pattern varies with the season of the year.

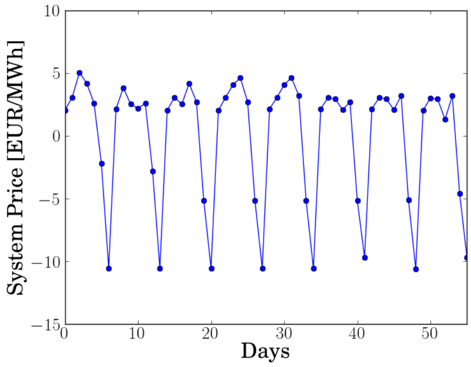

Figure 5.2: Typical weekly pattern of the Omel system price, beginning with a monday. It is seen that the weekends have lower system prices.

weekly patterns with fluctuations around an approximate zero mean. As an illustrative example, a subsequence of such a time series, using the OMEL data set, is shown in Fig. 5.2.

Finally, the time series of typical weekly patterns is subtracted from the original time series of daily mean prices. Using the OMEL time series as an example, the PSD of the original time series can be compared to the PSD of the "deseasonalised" time series, see Fig. 5.3 (a) and (b). It is seen that the weekly periodicity, including frequencies $f = 2/7$ days^{-1} and $f = 3/7$ days^{-1}, is considerably reduced. It is also seen that other characteristics, such as the annual periodicity, are left as good as untouched.

5.2 Volatility

The daily volatilities, defined in Section 2.6.3, based on respectively hourly system prices, daily mean system prices and deseasonalised daily mean system prices, are given in Table 5.1. As seen in Table 5.1, there are large differences between these three measures of the volatility. It is seen that the daily volatilities based on the hourly system prices (calculated from the 24 hour return series) are considerably higher than the daily volatilities based on the daily system prices. This may be explained by at least two factors. First of all, as previously seen, the daily consumption pattern is differently distributed on working days than in the weekends, e.g. the 24 hour return from Friday at 8 a.m. to Saturday at 8 a.m. is on average expected to be a considerable negative return value (The demand is lower on a Saturday morning than on a Friday morning). This difference in the demand distri-

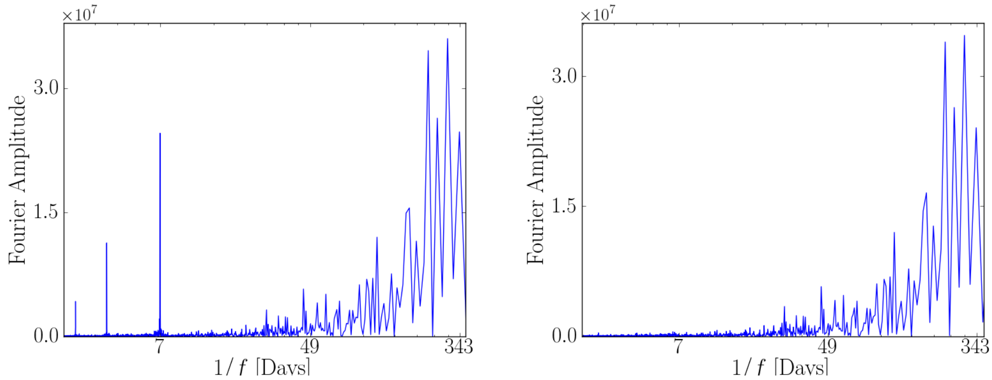

(a) Original (b) Deseasonalised

Figure 5.3: PSDs of (a) the original OMEL time series and (b) the deseasonalised OMEL time series.

Table 5.1: Volatility of the analysed time series from the different electricity markets, as well as the volatility of the closing values of the Dow Jones Index.

Market	**Daily Volatility**		
	Hourly	Daily	Deseas. daily
Nord Pool	0.17	0.12	0.09
OMEL	0.58	0.18	0.15
TGE	0.21	0.11	0.07
EEX	0.74	0.28	0.19
CalPX	0.76	0.23	0.20
PJM	n/a	0.19	0.16
Dow Jones	n/a	0.01	0.01

bution will cause fluctuations in the return series, and therefore increase the volatility. Since the distribution of the demand throughout a day is not captured in the daily mean prices, the daily mean prices will have a lower expected volatility. The second factor that causes the hourly prices to have a higher volatility than the daily prices, is that the price spikes present in the hourly price series are partly averaged out in the daily mean price series. This make the daily return values, and the volatility, smaller for the daily mean price series than for the hourly price series.

To summarize, it is the weekly periodicity caused by the difference in the demand distributions throughout working days compared to the weekends that is captured by the hourly data but not the daily data. However, as it was seen in Section 5.1, there is still another weekly periodicity in the daily price data. This periodicity is removed in the deseasonalised series of daily mean prices, which is the reason why the deseasonalised time series have less volatile daily returns than the original daily mean time series.

Whether the periodicity is taken into consideration or not, Table 5.1 shows that the analysed electricity markets all exhibit a strong volatility compared to other financial and commodity markets (see Table 2.1). The failed Californian PX is found to have the highest "deseasonalised volatility", with $\sigma_{\Delta t=1\mathrm{day}} = \mathrm{STD}(r_{\Delta t=1\mathrm{day}}) = 0.20$, while the Polish TGE is found to have the lowest deseasonalised volatility, with $\sigma_{\Delta t=1\mathrm{day}} = 0.07$, which still is seven times the volatility of the Dow Jones Index closing values.

5.3 Characteristics of the Returns

The (logarithmic) returns for the different time series are calculated as described in Section 2.6.2. The daily returns are chosen to be investigated because, as already argued, then the daily periodicity are phased out. The daily returns are calculated both for the hourly price data (with $\Delta t = 24\mathrm{h}$) and, for comparison, for the daily mean price data. Furthermore, the daily returns of the deseasonalised data are calculated.

In order to have a measure of how large a return value is compared to the "normal" fluctuations of the returns, the return series are normalized with their standard deviations[5]. The normalized return series based on the original daily mean system price series are plotted for the different markets in Fig. 5.4. The normalized return series based on the hourly data and the deseasonalised daily data look quantitatively similar to the series based on the daily mean data. However, when the ACFs, introduced in Section 2.4.2, of the daily returns based on the daily data are compared to those of the deseasonalised daily data, Fig. 5.5, the periodicity becomes apparent. While the periodic daily mean prices yield an apparent correlation at time lags 7 days, 14 days, 21 days, etc., the deseasonalised daily mean prices seem to have an autocorrelation with lags up to between two and 9 days.

Another feature that can be seen in Fig. 5.5, is that the daily returns of the daily system prices (both original and the deseasonalised) are negatively correlated to the previous day or the day before that. This is an indication of mean-reversion, which will be studied in detail in Section 5.5.

Empirical evidence has long suggested that the returns of financial assets are

[5]Which equal the daily volatilities of the price series and can thus be found in Table 5.1.

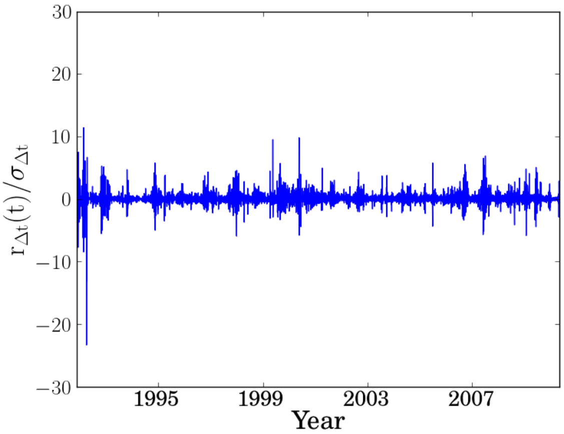

(a) Nord Pool

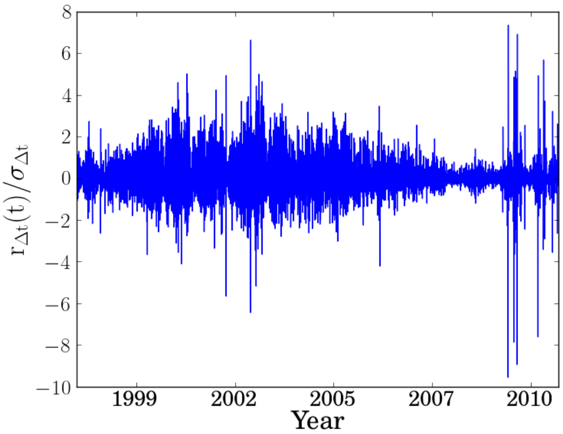

(b) OMEL

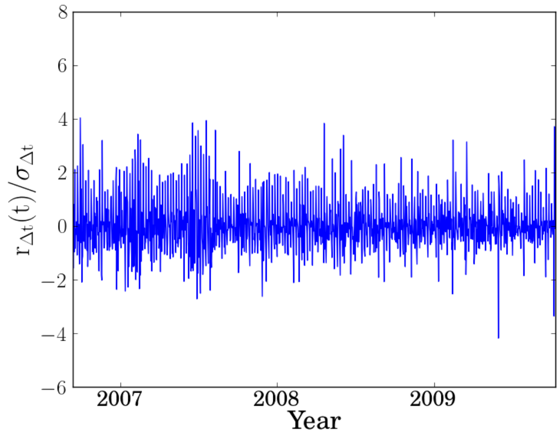

(c) EEX

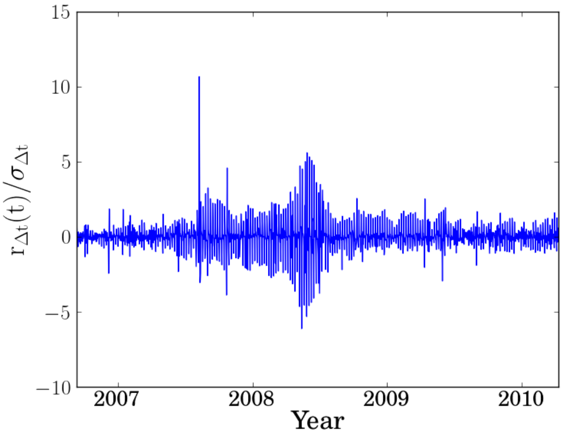

(d) TGE

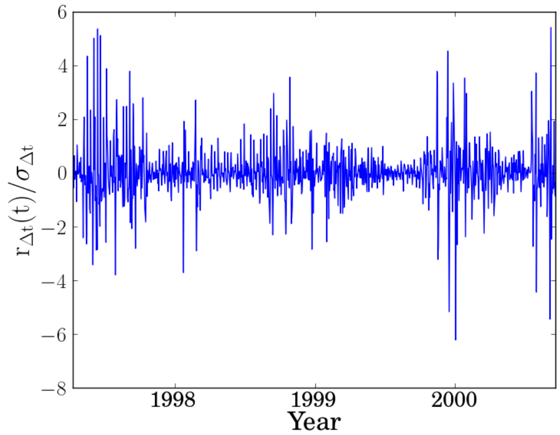

(e) CalPX

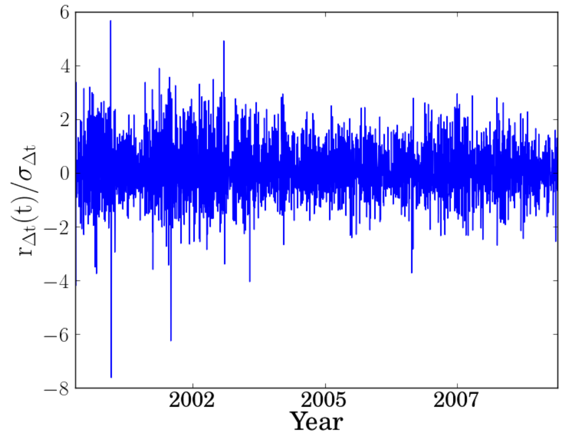

(f) PJM

Figure 5.4: Normalized logarithmic returns based on the daily mean prices with $\Delta t = 1$ day.

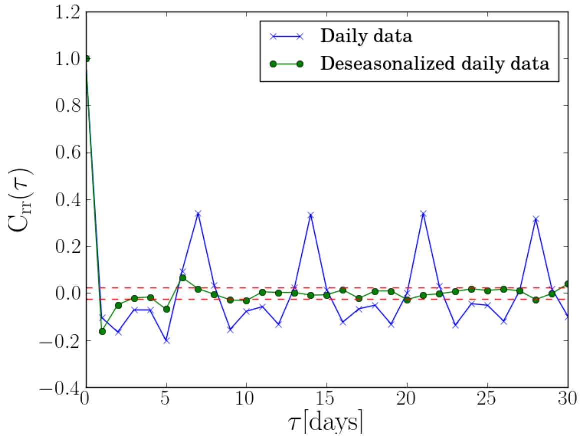

(a) Nord Pool

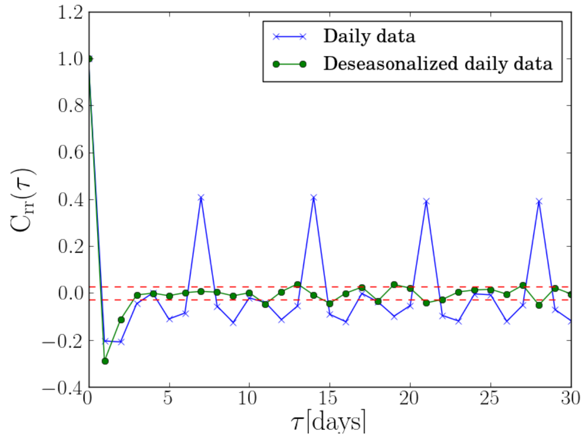

(b) OMEL

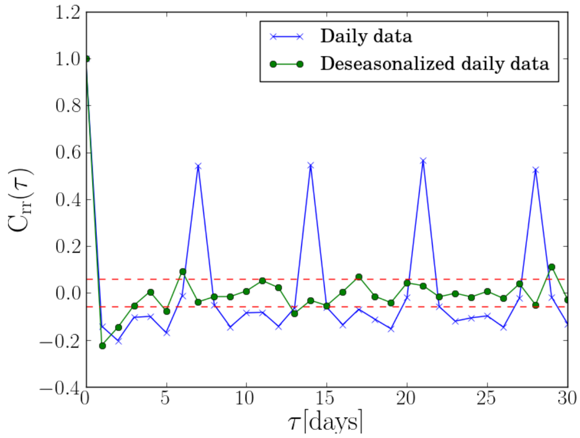

(c) EEX

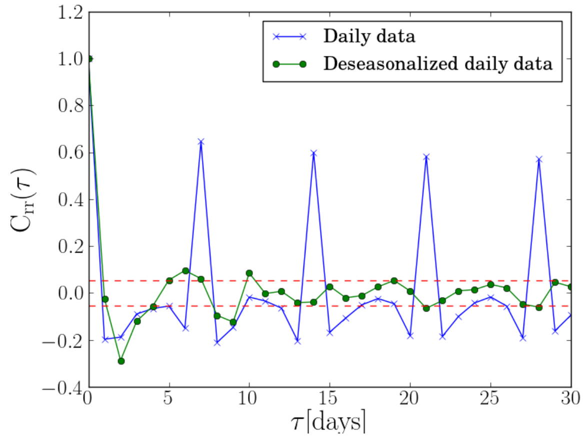

(d) TGE

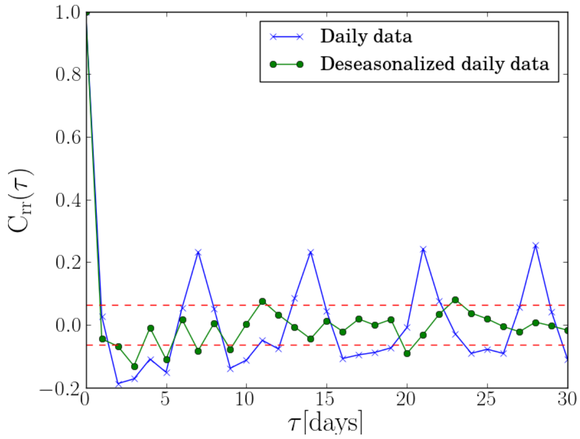

(e) CalPX

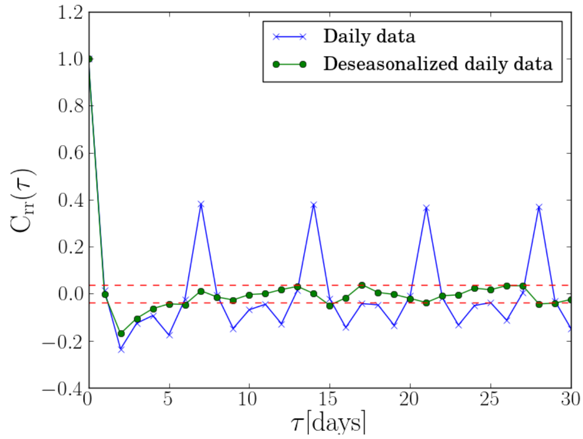

(f) PJM

Figure 5.5: ACFs of the normalized, logarithmic daily returns for daily data and deseasonalised daily data.

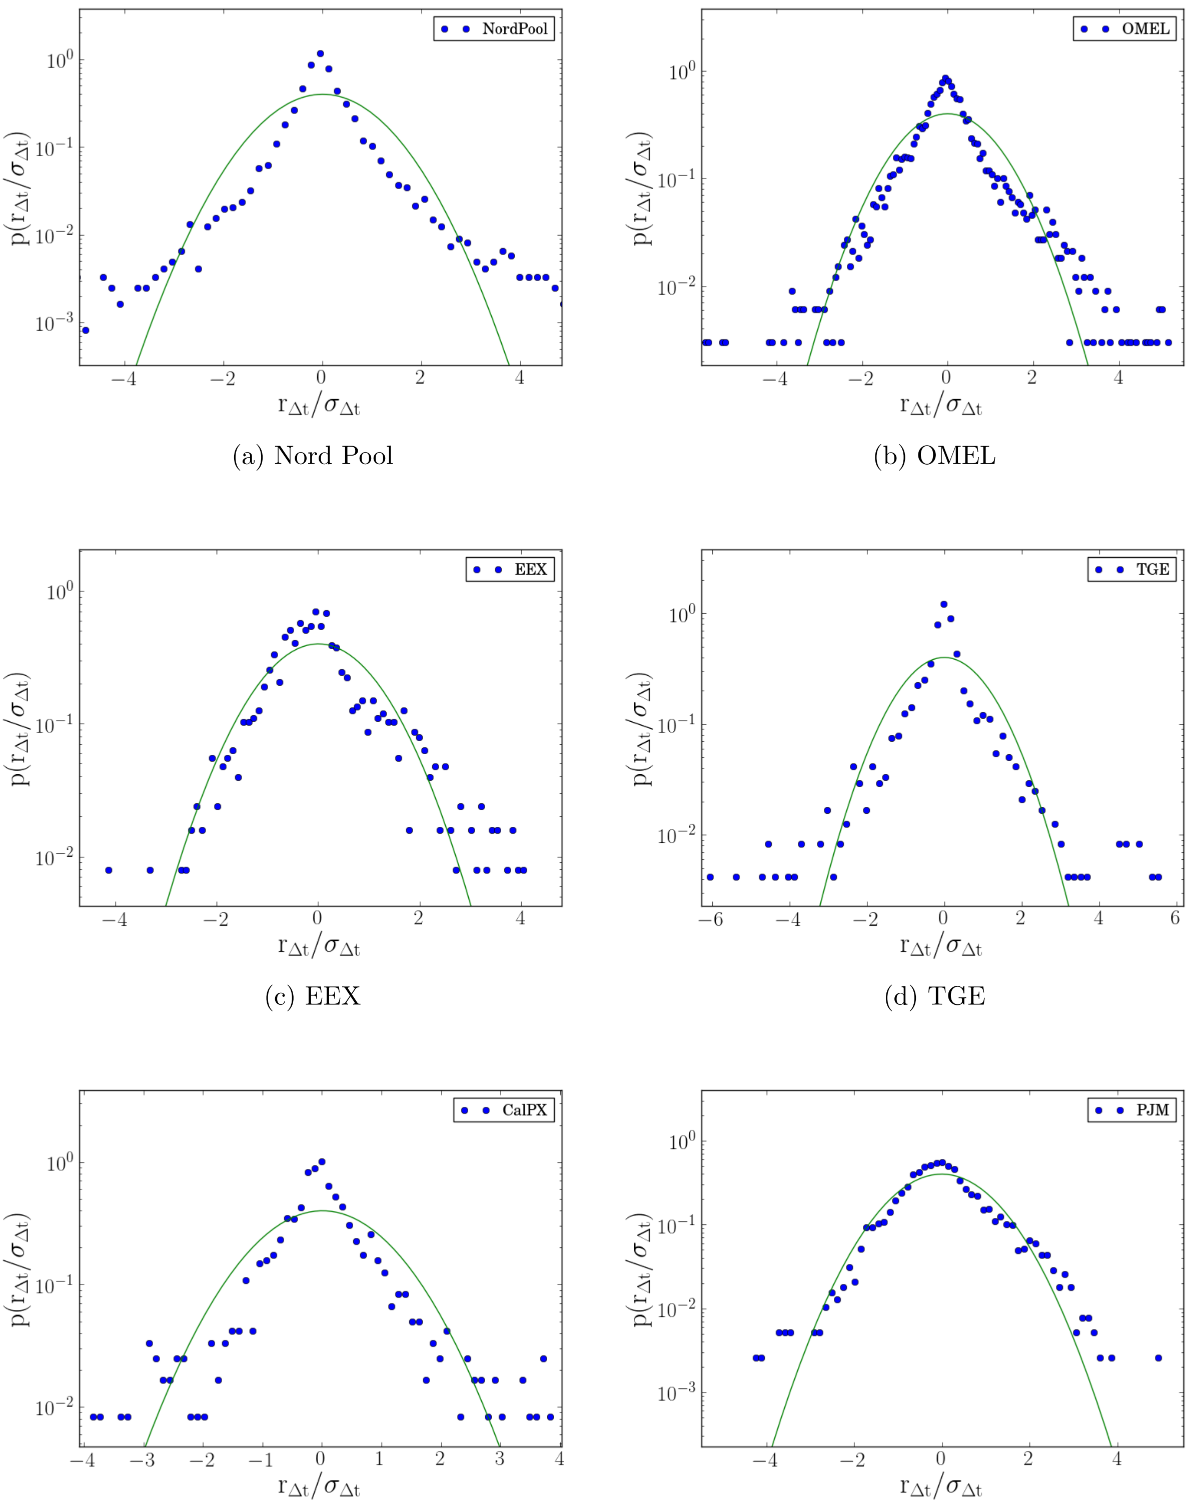

(a) Nord Pool

(b) OMEL

(c) EEX

(d) TGE

(e) CalPX

(f) PJM

Figure 5.6: Probability density functions of the logarithmic returns. The superimposed lines are Gaussian distributions with equal mean and standard deviation as the empirical data.

not normally distributed [2]. Looking at the normalized daily return distributions of the daily mean system prices, see Fig. 5.6, this is evident also for electricity system prices. Table 5.2 summarizes the kurtosis and skewness[6] of the distributions, including the daily return distributions based on the hourly and deseasonalised data. In summary, all the distributions are found to be leptokurtic, meaning that they have excess kurtosis, and are characterised by fat tails. This indicates that large daily variations in electricity prices are more common than what one could expect from a process where the returns were normally distributed[7]. This result is in agreement with earlier reports [2].

Table 5.2: Kurtosis and skewness for the probability distributions of logarithmic daily returns in the respective markets. The resolution of the data employed in the calculations is indicated as "Hourly" or "Daily" data. The columns labelled "Deseas. daily" indicate that any weekly periodicity of the daily mean time series is removed before the time series of returns are created.

Market	**Kurtosis**			**Skewness**		
	Hourly	Daily	Deseas. daily	Hourly	Daily	Deseas. daily
Nord Pool	440.2	87.5	49.8	-1.10	0.03	-0.32
OMEL	100.8	9.04	46.2	0.00	0.11	-0.85
TGE	1640.8	15.4	76.8	0.76	0.86	3.65
EEX	-	-	-	-	-	-
CalPX	83.5	8.07	17.9	0.26	0.25	-0.37
PJM	-	3.04	7.20	-	0.18	-0.43
Dow Jones	-	24.8	23.1	-	-0.60	-0.52

Volatility Clustering

The daily returns, shown for the daily data in Fig. 5.4, indicate that the volatility varies with time and that the series may be categorized into *regimes* with different volatility. It seems that there are periods when the returns fluctuate strongly, while other periods may be calm with only small fluctuations. This phenomenon is termed *volatility clustering* and, as Simonsen, Weron and Mo [6] points out, it is reminiscent of the intermittent patterns often found in time series of the velocities of turbulent fluids [45, 46]. In order to take a closer look on the volatility clustering, an estimate of the time-dependent daily volatility should be investigated.

[6]Defined in Appendix A.
[7]Gaussian distributions have a kurtosis of 0.

The time-dependent volatility could be estimated in several ways, but the simplest and most intuitive choice would be to use a time interval with length T, perfectly embracing the point of evaluation, as a basis to calculate the time-dependent volatility as

$$\begin{aligned}\hat{\sigma}_{\Delta t}(t,T) &= \mathrm{STD}\Big[r_{\Delta t}(\tau)\Big], \qquad \tau \in \Big[t - T/2, t + T/2\Big] \\ &= \left(\frac{1}{N-1}\sum_{\tau=t-\frac{T}{2}}^{t+\frac{T}{2}}\Big(r_{\Delta t}(\tau) - \overline{r}_{\Delta t}\Big)^2\right)^{1/2}. \end{aligned} \tag{5.1}$$

Here, $\overline{r}_{\Delta t}$ denotes the sample mean of the returns in the interval $\left[t-\frac{T}{2}, t+\frac{T}{2}\right]$ and N is the number of data points in the same interval. Fig. 5.7 displays the results of employing (5.1) with $T = 100$ days as a moving filter on the normalized daily returns series of Fig. 5.4. It is seen that the time-dependent volatilities are varying quite drastically in time. For both Nord Pool, OMEL, TGE and CalPX, the maximum volatility levels are about five times the minimum volatility levels. The variations in volatility at the PJM and EEX markets seem to be at a more moderate level.

It should be mentioned that there exist more sophisticated ways of estimating the time-dependent daily volatility. An example of this would be to employ an exponentially weighted moving average (EWMA) instead of the non-weighted approach of (5.1). Such a weighted approach would put more emphasis on the values close to the point of evaluation. However, for the purpose of being able to qualitatively conclude that the volatility is time-varying, the non-weighted approach given here is considered adequate.

There are many factors that cause the volatility to change. For instance would non-stability in the production units cause more volatile system prices. Furthermore, it is not hard to imagine that certain seasons of the year have large fluctuations in the temperature, which again would lead to higher volatility. In Fig. 5.8 the time-dependent volatilities calculated with $T = 30$ days in (5.1) are plotted as functions of the mean system price of the same $T = 30$ days. In order to give a clearer picture of any trend, median filters of length $N = 25$ are employed on the scattered data, as described by Simonsen [47]. To better visualize any trend, either a linear or an exponential curve is fitted to the median filtered data. It is seen that the volatility at the Nord Pool market tends to be at a higher level when the system price is low, but is not dependent on the system price when the prices are above a certain level. This effect is also seen weakly in the OMEL data, but is not apparent for the other markets.

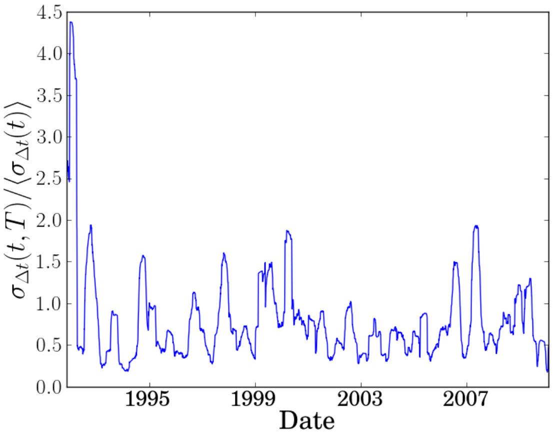

(a) Nord Pool

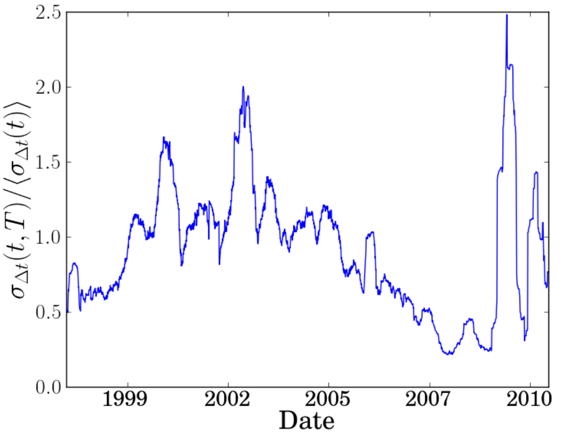

(b) OMEL

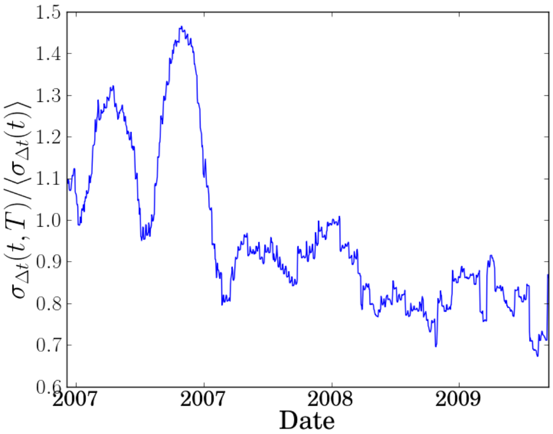

(c) EEX

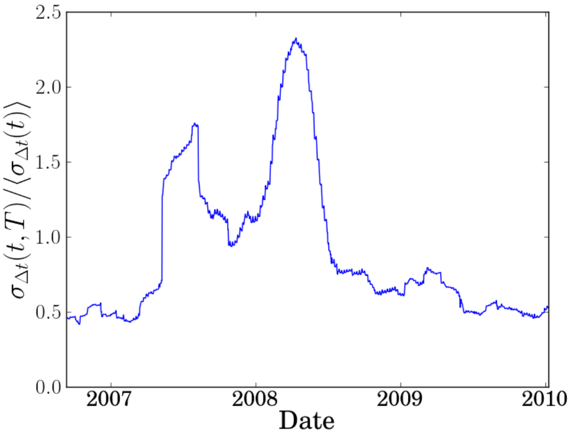

(d) TGE

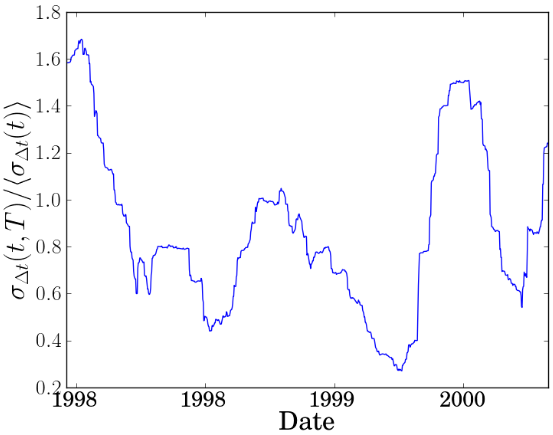

(e) CalPX

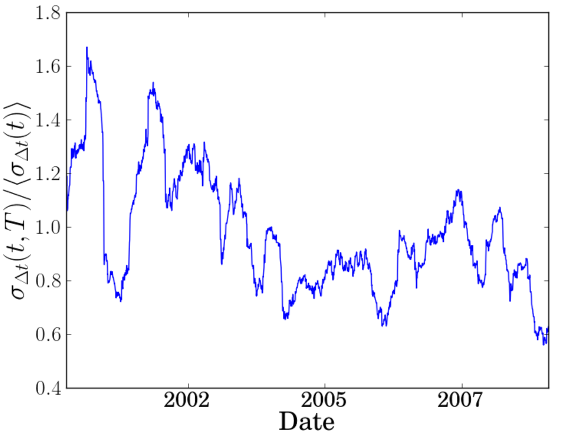

(f) PJM

Figure 5.7: The normalized time-dependent volatility of the daily mean prices, defined in (5.1) with $T = 100$ days. The normalization factor is the overall volatility given in Table 5.1.

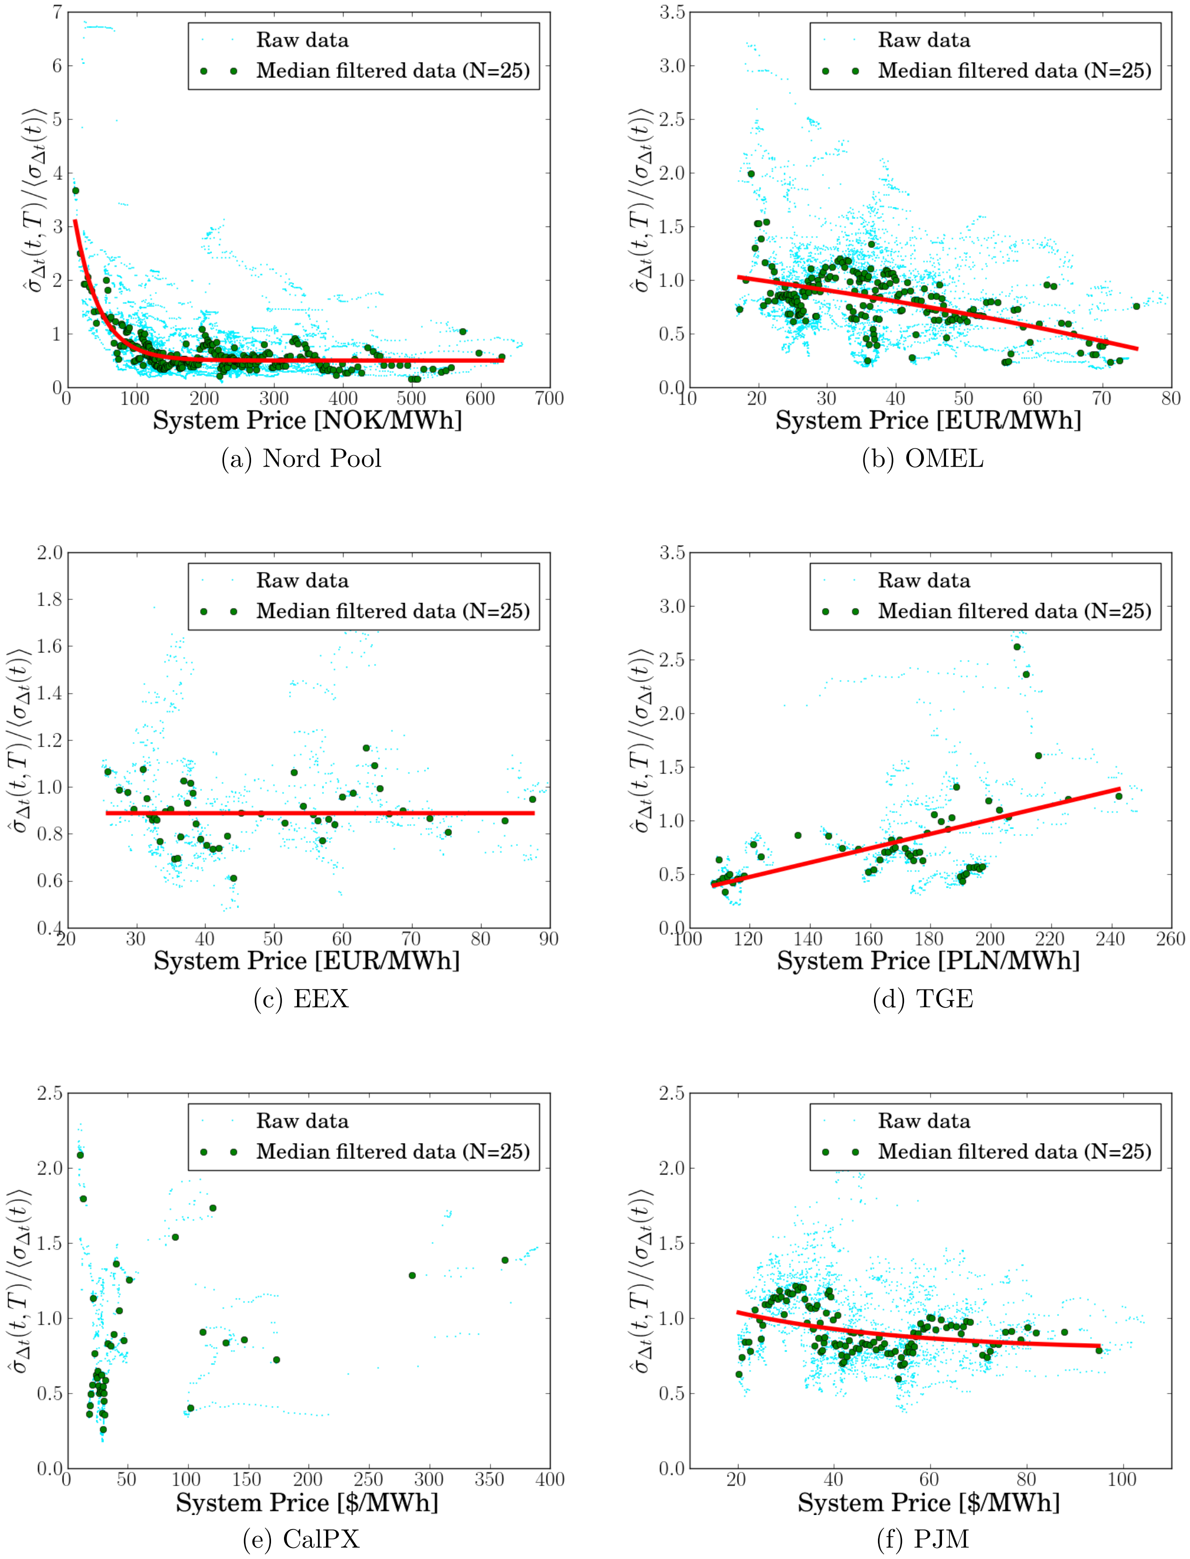

(a) Nord Pool (b) OMEL

(c) EEX (d) TGE

(e) CalPX (f) PJM

Figure 5.8: Scatter plot of the time dependent volatility (5.1) versus the corresponding mean system price. The circles are the result of applying a median filter of length N=25 to the raw data, as described in the text. The solid, red lines are either exponential or linear curves meant only as a guide to the eye.

It is not entirely clear what causes this effect. At first, one could suspect that the weekly seasonality could be one of the reasons. However, the deseasonalised data yields qualitatively the exact same results. Simonsen [47] speculates that the effect might origin from what he refers to as *forced production*. In seasons with low demand, which often is reflected in low system prices, some production companies still need to produce a certain amount of electric energy. In the Nordic countries, for example, overfull water reservoirs will force the production company to either produce electricity or to drain water without producing electricity. The profit-optimizing option will then often be to produce more electricity even if the demand, and price, is low. Then, when the filling fraction of the reservoir again is within the safety limit, the profit-optimizing option could be to wait for times of higher system prices. This leads to a higher volatility, as is seen in Fig. 5.8. It is not clear what causes the higher volatility in times of low system prices in the Spanish OMEL market. An interesting observation is also the absence of the effect described above for the other markets. In fact, for the Polish TGE market the opposite effect is apparently observed. However, it is questionable whether the size of the data set is large enough to draw any conclusions about this[8].

5.4 Spikes

Using the spike definition of Section 2.6.1, an intuitive measure of the "spikedness" of the electricity markets is the frequency of occurring spike events, e.g. the number of spikes occurring on average per year. The results are given in Table 5.3. Since the duration of a price spike is usually a few hours, only the hourly time series are investigated here. Furthermore, since the spiky nature of Nord Pool apparently changed around 1998-2000[9], it is chosen to use the 11 year period from year 2000 until the end of 2010 as basis for the investigation. Consequently, the same period of time is chosen from the OMEL data. For all datasets, only full years are taken into consideration. It is seen that Nord Pool and OMEL had approximately 10 spikes per year, while the other three markets seem to have a more spiky nature.

Weron [2] argues that price spikes are a consequence of the bidding structure. Recall from Section 2.2.4 that in uniform price auctions, market participants do not necessarily have to pay the price they bid. As long as their bids are above the system price, they can buy electricity for the system

[8]For the TGE dataset, there are 47 independent periods of $T = 30$ days.

[9]This was most likely caused by the market entrance of Finland in 1998 and Denmark in 2000.

Table 5.3: The number of spikes occurring in the hourly data on average per year.

Market (year)	Spikes [#/year]
Nord Pool (2000-2010)	10.4
OMEL (2000-2010)	9.4
TGE (2007-2009)	75.0
EEX (2007-2009)	44.7
CalPX (1999-2000)	71.0

price. Since electricity is a necessity for many of the market participants, they regularly place bids at the maximum allowed price level. Furthermore, the marginal production cost, i.e. the cost of producing one more unit of electricity, is growing rapidly at high production levels [6]. In times of high demand, e.g. peak hours during extreme cold winter days, the offers of the more expensive production utilities are accepted at the auction and the system price gets higher. Of course, the power generation companies also know that some buyers place bids at high price levels. If they anticipate that the demand at some time will be high, the profit optimizing option for them could be to try to force increases in the price by placing offers at a high price level. This kind of speculation is probably most common in generation companies with flexible generators. For other generators, the price of shutting down the production may be so high that they do not speculate in forcing the prices up, since this would expose them to expensive risk of being forced to reduce the production.

An interesting issue related to the spikes are the abuse of market power to force the prices up. If a generation company with a large market share anticipate high demand, then they may single-handedly be able to force an increase in the prices. In Section 2.3.3, it was mentioned that EEX may have a lack of competition since the two largest production companies have a combined market share of 55%. This could be one of the reasons why EEX are more spiky than for example Nord Pool. However, the underlying mechanisms of price spikes are probably complex. It is reasonable to believe that other factors, such as the amount of available production capacity and the type of power plants available, are also affecting the number of price spikes in a market.

Table 5.4: Estimates of the Hurst exponent using the AWC and the DFA methods. The level of significance at which a null hypothesis of H=0.5 can be rejected for the DFA estimates are put in parenthesis. The $\overline{\text{RMS}}(n)$ (3.4) of the PJM data did not obey a power law and the DFA could therefore not estimate H for this series.

Method	**AWC**	**DFA (daily data)**	
Market	Hourly*	Daily	Deseasonalised daily
Nord Pool	0.43	0.38 (99%)	0.43 (90%)
OMEL	0.24	0.24 (99%)	0.29 (99%)
TGE	0.26	0.21 (99%)	0.35 (99%)
EEX	0.23	0.13 (99%)	0.22 (99%)
CalPX	0.33	0.32 (99%)	0.36 (95%)
PJM (daily)	0.02	-	-
Dow Jones	0.57	0.52 (<90%)	0.52 (<90%)

*Hourly data is employed, but only time scales corresponding to above one week are used as basis for the estimation of the Hurst exponent, see Fig. 5.9.

5.5 Hurst Analysis

The Hurst exponents for the increments of the different time series are estimated with the DFA and AWC methods, described in Section 3.2 and 3.3.4, respectively. The results are given in Table 5.4. In addition, the wavelet spectra of the AWC method are presented in Fig. 5.9. The regression lines that estimated the Hurst exponents, found by weighted linear least squares fits to the logarithmic values, are indicated in the figure. For the hourly price data, a crossover is observed between the intra-day time scale and the time scale above a day. Furthermore, another crossover corresponding to the weekly time scale is observed in all datasets, but this crossover is less pronounced. These crossovers are most likely caused by the periodicities in the hourly data, and demonstrate the reason why daily data, and not hourly, are employed in the DFA analysis. The AWC estimates of H, given in Table 5.4, are valid for the scaling regimes corresponding to time scales larger than a week.

As can be seen by the differences between the AWC and DFA Hurst estimates in Table 5.4, determination of the Hurst exponent of a time series is a non-trivial business. For many purposes, however, it is sufficient to establish whether H is smaller than, equal to or bigger than 0.5. By using the empir-

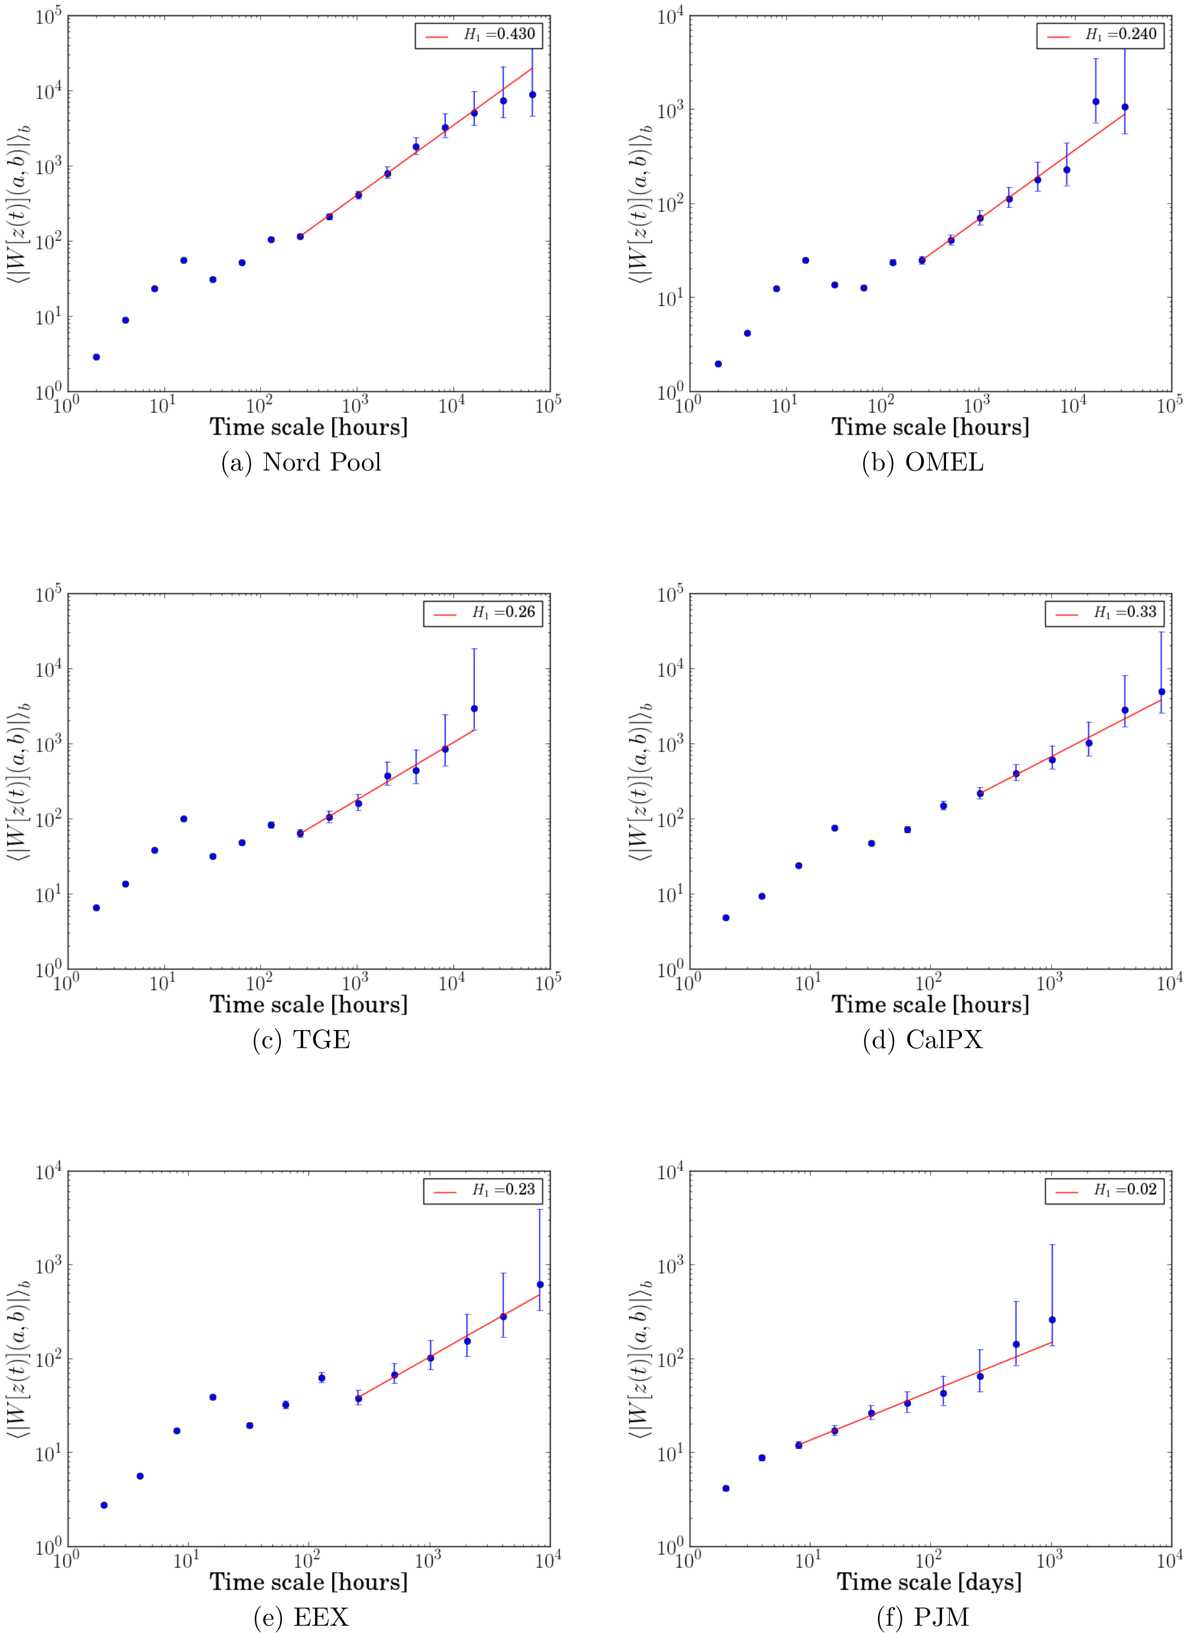

(a) Nord Pool

(b) OMEL

(c) TGE

(d) CalPX

(e) EEX

(f) PJM

Figure 5.9: Wavelet spectra and estimation of the Hurst exponent.

ical confidence intervals for $H = 0.5$, which Weron [41] established for the DFA, the null hypothesis of $H = 0.5$ may be tested at 90%, 95% and 99% significance levels. Employing DFA to the daily mean system price increments reveals that all[10] of the analysed series have a Hurst exponent significantly less than $H = 0.5$, meaning that the system prices are anti-persistent. When analysing the deseasonalised daily system prices the same conclusion can be drawn, but the significance levels for the Nord Pool and CalPX data sets are then reduced to 90% and 95%, respectively. For comparison, it is also seen that the Hurst exponent of the Dow Jones data set can not be distinguished from $H = 0.5$ using this DFA method.

[10]The PJM series did not reveal any power law scaling of the $\overline{\mathrm{RMS}}(\mathrm{n})$ of (3.4), and H for this series could hence not be estimated using DFA.

Chapter 6

Conclusions

In this thesis, the market structures of six electricity markets have been studied. These markets are the Nord Pool (the Nordic countries), OMEL/MIBEL (Spain), EEX (Germany), TGE (Poland), CalPX (California, US) and PJM (the American East) markets. Time series of the corresponding system prices for electricity have been analysed, in order to identify characteristics of electricity prices and differences between the markets.

It has been observed that time series of electricity system prices exhibit periodic patterns both on the daily, weekly and annual time scale. These periodicities can affect other measures of the time series. It has been argued how the daily periodicity can be eliminated and in which cases it is reasonable to neglect the annual seasonality. Furthermore, a method to remove the weekly periodicity in daily mean electricity prices has been employed. Based on the power spectral densities, it is seen that the method succeeds with removing the weekly periodicity. To what extent the method also removes other characteristics should be investigated further in future works.

The daily and weekly periodicities amount to a considerable amount of the daily volatilities. Still, the daily volatilities of the deseasonalised data have all been found to be higher than the daily volatilities seen in other types of commodity markets. Hence, it is affirmed that electricity system prices are very volatile. Furthermore, it is concluded that the high level of volatility can not be attributed to the periodic fluctuations alone.

The logarithmic return series exhibited intermittent patterns, which is an indication of volatility clustering. By employing a time-dependent measure of the daily volatility, it has been observed that the level of volatility is fluctuating. In particular, for Nord Pool and OMEL, it has been found that the volatilities are high in times of low system prices. It was concluded that this effect can not be attributed to either the daily nor weekly periodicity.

It has been concluded that, for all the markets investigated, the probability density functions of the logarithmic returns are leptokurtic.

The hourly system prices from TGE, EEX and CalPX have been found to be more spiky than the series from Nord Pool and OMEL. Still, spikes are occurring frequently in all of the analysed time series.

Finally, two different methods of measuring the long-term correlations in the increments of the system prices have been employed. It was concluded that the series from all the electricity markets are anti-persistent, i.e. the Hurst exponents of the increments were less than 0.5. Furthermore, the AWC method revealed that there were up to three different scaling regimes in the hourly system prices.

Further Work

First of all, different ways of removing the weekly periodicity, and their impact on for instance the volatility, could be compared to the method employed in this work. Also, other measures of the spikes could be considered. Furthermore, in this report it was discussed how the amount of flexible generators could influence some of the characteristics, such as the number of spikes. It would have been interesting to investigate the market share of different power plants in the different markets, and compare the findings to the mentioned discussion in this report.

When it comes to the Hurst analysis, the impact of the three different scaling regimes observed with the AWC method could be investigated. It is also believed that there is a potential in further development of the AWC method. Instead of measuring the fluctuations at the different time scales with the absolute mean of the wavelet coefficients, the fluctuations could be quantified by the median value of the wavelet coefficients at the different time scales. This approach would decrease the influence of the price spikes, which can be desired when investigating the influence of the more moderate fluctuations of the time series.

It is desirable to establish statistical distributions for the absolute mean of the wavelet coefficients. Developing confidence intervals for the Hurst estimates of the AWC method would increase the utility value of the method. If this is not possible, a Monte Carlo study of the AWC estimates could be performed in a similar fashion as Weron did for the DFA estimates.

Finally, it has been briefly mentioned that indications were found, during the work on this thesis, that the electricity system prices were multi-affine. This property could be investigated further, for instance by means of the wavelet transform or by employing generalized versions of the DFA.

Bibliography

[1] Newbery DM. Problems of liberalising the electricity industry. *European Economic Review*. 2002;46(4):919-927.

[2] Weron R. *Modeling and forecasting electricity loads and prices, A statistical approach*, Chichester: John Wileys & Sons Ltd.; 2006.

[3] Oren S. Capacity payments and Supply Adequacy in Competitive Electricity Markets. *Proceedings of the VII Symposium of Specialists in Electric Operations and Expansion Planning (SEPOPE VII)*. Curitiba, Brazil, May 21-26 (2000).

[4] Web pages of Nord Pool Spot AS [Internet]. Visited May 20, 2011. Available from: `www.nordpoolspot.com`.

[5] Allen WB, Weigelt K, Doherty N, Mansfield E. *Managerial Economics: Theory, Applications, and Cases*, 7th edition. New York: W. W. Norton & Company; 2009.

[6] Simonsen I, Weron R, Mo B. Structure and stylized facts of a deregulated power market. Unpublished paper. Available from: `http://mpra.ub.uni-muenchen.de/1443/`.

[7] Web pages of Compania Operadora del Mercado Español de Electricidad (OMEL). Visited May 20, 2011. Available from: `www.omel.com/en`.

[8] Web pages of the European Energy Exchange (EEX). Visited May 20, 2011. Available from: `www.eex.com`.

[9] Web pages of the European Power Exchange (EPEX). Visited May 20, 2011. Available from: `www.epexspot.com`.

[10] Viehmann J. Risk premiums in the German day-ahead Electricity Market. *Energy Policy*. 2011;39(1):386-394.

[11] Rademaekers K, Slingenberg A, Morsy S. Review and analysis of EU wholesale energy markets – Historical and current data analysis of EU wholesale electricity, gas and CO2 markets. Final Report prepared for the European Commission DG TREN, Rotterdam, 2008.

[12] Weigt H, von Hirschhausen C. Price formation and market power in the German wholesale electricity market in 2006. *Energy Policy*. 2008;36(11):4227-4234.

[13] Web pages of the Towarowa Gielda Energii (TGE). Visited May 20, 2011. Available from: `www.tge.pl/en`.

[14] Web pages of the PJM Interconnection. Visited May 20, 2011. Available from: `www.pjm.com`.

[15] Joskow PL. California's electricity crisis. *Oxford Review of Economic Policy*. 2001;17(3):365-388.

[16] Free Statistics Software, Office for Research Development and Education, version 1.1.23-r6 [Internet]. Wessa P. Visited May 5, 2011. Available from: `http://www.wessa.net/`.

[17] Wei WWS. *Time Series Analysis*. 2nd ed. Boston: Pearson Education, Addison-Wesley; 2006.

[18] Bradley JN, Brislawn CM and Hopper T. FBI wavelet/scalar quantization standard for gray-scale fingerprint image compression. *Visual Information Processing II*. 1993;1961(1);293-304.

[19] Graps A. An introduction to wavelets. *Computational Science Engineering, IEEE*. 1995;2(2):50-61.

[20] Jawerth B and Sweldens W. An Overview of Wavelet Based Multiresolution Analyses. *SIAM Review*. 1994;36(3):377-412.

[21] Press WH, Teukolsky SA, Vetterling WT and Flannery BP. *Numerical Recipes in C* 2nd ed. Cambridge: Cambridge University Press; 1992. pp.591-606.

[22] Combes JM, Grossmann A and Tchamitchian P. *Wavelets*. Berlin: Springer-Verlag; 1989.

[23] Shensa MJ. The discrete wavelet transform: wedding the a trous and Mallat algorithms. *Signal Processing, IEEE Transactions on*. 1992;40(10):2464-2482.

[24] Trueck S, Weron R and Wolff R. Outlier Treatment and Robust Approaches for Modeling Electricity Spot Prices. *Paper presented at the 56th Session of the International Statistical Institute, Invited Paper Meeting IPM71. Statistics of risk aversion*, Lisbon, Portugal, Aug 2007.

[25] Lapuerta C, Moselle B, 2001. Recommendations for the Dutch electricity market. Report from *The Brattle Group, Ltd.* London. Available from: `http://www.brattle.com/_documents/UploadLibrary/Upload338.pdf`.

[26] Bierbrauer M, Trück S, Weron R. Modeling electricity prices with regime switching models. *Lecture Notes in Computer Science.* 2004;3039:859-867.

[27] Clewlow L, Strickland C. *Energy Derivatives - Pricing and Risk Management.* London: Lacima Publications; 2000.

[28] Borovkova S, Permana F. Modelling electricity prices by the potential jumpdiffusion. *Proceedings of the Stochas- tic Finance.* Lisbon, Portugal, 2004.

[29] Stevenson M. Filtering and Forecasting Spot Electricity Prices in the Increasingly Deregulated Australian Electricity Market. Research Paper, Quantitative Finance Research Centre, University of Technology, Sydney. 2001;vol. 63.

[30] Hull JC. *Options, futures, and other derivatives 4th ed.*, New York: Prentice Hall International; 2000.

[31] Hurst H.E. Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers.* 1951;116:770-808.

[32] Weron R, Przybylowicz B. Hurst analysis of electricity price dynamics. *Physica A.* 2000;283(3-4);462-468.

[33] Feder J. *Fractals*, New York: Plenum Press; 1988.

[34] Mandelbrot BB and van Ness JW. Fractional Brownian Motions, Fractional Noises and Applications. *SIAM Review.* 1968;10(4):422-37.

[35] Percival DB, Walden AT. *Wavelet Methods for Time Series Analysis.* New York: Cambridge University Press; 2000.

[36] Bassingthwaighte JB and Raymond GM. Evaluating Rescaled Range Analysis for Time Series. *Annals of Biomedical Engineering.* 1994;22(4):432-444.

[37] Simonsen I, Hansen A, Nes O. Determination of the Hurst exponent by use of wavelet transforms. *Phys. Rev. E.* 1998;58(3):2779.

[38] Mielniczuk J and Wojdyllo P. Estimation of Hurst exponent Revisited. *Computational Statistics & Data Analysis.* 2007;51(9):4510-4525.

[39] Taqqu MS, Teverovsky V and Willinger W. Estimators for long-range dependence: an empirical study. *Fractals.* 1995;3(4):785-788.

[40] Peng C-K, Buldyrev SV, Havlin S, Simons M, Stanley HE, Goldberger AL. Mosaic Organization of DNA nucleotides. *Physical Review E.* 1994;49(2):1685.

[41] Weron R. Estimating long-range dependence: Finite sample properties and confidence intervals. *Physica A.* 2002;312(1-2):285-299.

[42] Barabasi A-L, Szepfalusy P, Vicsek, T. Multifractal spectra of multi-affine functions. *Physica A.* 1991;178(1):17-28.

[43] Uhm W and Kim S. Discrete Wavelet Analysis of Multifractal Measures and Multi-Affine Signals. *Journal of the Korean Physical Society.* 1997;32(1):1-7.

[44] Der Spiegel Online. *Germany to Phase Out Nuclear Power by 2022.* Web page visited May 30, 2011. Available from: `http://www.spiegel.de/international/germany/0,1518,765594,00.html`.

[45] Mantegna RN and Stanley HE. *An Introduction to Econophysics: Correlations and Complexity in Finance,* Cambridge: Cambridge University Press; 2000.

[46] Frisch U. *Turbulence: The Legacy of A. Kolmogorov,* Cambridge: Cambridge University Press; 1997.

[47] Simonsen I. Volatility of power markets. *Physica A.* 2005:355(1);10-20.

[48] Forbes C, Evans M, Hastings N and Peacock B. *Statistical Distributions.* 4th ed. Hoboken, New Jersey: Wiley. 2011.

[49] Kreyszig E. *Advanced Engineering Mathematics.* 9th edition. Hoboken, New Jersey: Wiley. 2006.

[50] More JJ. The Levenberg-Marquardt algorithm: Implementation and theory. *Numerical Analysis. Lecture Notes in Mathematics.* 1978;630:105-116.

This bibliography is written in the Vancouver style.

Appendices

Appendix A

Basic Statistics and Probability Theory

A.1 The Probability Density Function

The probability density function (PDF) $f(x)$, also termed *probability distribution* or simply *distribution* of a continuous random variable X, is implicitly defined through the integral

$$F(x) = P(X \leq x) = \int_{-\infty}^{x} f(x^*)\mathrm{d}x^*. \tag{A.1}$$

Here, P denotes the probability of realizing whatever is inside the brackets. The PDF can explicitly be defined as the derivative of the *cumulative* distribution function $F(x)$

$$f(x) = \frac{\mathrm{d}F(x)}{\mathrm{d}x}. \tag{A.2}$$

A.2 The Normal Distribution

A random variable X is said to be normally distributed or Gauss distributed if its PDF are given by

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{ -\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2 \right\}. \tag{A.3}$$

Here, μ is the mean and σ is the standard deviation of the process.

A.3 The Chi-square Distribution

If $\{X_t : t = 1, 2, ..., n\}$ is a sequence of independent and identical normally distributed random variables with mean μ and standard deviation σ, then the random variable

$$Y = \frac{\sum_{t=1}^{n} \left(X_t - \overline{X}\right)^2}{\sigma^2} \tag{A.4}$$

is chi-square distributed with $n-1$ degrees of freedom. Here, $\overline{X}$ denotes the sample mean of the sequence $\{X_t\}$, i.e. $\overline{X} = \frac{1}{n}\sum_{t=1}^{n} X_t$. The PDF of a chi-square distribution is given by [48]

$$f_m(y) = \begin{cases} \frac{y^{(m-2)/2}\exp[-y/2]}{[2^{m/2}\Gamma(\frac{m}{2})]} & \text{if } y \geq 0 \\ 0 & \text{if } y < 0, \end{cases} \tag{A.5}$$

where $\Gamma(\frac{m}{2})$ is the *gamma function* [49] and m is the number of degrees of freedom.

A.4 The Chi distribution

If $\{X_t : t = 1, 2, ..., n\}$ is a sequence of independent and identical normally distributed random variables with mean μ and standard deviation σ, then the random variable

$$Z = \sqrt{\frac{\sum_{t=1}^{n} \left(X_t - \overline{X}\right)^2}{\sigma^2}} \tag{A.6}$$

is chi distributed with $n-1$ degrees of freedom. The PDF of a chi distribution is given by [48]

$$f_m(z) = \begin{cases} \frac{z^{(m-1)}\exp[-z^2/2]}{[2^{(m-2)/2}\Gamma(\frac{m}{2})]} & \text{if } z \geq 0 \\ 0 & \text{if } z < 0, \end{cases} \tag{A.7}$$

where $\Gamma(\frac{m}{2})$ is the *gamma function* [49] and m is the number of degrees of freedom.

If Z is chi distributed with m degrees of freedom, then $Y = Z^2$ is chi-square distributed with m degrees of freedom.

A.5 Kurtosis

The "peakedness" of a PDF of a real-valued random variable X may be measured by the *kurtosis*, defined as

$$\mathrm{Kurt}\Big[X\Big] = \gamma_2 = \frac{\mathrm{E}\Big[(X-\mu)^4\Big]}{\mathrm{E}\Big[(X-\mu)^2\Big]^2} - 3 = \frac{\mu_4}{\sigma^4} - 3. \tag{A.8}$$

Here, μ_4 denotes the fourth order moment about the mean, while σ denotes the standard deviation. The subtraction of 3 is a convenient correction that makes it simpler to calculate the kurtosis of a sum of random variables. With this definition, the kurtosis of the normal distribution is 0.

Distributions with negative valued kurtosis are called *platykurtic*, distributions with zero kurtisos are called *mesokurtic* and positive valued kurtosis are referred to as *leptokurtic*.

A.6 Skewness

The skewness is a quantification of the *asymmetry* of a PDF. A skewness greater than zero signifies that the tail at the right hand side of the mean of the PDF is more pronounced (longer) than the left tail, while a negative-valued skewness signifies the opposite. The definition of skewness is given in terms of the third standardized moment and the standard deviation of the distribution:

$$\gamma_1 = \frac{\mathrm{E}\Big[(X-\mu)^3\Big]}{\mathrm{E}\Big[(X-\mu)^2\Big]^{3/2}} = \frac{\mu_3}{\sigma^3}. \tag{A.9}$$

A.7 The Least Squares Principle

When a line $y = \alpha + \beta x$ is fitted to a sample of n paired observations $(x_1, y_1), ...(x_n, y_n)$, α and β should be chosen such that the sum of the squared, vertical distances from the points to the line is minimized, i.e.

$$\underset{\alpha,\beta}{\mathrm{argmin}}\left(\sum_{i=1}^{n}\Big|y_i - \alpha - \beta x_i\Big|^2\right). \tag{A.10}$$

Furthermore, if the uncertainties associated with the different observations are not equal, each observation could be attributed a weight w_i. Then, the optimization problem reads

$$\underset{\alpha,\beta}{\operatorname{argmin}}\left(\sum_{i=1}^{n} w_i \Big| y_i - \alpha - \beta x_i \Big|^2\right). \tag{A.11}$$

This minimization problem could be solved numerically by means of the Levenberg-Marquardt algorithm [50].

Appendix B

Smoothing by Means of the Wavelet Transform

When smoothing a time series, the aim is to obtain a time series where rapid fluctuations below a certain time scale are removed. As already mentioned, there exist several smoothing algorithms. In this chapter, a simple smoothing algorithm based on the wavelet transform, introduced in Section 2.5, is described. As an example, the time series of daily mean prices in OMEL will be employed.

Recall from Section 2.5 that the discrete wavelet transform (DWT) decomposes the data into a new basis of functions, namely the *wavelet* basis. Since these functions are defined at characteristic dyadic time scales, the fluctuations of the original time series will be decomposed into dyadic time scales. The underlying idea of using wavelets as a smoothing tool is to first transform the series into the wavelet domain, then set all wavelet coefficients corresponding to a certain time scales lower than a_{smooth} to zero, and finally transform the coefficients back to the time domain. This will result in a smooth time series where fluctuations at time scales smaller than a_{smooth} are averaged out over larger time scales.

Employing the DWT, problems may arise at the boundaries of the time series. In order to overcome any problems, a certain number of zero valued data points should be added at each end of the series. The only punishment associated with adding many zero values are longer calculation time. Since the calculation time for the purposes of this report are short anyway, at least a couple of hundred zero values are added. It should also be kept in mind that in order to employ the DWT with the pyramidal algorithm, the number of data points must be 2^n, for integer n.

The smoothing algorithm for averaging out fluctuations corresponding to time scales smaller than $a_{\text{smooth}} = 2^S$ could be summarized as follows:

1. Add zero valued data points to the beginning and the end of the time series, until the length of the time series is a power of two, see Fig. B.1 (a). At least 100 zero valued points should be added at each end.

2. Employ the DWT to the time series.

3. Set all wavelet coefficients corresponding to time scales smaller than a_{smooth} to zero.

4. Employ the inverse DWT to transform the series back to the time domain, see Fig. B.1 (b).

5. Cut off the values corresponding to the zero valued data points that were added at step 1, see Fig. B.1 (c).

This smoothing procedure is employed in Section 5.1.1 in order to extract the weekly fluctuations of the time series (which is achieved by subtracting the smoothed time series from the original time series). In such a case, a_{smooth} should be chosen such that at least in the time scales in the order of a week are smoothed. For the results given in this report, $a_{\text{smooth}} = 2^5$ days, meaning that fluctuations at time scales $a = 2^4$ days $= 16$ days and below are averaged out.

Finally, it should be mentioned that another way of extracting the weekly fluctuations could be to set the wavelet coefficients corresponding to scales *larger* than a_{smooth} to zero in step 3. This illustrates the fact that there are many ways of obtaining the small scale fluctuations, and it is difficult to measure which methods that give the best results. Furthermore, the effect of employing different wavelet bases, is not investigated here. The results given here are based on the Daubechies wavelet of fourth order.

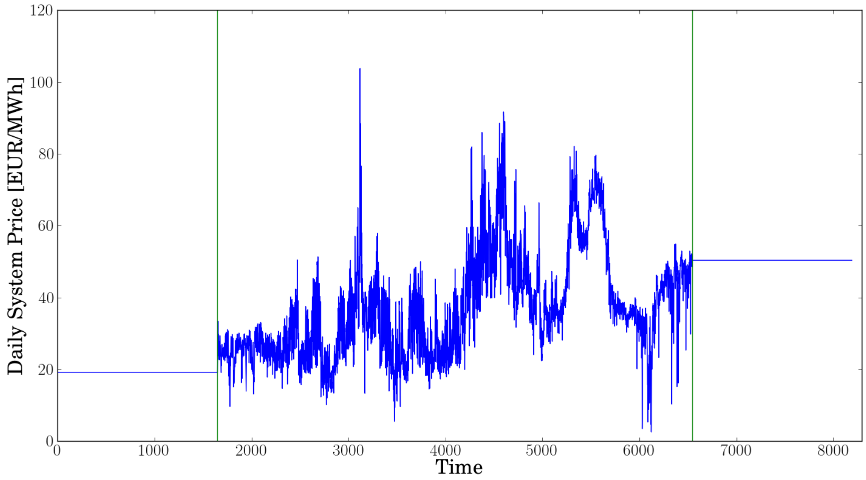

(a) Step 1: Adding zeroes.

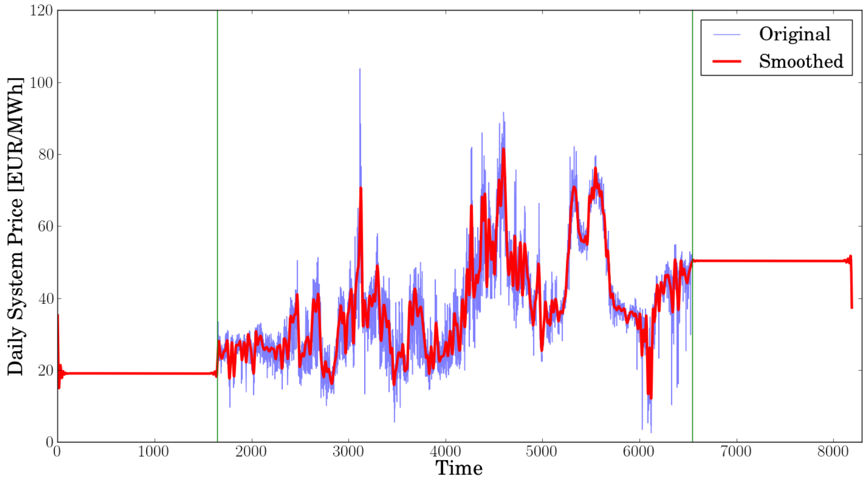

(b) Step 4: Transforming back to the time domain.

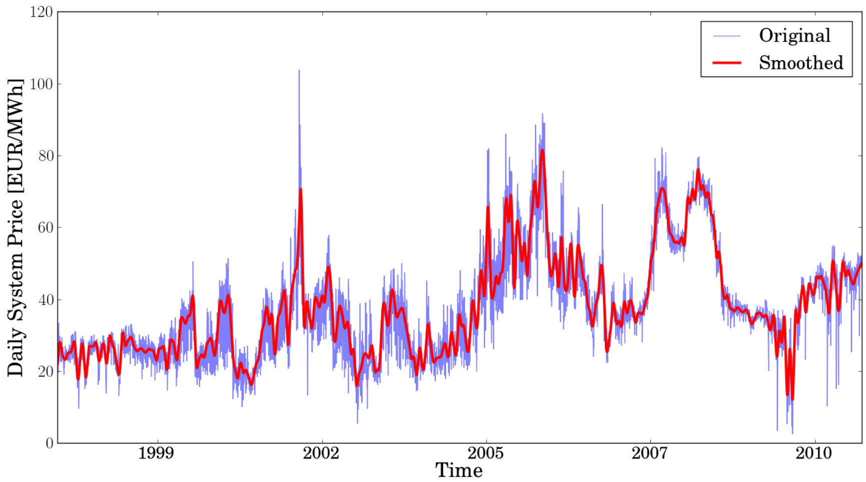

(c) Step 5: Cutting off the artificial parts.

Figure B.1: Step 1, 3 and 4 of the smoothing procedure with $a_{\text{smooth}} = 2^5$ days. The time series is the OMEL daily mean prices.