Forecasting daily elspot prices - whether weather will work

Oscar Egnell

Norwegian University of Life Sciences
Department of Economics and Resource Management (IØR)
Master Thesis 30 credits 2012

PREFACE

This master thesis marks the end of my two year master program at the Department of Economics and Resource Management at the Norwegian University of Life Sciences in Ås. During the program I have taken the chance to increase my knowledge of energy economics. This is my final work which aims at proving my skills within this field.

I would like to thank my advisor, Professor Olvar Bergland, for assisting me with his expertise and knowledge. He has guided me with great support and encouraging discussions. I would also like to thank Magnus Baltscheffsky at WeatherTech, who has offered supporting thoughts regarding the weather variables. Special thanks to all family and friends who have helped with proofreading.

Oslo, May 2012

Oscar Egnell

ABSTRACT

It is always important for participants in a market to receive and process all available information in order to make rational decisions. At Nordpool this is of extra relevance, since participants have to hand in bids on a day-ahead basis. To have an accurate forecast of the prices which regards the relevant information is therefore crucial. This thesis' objective is to forecast Elspot daily prices and testing if weather variables will increase the accuracy of the forecast. This is done by implementing different ARIMA models and including explanatory variables. The variables used are temperature and precipitation.

We find that out of the different ARIMA versions, the SARIMA stands out as the most well defined. This is a seasonal ARIMA model accounting for weekly seasonality. However, after much testing it was not possible to find a model which controls for all serial correlation, there were still some left in the residual. When including the weather variables the model got slightly improved. Most of the improvement originates from the inclusion of temperature.

The forecast does a good job of predicting prices one day ahead. It also gives an ok indication of the next four to five days price movements. The inclusion of weather variables does not improve the forecast as much as expected. It increases the certainty of the forecast only slightly in volatile periods. We conclude that the model does a god job forecasting, but it can be developed further. Future research should focus on completely removing the serial correlation, and extracting all possible information from the weather variables.

SAMMANFATTNING

Det är viktigt för alla aktörer i en marknad att ha tillgång till all tillgänglig information för att ha möjlighet att ta rationella och genomtänkta beslut. Detta är extra viktigt när det gäller Nordpool, eftersom det är en marknad där aktörerna lägger bud på elektriciteten dagen innan. Det är därför nödvändigt att ha tillgång till prisprognoser som ger noggranna uppskattningar baserat på relevant information. Uppsatsens syfte är att göra en prisprognos på det dagliga Elspotpriset, samt att testa om prisprognosen förbättras med hjälp av vädervariabler. Detta görs med hjälp av olika ARIMA - modeller och förklarande variabler. Vädervariablerna består av temperatur och nederbörd.

Av de olika ARIMA – varianterna var det SARIMA som var den mest väldefinierade modellen. Det är en ARIMA med säsongsvariation som tar hänsyn till det veckovisa mönstret. Efter att ha testat många olika modeller så har det inte varit möjligt att identifiera en modell som förklarar all seriekorrelation. Efter att ha inkluderat vädervariablerna förbättrades modellen något. Det mesta av förbättringen härstammar från temperaturvariabeln.

Modellens prisprognos fungerar bra när den förutsäger morgondagens pris. Den ger också en ok indikation för de nästa fyra till sex dagarnas prisrörelser. Användandet av vädervariablerna i prognosen ger inte en lika stor förbättring som var antaget. Dem ökar prisprognosens säkerhet lite under oroliga och volatila perioder. Sammanfattningsvis ger modellen en bra prognos för morgondagen pris, men det finns utrymme för att utveckla den till att få kontrollerat all seriekorrelation och att kontrollera för all väderinformation.

INDEX

1. Introduction ... 1
1.1 Objective ... 1
1.2 Limitations ... 2
1.3 Previous literature ... 2
1.4 Structure ... 3
2. The Nordic Electricity Market ... 4
2.1 System price and price areas ... 4
2.2 Price mechanism ... 6
2.3 Power production and supply ... 7
3. Model and Data ... 9
3.1 ARIMA ... 9
3.1.1 The three parts of ARIMA ... 10
3.1.2 Different versions of ARIMA ... 11
3.1.3 Tests ... 12
3.1.4 Forecasting ... 13
3.2 Data ... 13
3.2.1 Descriptive Statistics ... 14
4. Results and Analysis ... 18
4.1 Identification ... 18
4.2 Parameter Estimation and Testing ... 21
4.3 Forecasting ... 25
5. Conclusion ... 28
References ... 29
Appendix A ... 30
Appendix B ... 31

1. Introduction

Norway's electricity market was one of the first to be liberalized in the beginning of the nineties. This was soon followed by the inclusion of the neighboring countries into Nordpool, the market place. Today, Nordpool is the main market where the Nordic countries trade electricity in a day-ahead physical market. The prices can move from high maximum values to low minimum values over a year, making it more volatile for specific periods.

Because of the day-ahead practice, buyers and sellers need to report their bids the day before they go into effect. For this reason it is important for the participants to use production and consumption forecasts, together with forecasts of the electricity price, in order to collect all the available information and make rational decisions when handing in their bids. If participants are wrong in their bids it can be costly since one has to correct the mistakes in Elbas, an intraday trade market, which is used for controlling and maintaining the balance between supply and demand. There is also a market for financial contracts, such as futures, which is based on the underlying commodity, i.e. electricity. These contracts are used for risk management, i.e. hedging the prices of electricity. There is a great need for information and good forecast when acting on this market as well.

Due to the nature and landscape, Norway is a country which has more than 90% of its electricity production from hydropower and most of it is storable in reservoirs. The water in the reservoirs originates from precipitation and has a strong seasonal pattern. Because of a history with easy access to natural resources for electricity production, most of the heating and energy-demanding services in Norway are based on electricity. It follows that weather information is an important price driver both for electricity input and output. Weather can thus be used as information when forecasting the electricity prices in order to improve it.

1.1 Objective

The objective of this master thesis is to explore the ARIMA models ability to forecast daily electricity prices in the Norwegian market, mainly the price area NO1, i.e. the Oslo area. The forecast is for short-term, i.e. the next-coming days, It is also the objective to test different weather variables, to see whether they improve the forecast or not. Including weather variables is assumed to improve the forecasting of the model, since weather is regarded as a price driver in a hydropower-based electricity market. There are many different time series models to choose from, and our objective is not to prove that ARIMA is the best one for modeling the Norwegian electricity prices, but rather to test explanatory variables capability to improve the forecast. This will be done by performing a time series analysis using ARIMA and its extensions, to find a well suited model for the Norwegian electricity prices. By using data consisting of electricity price quotes, the ARIMA will use the prices own information and history to build a model. This will be mixed with weather information aiming at improving the model. The prices are daily prices from the Norwegian price area NO1.

1.2 LIMITATIONS

Since the prices are calculated and reported once a day, we have decided to predict daily prices instead of hourly. Traders and participants may prefer information based on the traded asset, which is hourly. However, it is of interest to find out how the prices are moving and predict tomorrow's movement in the daily prices, since this gives an indication of the total trends for all the 24 hours. In order to include weather variables in the model, the prices of NO1 have been chosen. If the system price were chosen, it would be needed to collect weather data for the whole area. And as will be shown, the distribution between NO1 and System prices does not differ too much, which implies that a model describing NO1 could be used on system prices as well.

1.3 PREVIOUS LITERATURE

This section will give a brief insight into the literature concerning the forecasting of time series data. It will present an overview of already used methods to forecast electricity prices, and also other research done with ARIMA models or similar.

There are many different efforts to model the Elspot prices and forecast it. Several of these include stochastic dual dynamic programming (SDDP) and/or more advanced algorithms. The SDDP analysis is based on forecasting the cost function, therefore creating the future cost function (FCF), for thermal plants and hydropower plants (Pereira, et al., 1999). This is a model applied by energy analysis firm Point Carbon. The PoMo model is developed by EME Analys to forecast the weekly electricity prices. It bases its analysis on hydro power information and estimates marginal cost curves for thermal power plants (Fridolfsson & Tangerås, 2008).

In "ARIMA Models to predict Next - DAY Electricity Prices" (2003), the authors propose an ARIMA model to predict the next day 24 hourly prices. They apply the models on the Spanish as well as the Californian markets. Their results are models that include the last 5 hours for Spain and 3 hours for California and then 24 hour intervals to include the daily movements, and also weekly. This is included in the AR part, but they also include the daily and weekly trends in the MA parts (Contreras, et al., 2003). Their models are somewhat complicated and it could be argued that they do not hold up to the unwritten rule of a parsimonious model.

A discussion paper written at NHH, *Norwegian School of Economics*, investigates the implications of wind power from Denmark on the Norwegian price volatility. This study uses ARIMA to model the price volatility in Norway, and add wind power as an independent variable. It implements seasonal adjustments in the model, both as an extended ARIMA model (SARIMA) and as a model with dummy variables such as day of the week (Mauritzen, 2010). The paper models hourly prices as well as daily.

Niels Haldrup and Morten Nielsen suggest a fractional integration method to model and forecast the Nordic Elspot hourly prices (Haldrup & Nielsen, 2004). They use data from only 2000-2003. The specified model they use is an extension of ARIMA called ARFIMA, or more precisely *auto regressive fractional integrated moving average*. This means that instead of integrating the time series with 1 or 2, which is normal, the model finds a fraction from {-0.5 – +0.5}.

In a master thesis by Kristian Hjelset and Line Monsbakken, the objective is to find a model to describe the Norwegian electricity prices and the prices for contracts for differences (Hjelset & Monsbakken, 2005). They test different time series models, amongst others ARIMA, and find that a GARCH-like model gives the best fit. They discuss that the addition of MA to an AR model only gives slightly improved results, and that it was not the best model for modeling daily prices. However, it does not appear that they have utilized all the options available in ARIMA modeling.

1.4 Structure

This paper will start off by describing the background of the Nordic electricity market and the mechanisms surrounding it. This is important in order to understand the prices of electricity. Chapter three will have a description of the applied time series model, followed by a section on the data set. This section will also include descriptive statistics. In chapter four the results will be presented, together with the analyses. The final chapter concludes the work, and offers some suggestions for further research.

2. THE NORDIC ELECTRICITY MARKET

Since the Norwegian electricity market liberalization started in 1991 the other Nordic countries have followed, and today they are integrated into the same system via Nordpool. This is one of Europe's biggest power markets, and it is the place where participants can sell and buy electricity. It is owned jointly by the Nordic transmission system operators –TSO–. According to Nordpool there are 350 companies from 20 countries trading on the market, and it had a turnover of 316 TWh in 2011. The market is a net pool, which means participants can buy and sell electricity via the pool, or they can agree to bilateral contracts outside of Nordpool's market (Green, 2005). 74% of the Nordic electricity production is traded in the pool, and the rest is traded through bilateral contracts (Nord Pool Spot AS, 2012).

The market is based on a day-ahead system, where the buyers and sellers report their expected consumption/production of electricity for every hour of the next day as well as the price they are willing to buy/sell it for. The bids are due 12:00 CET and Nordpool then aggregates the bids and calculates the intersection of the sell and buy curves, which gives the system prices for the 24 hours of the next day. The prices are generally reported to the market around 12:30 to 12:45 CET. These are the spot prices, the actual price for electricity the next day, and they are set at Nordpool Elspot.

There is also a financial market where it is possible to trade financial contracts based on Elspot and it is handled by Nasdaq OMX Commodities. These contracts include futures, forwards, and contracts for differences –CfD–, and are used for hedging and risk management. All these contracts are based on the system price, and CfD's are contracts that are based on the difference between system price and the area prices. If the balance between demand and supply from the spot market is not maintained in the following day, it has to be corrected for in the Elbas intraday market (Nord Pool Spot AS, 2012). The offset from the spot markets balance can be due to power plant failure or unexpected demand shifters.

2.1 SYSTEM PRICE AND PRICE AREAS

Electricity is different from other commodities because it experiences different physical laws. The production has to be met by demand at all times, and is delivered momentarily. One cannot buy electricity and store it (at least not in a significant amount), but instead it has to be delivered through the grid. The grid´s service is to transport electricity from the place of production to the place of consumption, and it is constrained by the capacity limits of the cables. This means that electricity can be transported from the hydropower plant to the households or industrial factories, but only to a certain level. When these limitations, called bottlenecks, are exceeded for a long period of time one can control for this by dividing the market into different areas so that there are individual prices for each one (Nord Pool Spot AS, 2012). The price difference will then represent the cost of the capacity limits. If there were no limits within the whole Nordic market it would be possible to set one price for the region. When Nordpool receives all the bids for the next day, they first calculate one price without the price areas, called the system price. Then the expected power flows between the areas are accounted for and if there are bottlenecks, different prices will emerge.

The Nordic market is divided into several different price areas; Norway NO1-5, Sweden SE1-4, Denmark DK1-2 and Finland FI. The Swedish areas where introduced in November 2011, and before that it was one price for Sweden. Because of the one price, bottlenecks emerged and were handled by counter trade before the introduction of price areas. Countertrade is also applied if bottlenecks occur within a price area for a short period (Nord Pool Spot AS, 2012).

The Norwegian areas have changed a few times since Nordpool was introduced. Figure 2.1 and 2.2 show the price areas changes from 2000. In the first picture there are only two price areas, basically dividing the south and the north. Following was a decision to divide the country into more zones, in December of 2002 NO1 was split into three areas (Ministry of Petroleum and Energy, 2003). In 2010 the NO1-5 was introduced, which is what is currently in use. These changes have reshaped the NO1 price area from 2000 to present, from being one area of the south to being one of three areas in the south. The size has decreased significantly.

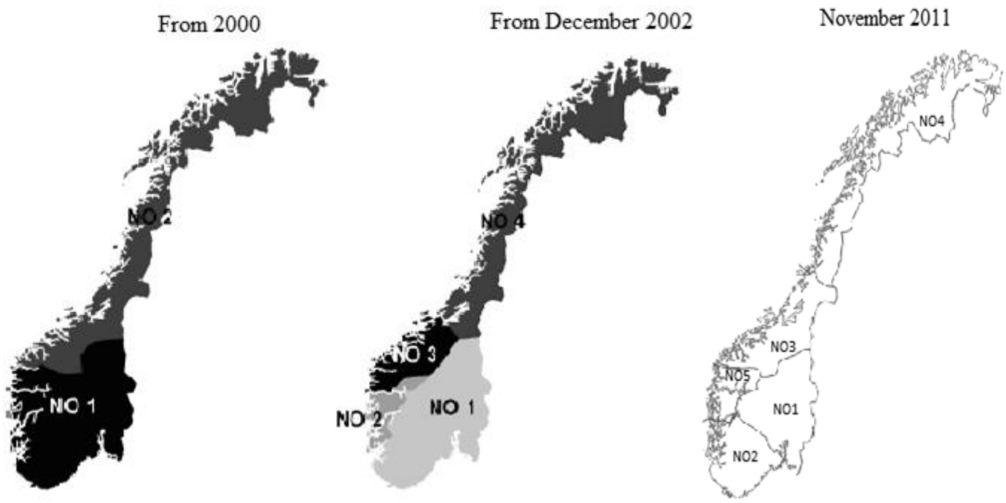

FIGURE 2.1 NORWEGIAN PRICE AREA CHANGES, SOURCE: MINISTRY OF PETROLEUM AND ENERGY

The Nordic countries are connected via the grid, making it possible to export and import power. This is important when there are local/national shortages and surpluses of power supply. The Nordic countries are also connected to other close neighbors such as Germany, Estonia and the Netherlands. The whole grid system is shown in figure 2.3. NO2 area is connected to the Netherlands and Denmark, via NorNed. The NO1 is connected to Sweden, and there is a cable planned to deliver power between NO1 and south of Sweden called Sydvestlinken (Statnett, 2011). These connections increase the efficiency of the Norwegian power market. They give rise to a higher degree of security in terms of being able to supply power at all times for all the demand. This is because the countries that are connected through the grid have different power producing technologies, which gives a desired technology differentiation.

FIGURE 2.2 NORDEL TRANSMISSION GRID OF NORTHERN EUROPE, SOURCE: SWEDENERGY

2.2 PRICE MECHANISM

Nordpool receives the bids of the buy/sell offers from consumers/producers. This price setting mechanism is an auction based exchange and delivers an equilibrium price between supply and demand. This results in marginal pricing (Nord Pool Spot AS, 2012). The figure shows the sell/buy curves, and the interaction which leads to equilibrium.

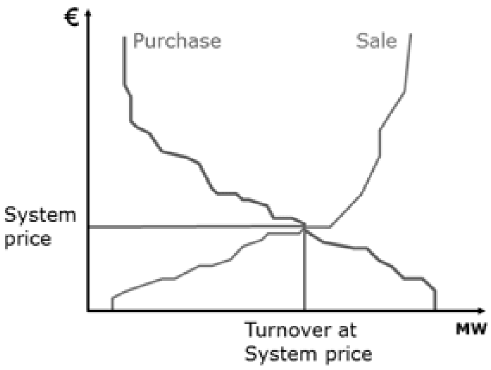

FIGURE 2.3 SUPPLY AND DEMAND, SOURCE: NORDPOOL

The purchase, or demand, often appears different depending of what time of the day. If demand is higher, it will lie further to the right in the graph. The demand depends on how much electricity is needed. During summer, there is often less economic activity and less need for heating, resulting in less demand, shifting the curve inwards. The sale, or supply, move to the right if there is offered more electricity.

Since Norway is hydropower based, the supply should be represented by the marginal cost of producing electricity from hydropower. Once the investment of the hydropower-plant is made, it does not cost more to produce more electricity. Therefore, the marginal cost is often zero. However, there is an intertemporal opportunity cost of using the water value today or saving it for the future (Førsund, 2007). The reservoirs help the producers to transfer the water value from spring and summer to late autumn and winter, maximizing the consumer's marginal willingness to pay.

2.3 Power production and supply

Most of the Norwegian electricity production is based on hydropower. In 2009, 96% of Norway´s total electricity production derived from hydropower (NVE, 2010). The hydropower is mostly based on water reservoirs which are filled up during the snow melting, combined with summer and autumn's rain falls. The inflow of water is a stochastic variable that varies with time. This means that different years can produce different amounts of hydropower-based electricity (see figure 2.4). Since Norway´s production is mostly based on hydropower, it is exposed to the changes in yearly water inflow. During years of low inflow, electricity is imported from neighboring countries such as Sweden and Denmark. In Denmark most electricity production results from thermal power, but has a larger share of wind power than most (Danish Energy Agency, 2011). Sweden on the other hand has more of a mix, which contains nuclear, hydro, thermal and small amount of wind power.

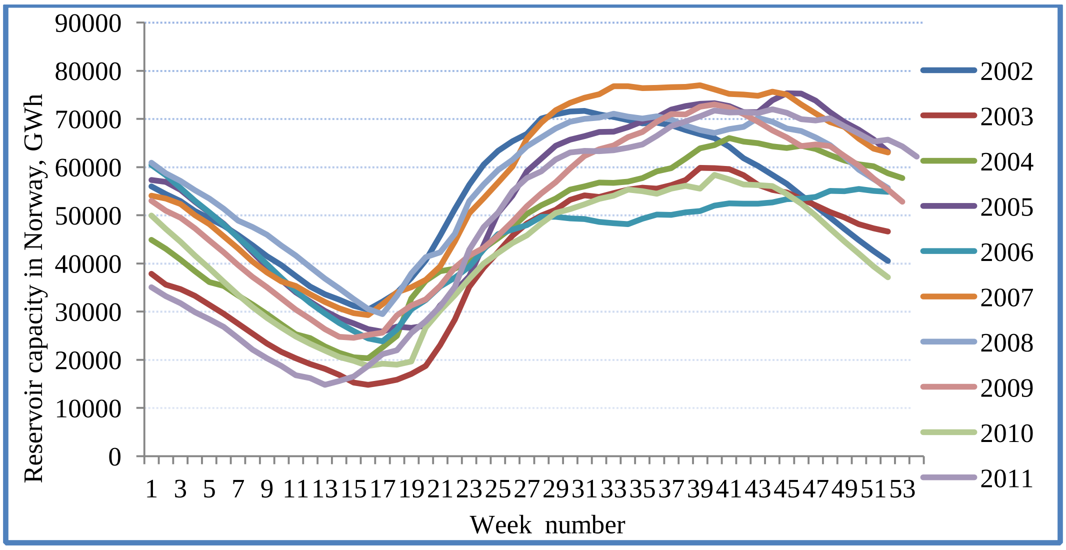

Figure 2.4 Yearly inflow 2002-2011, Source: NVE

Because of the high dependence of hydropower in Norway, the main production input is water inflow. This is basically rain and snow melting. Therefore, one of the main price drivers of the electricity prices is precipitation. Rain or snow increases the water inflow to the reservoir, which in turn increase the supply. If there is heavy rainfall in a short period of time, reservoirs are filled up more rapidly. The reservoirs are limited by environmental constraints such as maximum and minimum water levels (Førsund, 2007). If heavy rainfall leads to maximum reservoir levels, producers must produce electricity, even if the prices are very low and the water value is indicating saving the water for coming periods. Other price drivers are energy fuels, installed capacity, transmission constraints and demand factors such as weather sensitivity (Hughes & Parece, 2002).

There was recently an introduction of electricity certificates in Norway, which is already in use in Sweden. The certificates will promote renewable energy sources, mostly hydro- and wind power. When it comes to hydropower, Norway have unused water resources. The ability to utilize these are however limited, due to natural interference and other factors (NVE, 2010). The wind power potential is regarded as good in Norway due to the great Atlantic coastline, but the investment cost exceeds that of building out the potential hydropower. Consequently wind power is not projected to increase much in percentage of total electricity production, until hydropower is explored to its full potential.

3. Model and Data

Since the prices of electricity are given for different times, they need to be modeled with time series analysis. The time series model which has been chosen for this objective is ARIMA. It is a model that uses the history of the time series own data to model and predict, i.e. past observations are used to describe today's electricity prices. This chapter will describe the model and will be followed by a section describing the used data, software etc.

3.1 ARIMA

In order to forecast the Elspot prices, ARIMA, *Autoregressive Integrated Moving Average*, has been chosen. In its native form, the model is a univariate model, meaning it uses past observations or realizations to explain the movement of the time series. Basically, ARIMA is an autoregressive process with added moving average. An AR process is when a time series is described by its past values, known as a difference equation (Enders, 1994). AR (1) means that the lagged value describes today prices. MA (1) uses the last observation errors to describe today prices. The model with the parameters is often written as ARIMA (p, d, q), where p is the AR lags, d is the order of integration and q is the MA. Many time series do not have to be integrated in order to fulfill the stationary condition, in which case the model is just an ARMA.

This is a model which is good for short-run modeling and forecasting. It takes into account the previous prices and the errors of the model. It can be extended to account for seasonal effects called SARIMA, Seasonal Autoregressive Integrated Moving Average. The model can also be extended to include independent variables, in which case it is called ARMAX. If the series has a long-memory process, ARFIMA could be used in order to integrate using a fraction (Stata Corp, 2011). When the right model is identified, using the recently stated options, it is possible to use it to forecast prices.

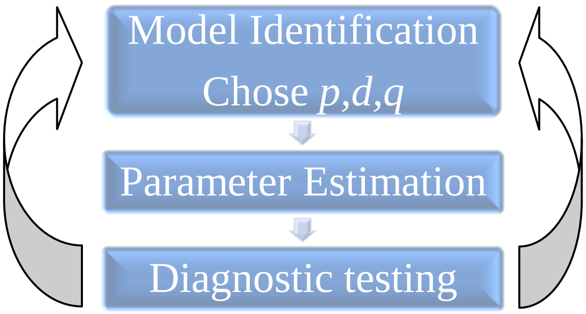

Figure 3.1 Illustration of Box-Jenkins

The figure describes Box-Jenkins proposed method to perform ARIMA modeling. The first step includes identifying the AR (p), if there should be any differentiating I (d), and how many MA (q) there should be. The second step is the parameter estimation. Stata´s ARIMA equation is used to calculate the coefficients for the AR and MA parts, together with other statistical information. In the third and last part, tests are performed in order to validate the model. This procedure is redone until one find the model best suited for describing the price movements (Enders, 1994).

3.1.1 The three parts of ARIMA

Stationary time series is a requirement in the ARMA model. However, if the series is not stationary, it may be corrected by integration of various degrees (then making it an ARIMA model). This is important because if the series is stationary, or made stationary, the mean, variance and autocorrelations can be approximated (Enders, 1994). Stationary is described by Enders as:

> *"…a time series is covariance stationary if its mean and all autocovariances are unaffected by change of time origin /…/ there is no ambiguity in using the terms stationary and covariance stationary interchangeably."*
>
> (Enders, 1994, p. 69)

However, if the time series is large enough and the AR (1)-coefficient is less than one, there will be a convergence towards the mean and therefore making it stationary (Enders, 1994). If the series is non-stationary, it can be corrected by integration in the ARIMA specification (*p*, *d*, *q*), where *d* will take on an integer. To take logs of the time series or control for inflation might also help. If the time series is stationary by origin, the model does not need parameter *d*, i.e. I (0), hence ARMA. From here on, if nothing else stated, the base model will be ARMA since I (*d*) is not a parameter, but an action applied on the time series in order to fulfill the ARMA requirements.

An AR process is when you use lags of the dependent variable to describe it, or put differently, when previous prices affects today prices. The AR equation with *p* lags is given by:

$$y_t = \alpha_o + \sum_{i=1}^{p} \alpha_i y_{t-i} + \varepsilon_t \qquad \text{Eq. 3.1}$$

where y_t is the observation at time t, α_0 is the intercept, α_i is the coefficient of the lag and ε_t is a white noise process of the errors (Enders, 1994).

Using **moving average** is done by extending the previous equations to incorporate previous errors (or when estimated, residuals) to improve the model. Practically, you lag the residuals to get describing variables of the price. This is an advantage because you use the history of the time series to correct the model. The MA equation looks like the following:

$$x_t = \sum_{i=0}^{q} \beta_i \varepsilon_{t-i} \qquad \text{Eq. 3.2}$$

where x_t is the moving average of order *q*, β_i is the coefficient of the lagged residuals (Enders, 1994).

Equations 3.1 and 3.2 can be combined to create the ARMA equation:

$$y_t = \alpha_o + \sum_{i=1}^{p} \alpha_i y_{t-i} + \sum_{i=0}^{q} \beta_i \varepsilon_{t-i} + \varepsilon_t \qquad \text{Eq. 3.3}$$

In order to identify ARIMA (*p, d, q*), it is common to use the Autocorrelation Function (ACF) and Partial Autocorrelation Functions (PACF). Autocorrelation is when a variable is correlated over time (Woldridge, 2009). ACF is the autocorrelation between the lags, often shown as a graph. The difference to PACF is that PACF shows the correlation only for the specified lag. The PACF for lag 3 is therefore only the correlation for lag 3, discarding the correlation for lower order of lags. The ACF for lag 3 will include the correlation for lower order lags (Nau, 2005). By investigating the functions it is possible to get an understanding of the properties of the time series one is modeling. If the ACF decays geometrically to zero, it is stationary, i.e. I (0). This is a result from what was stated previously, namely that the AR (1) coefficient must be less than 1 for stationary time series. Testing for stationarity can be somewhat troublesome because there can be a fine line between processes that are non-stationary, stationary and processes which needs fractional integration. Except from studying the ACF, there are statistical tests commonly used for stationarity, namely the augmented Dickey-Fuller and Phillips-Perron tests. They both have a null hypothesis of non-stationarity i.e. one unit root; therefore a rejection of the null hypothesis will indicate a stationary process (Stata Corp, 2011).

When it comes to identifying the numbers of p and q, PACF is used for AR (p) and ACF for MA (q). If there is a clear spike at lag 1 in PACF, one should maybe use AR (1). And if there is a significant spike at lag 5 in the ACF one might include MA (5). However, it is often not as clear cut as that in real time series, which makes it a try and retry process (Nau, 2005). In terms of choosing parameters, it can be tempting to include many variables in order to get a good model. However, this can lead to over-fitting, which might give good test scores, but will lead to a less accurate forecast (Enders, 1994). That is why parsimony is often recommended when identifying the ARMA model.

3.1.2 Different versions of ARIMA

Time series often include seasonality. It can be in the form of daily, weekly, monthly, quarterly, yearly or other time frames. For example, Christmas tree sales have a strong seasonal pattern for the month of December. In figure 2.4 there is a yearly seasonal effect in the reservoir levels. If the data spans over multiple seasons (e.g. day, month etc.) one will see the seasonality if it is plotted against the time line. It will also appear in ACF/PACF because of the serial correlation.

Seasonality can be controlled by using SARIMA $(p, d, q)\ (P, D, Q)_S$. This is a multiplicative model, where you use the base from ARMA, but extend it to include Seasonal Autoregressive (P), Seasonal Integration (D), Seasonal Moving Averages (Q) and a fourth part which describes the kind of seasonality (S). A SARIMA model could look like this: $(1, 1, 1)\ (0, 0, 1)_4$ and is interpreted as one autoregressive term, first differenced, one moving average extended with no seasonal autoregressive term, no seasonal differentiating and one lag 4 seasonal moving average. If this was monthly data, lag 4 will represent quarterly seasonality (Stata Corp, 2011).

ARFIMA is a variant where instead of choosing an integer degree of integration, d, the model can accept a fractional value. This can be of interest when a time series show signs of a long-memory process. This is when the ACF decays rather slowly, or slower than a short-memory process. If a long-memory process was analyzed using standard ARIMA, differencing when not needed could lead to over-differencing. Likewise, not differencing could lead to a non-stationary time series.

When applying ARFIMA, the fractional integration *d* captures the long-run effects, while the rest of the model, ARMA, captures the short-run effects (Stata Corp, 2011).

It is possible to extend an ARMA model to include other variables, such as variables that don´t originates from the time series itself. These would be explanatory or independent variables, and a model which includes them is called ARMAX. In such a model, the dependent variable is modeled as linear combination of the independent variables. As with previous versions, ARMAX can be used jointly with seasonal extension.

3.1.3 Tests

One of the steps in the Box-Jenkins procedure is diagnostic testing. Testing is important since one can try many different ARMA models and they need to be analyzed in order to realize which model is the most appropriate one. The tests that has been applied are AIC (Akaike Information Criterion), BIC (Bayesian Information Criterion a.k.a. SBC) and the squared correlation (of real and predicted observations).

The AIC and BIC are criterions which should be minimized. They are tests which are suited for likelihood estimation (which is used for estimating ARMA parameters).

$$AIC = T \ln(\text{residuals sum of squares}) + 2n \qquad \text{Eq. 3.4}$$

$$BIC = T \ln(\text{residuals sum of squares}) + n \ln(T) \qquad \text{Eq. 3.5}$$

where n = number of parameters estimated ($p + q$ + possible constant term);
T = number of usable observations. (Enders, 1994)

It is important that the tests are applied to the same sample size, since T appears in both equations. These two criteria regard the general idea of parsimony since a greater n will increase the score, therefore making it worse off. This applies to BIC more so than AIC.

The R^2 is usually applied in standard OLS regression. However, it is not well suited for likelihood estimation, and is often not calculated by statistical softwares. Instead, the squared correlation of y_t and the predicted y_t is applied. This gives a fraction 0-1 of the models fit, where 1 is a “perfect” fit.

All of the above tests are used jointly to determine the goodness of fit for the different models. None is used exclusively since it is not clear which models is best and it needs careful considerations to determine the “best” model. Apart from these tests, a few other general considerations have been applied. These include significance of the estimated parameters, and the residuals should be a white-noise process (Enders, 1994). That means that there should be no serial correlation left in the residuals. Also, Box-Jenkins proposed the model to be invertible. The consequence of invertibility is that the coefficient of the MA (q) should not be equal or greater than 1. If there are more than one moving averages, they should not sum up to 1, or be greater than 1 (Enders, 1994).

3.1.4 Forecasting

One of the advantages of using ARMA is its quality in forecasting. Once the identification and diagnostic testing is done and a model has been applied, the models coefficients can be used to perform out-of-sample forecasting, i.e. predict values further than the underlying observations time horizon. This is also known as dynamic forecasting. If the model includes describing variables, i.e. ARMAX, these will be included in the forecast equation. In order for them to help the forecasting, they need to exist in the data set for the time that is forecasted. The forecast equation with an ARIMA (1, 0, 1) model for period t+1 and t+2, respectively, is defined as follows;

$$Ey_{t+1} = \alpha_0 + \alpha_1 y_t + \beta_1 \varepsilon_t \qquad \text{Eq. 3.6}$$

$$Ey_{t+2} = \alpha_0 + \alpha_1 y_{t+1} + \beta_1 \varepsilon_{t+1} \qquad \text{Eq. 3.7}$$

where $Ey_{t+(1,2)}$ is the expected forecasted value for the two respective time periods.

The forecasting can be extended into further time periods and for different models. What is important is that y_{t+2} is based on the previous forecast of y_{t+1}. Therefore, one needs to be careful in forecasting far ahead, these models are best suited for short term forecasting (Stata Corp, 2011).

3.2 Data

All data used are from open sources, and can be found in order to replicate and validate the results from this paper. The price data is from Nordpool and Montel and is in NOK/MWh. It consists mainly of prices from the NO1 price area, i.e. the Oslo area, but also system prices. The time span of the prices is from 3rd of January 2000 to 14th of April 2012. It is in a daily format, which means the 24 hourly prices of the day are averaged to 1 price a day. The calculation is performed and presented by Nordpool and Montel, and is available at their websites. Since the prices are calculated and reported once a day, we have chosen to predict daily prices. This because we want to find out how the prices are moving and predict tomorrow’s movement in the daily prices, since this gives an indication of the total trend for all the 24 hours.

The weather variables include temperature and precipitation and are collected from the Norwegian Meteorological Institute´s climate database, eklima.met.no. For temperature a distinction has been made for the location, where one observation represents the price area. They have been collected as daily temperatures reported by the institute. All temperatures are in Celsius. Regarding precipitation, it is a difficult variable to measure because of the high variation in regard to area differences. Further, it is usual for the observations to be discontinuous, i.e. have missing values. To try and correct for this, observation points has been chosen around price area NO1[1] and averaged with only the existing data for that observation. Precipitation data is not as complete and accurate as one would wish, which introduces an uncertainty problem.

[1] See appendix A for details concerning the observation locations.

It is difficult to collect perfect weather data, especially precipitation. There are often incomplete observations, and you only get historic data for one specific location. It would be preferable to have data for a specific area, which you can get forecasts for. This has not been possible to find, and therefore the option was to collect precipitation data for many locations around the area and calculate the average of the observation for that time. It was also important to include many locations since precipitation has great location variations and often missing values. With temperature it was chosen to have one location as a representation for the area. This was possible because the temperature observations do not have missing values as precipitation does.

All the data has been processed using Stata/SE 12.1 which is a statistical software tool. It has specific utilities for ARIMA, SARIMA, ARFIMA and ARMAX that has been used extensively. Excel 2007 has been used for combining data sets.

3.2.1 Descriptive Statistics

Electricity Prices

Summary statistics are presented in table 3.1. There is a big difference from the minimum and maximum value. This is due to a few extreme outliers in the prices, which are visual as spikes in the plot in figure 3.2. There is quite a difference from the maximum and minimum values of system prices and prices in NO1. Since it is the transfer capacities that create the difference between system price and the price areas, and system price is the one set first, it is logical that the system price vary less than the NO1 price.

Table 3.1 Summary statistics

Variable	Obs.	Mean	Std. Dev.	Min	Max
Price System	4486	278.62	123.98	31.85	1090.02
Price Oslo	4486	273.02	129.50	16.61	1226.45

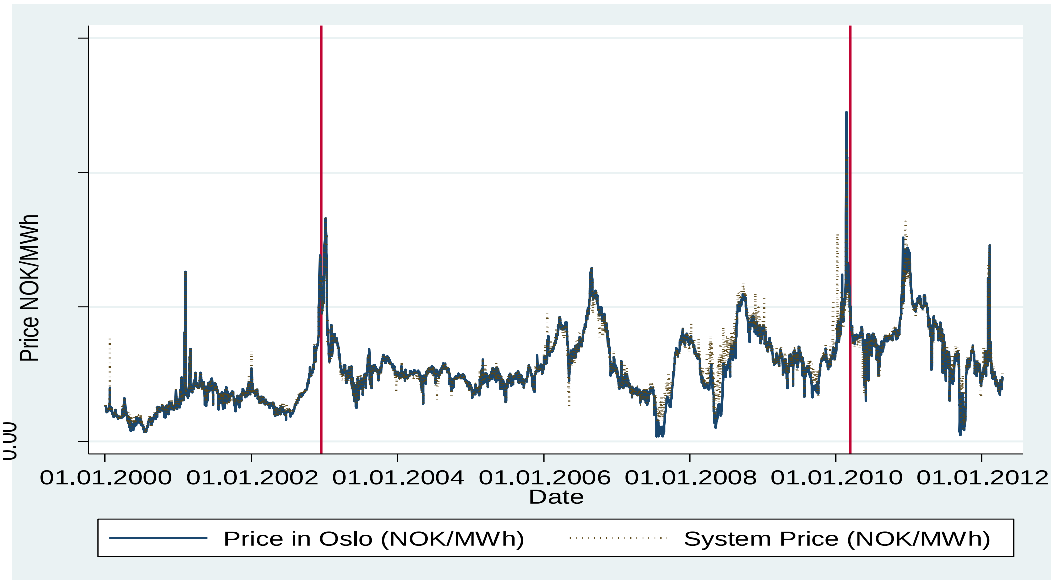

Figure 3.2 Plot of System and NO1 prices, 2000-2012

Figure 3.2 visualizes the price movements over the time span. The highest spike is from the winter of 2010, when prices rose above 1000 NOK/MWh, to 1226.45. Among the reasons for this was the shortage of water inflow from previous years, and exceptionally cold winter days. The lowest point was in 2007, when prices fell to only 16.61 NOK/MWh.

The plot includes both the system prices and the prices for NO1. The two follow each other most parts of the time span, but diverge during some parts. This is clearest around 2008 where system prices seem more volatile and to be higher than NO1. The vertical lines in the graph are changes to the price area NO1. In 2002 and 2010 the price areas were changed to be split up in more areas. These changes may have caused disturbances in the NO1 prices, but it should not affect the system prices, and therefore it could lead to greater differences between system and NO1 prices. However, it is not clear from the plot that these changes have affected the NO1 prices.

There are periods where the prices diverge from one another, and these periods appear more often after the changes. Having said that, there could be other reasons for the difference and after the last change in 2010, the prices are more similar. To further complicate the matter, there have been changes in import and export infrastructure. Knowing that these changes have been made within the selected time span and that they might affect the prices, it was still prioritized to have a long time series. We have therefore chosen to use the whole time series of January 2000 through April 2012.

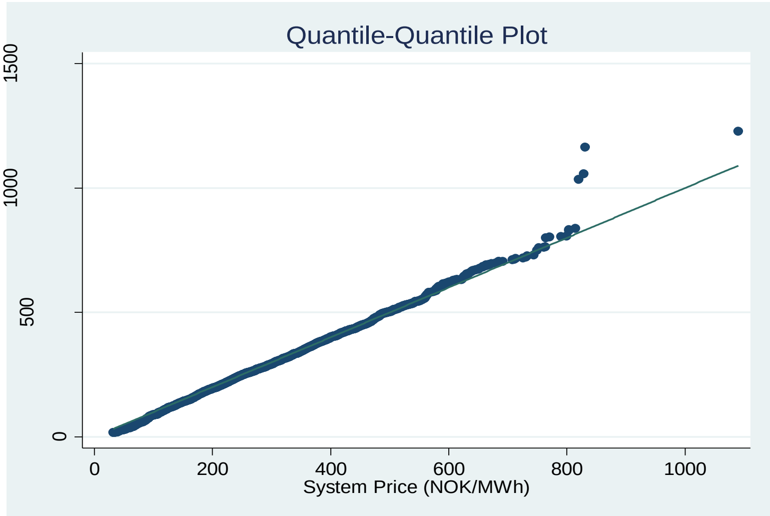

FIGURE 3.3 QQ-PLOT IF SYSTEM- AND NO1 PRICES

We have chosen to predict NO1 prices and not the system prices. But there is not necessarily a big difference between them, at least in relation to distributions. The Q-Q plot in figure 3.3 compares the NO1 and system price distributions. The four outliers of the higher order observation deviate from the distribution. Except from them, the observations lie along the linear line, suggesting similar distributions. A model that can fit NO1 prices can most likely fit system prices rather well.

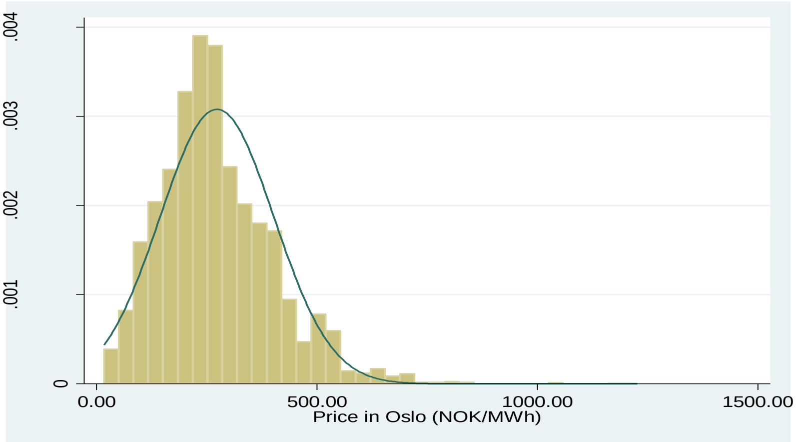

FIGURE 3.4 DISTRIBUTIONS OF THE NO1 PRICES

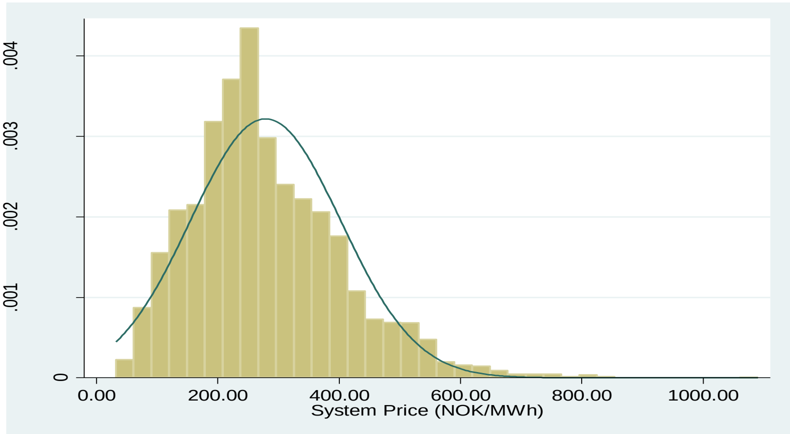

FIGURE 3.5 DISTRIBUTIONS OF THE SYSTEM PRICES

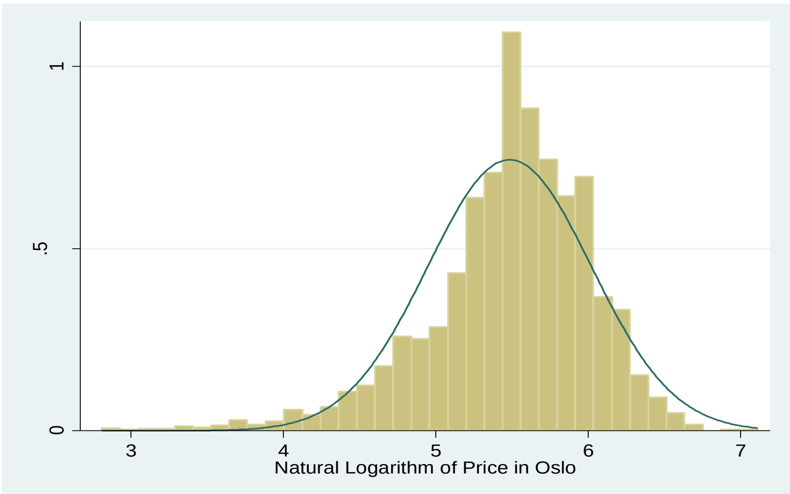

FIGURE 3.6 DISTRIBUTIONS OF THE LOG PRICES

The individual distributions are found in figure 3.4-6, where the price distributions are plotted together with the normal distribution curve. These indicate that the distributions are similar, but none of them follows the normal distribution. They are somewhat skewed towards the higher values of the prices. This is due to the spikes that often occur in the winter times. When testing if the natural logarithm might improve the distribution, it is shown that it still is skewed towards the higher values, but the tails are little more evened out. However, it still does not fit the normal distribution perfectly. The small improvement might indicate that log prices are more appropriate.

Explanatory variables

The explanatory variables are shown in table 3.2. The table displays the correlation between the variables and prices of NO1. These are only used as relative numbers, in order to compare and select. One temperature station is going to be used as a reference for temperature. Of these three, Rygge shows the highest correlation, although the difference is not massive compared to Gardermoen. Rygge will be the station referred to by temperature from here on.

TABLE 3.2 CORRELATION TABLE FOR PRICE NO1 AND WEATHER VARIABLES

Correlation	Price NO1	Description
Temp 1	-0,308	Gardermoen
Temp 2	-0,313	Rygge
Temp 3	-0,230	Blindern
Precipitation	-0,057	Averaged of several observations in NO1

Since temperature and precipitation are variables where you get forecasts for the next coming days, these are well suited to include in the forecasting of the next day prices. Temperature is a variable which affects the demand side. A decrease in temperature will increase the demand and therefore increase the prices. That is shown in the negative correlation. The same goes for precipitation, only it affects the supply side. When it rains the inflow to the water reservoirs will increase, therefore increasing the supply. When supply increases, all else equal, the prices will decrease giving rise to negative correlation.

For the explanatory variables to be valid for ARMAX, they need to be stationary and exogenous. Stationarity will be tested for in the next chapter. These weather variables are all independent from prices, since they are stochastic and determined by nature. The weather is not something electricity prices can affect.

4. Results and Analysis

In order to identify the best model, extensive testing has been performed. The most significant and relevant will be presented here. All of the mentioned ARMA-extensions have been applied in order to identify the best suited model for the electricity prices. The results will follow the Box-Jenkins method. First off is the identification of stationarity, seasonality, AR, and MA selection. It will be followed by parameter estimation and the test results from the most significant tests. Once a model is identified and selected it will be applied with weather variables, followed by forecasting of the prices.

4.1 Identification

Stationarity

The augmented Dickey-Fuller test will be conducted with different amounts of lags. There is a rule of thumb in order to determine the right lag-length (Schwert, 1989);

$$p_{max} = \left[12\left(\frac{T}{100}\right)^{\frac{1}{4}}\right] \qquad \text{Eq. 4.1}$$

where T = number of periods, e.g. years or months.

For this case, there are 4486 number of daily observations and with periods being weekly, it results in roughly 640 periods. The maximum lag-length is then approximated to 19. Other number of lags has also been tested. The null hypothesis for the test is unit-root, i.e. non-stationary. A small p-value will reject the hypothesis of non-stationarity.

Table 4.1 Augmented Dickey-Fuller Test for Unit Root

P-values ADF	Lags 5	Lags 5 w. trend	Lags 13	Lags 13 w. trend	Lags 19	Lags 19 w. trend
Price Oslo	0.0006	0.0004	0.0020	0.0018	0.0073	0.0089
Price System	0.0013	0.0006	0.0064	0.0048	0.0089	0.0071
Ln Oslo	0.0006	0.0004	0.0050	0.0057	0.0048	0.0050
Ln system	0.0051	0.0023	0.0273	0.0276	0.0223	0.0182

All the variables and all the lags show rejection of non-stationarity with 95% certainty. However, the natural logarithm for system prices shows higher p-values than the others. The test can include a trend part. When this has been included, the effect on Price Oslo and Price System was declining p-values. This may indicate that the time series is slightly trending.

TABLE 4.2 PHILLIPS-PERRON TEST FOR UNIT ROOT

P-values PP	Lags 5	Lags 13	Lags 19
Price Oslo	0.00	0.00	0.00
Price System	0.00	0.00	0.00
Ln System	0.00	0.00	0.00
Ln Oslo	0.00	0.00	0.00

To further investigate the question of stationarity, the Phillips-Perron test is also implemented. It has the same null hypothesis and it confirms the notion of stationarity. Table 4.3 indicates that the explaining variables are also stationary, which makes them valid.

TABLE 4.3 DICKEY-FULLER TEST OF THE EXPLANATORY VARIABLES

P-values DF	Lags 5	Lags 13	Lags 19
Temperature3	0.00	0.003	0.0024
Precipitation	0.00	0.00	0.00

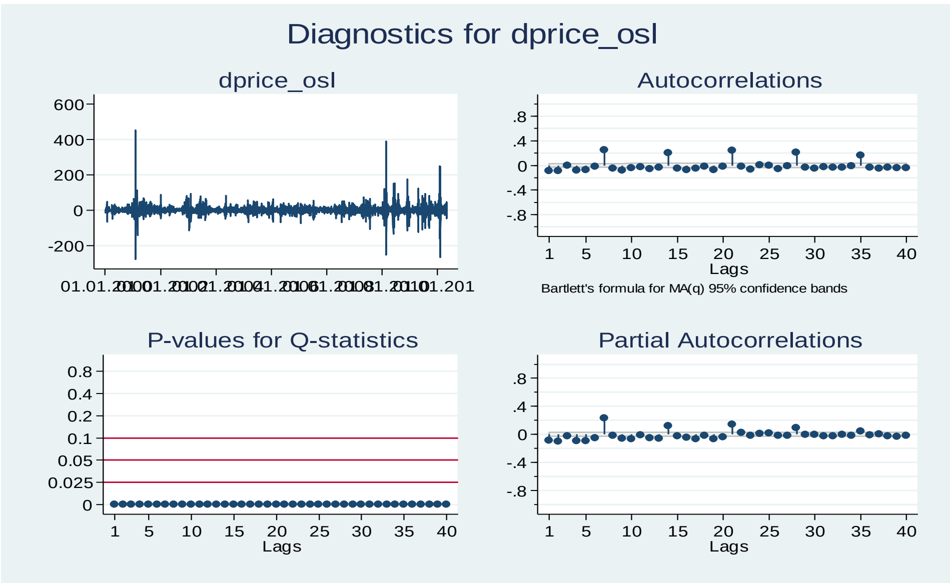

FIGURE 4.1 TIME PLOT, ACF AND PACF FOR PRICE OSLO I (1).

The above figure shows descriptive graphs for Price Oslo I (1), i.e. one differentiation. What is apparent is that the plot of the values shows a tendency of over-differencing, since they are consequently moving from positive to negative values. Also, there is no significant AR (1) spike (which is apparent in 4.2 further down) and the ACF does not decay as smoothly as it should. However, without the differentiation the ACF decays slowly (see figure 4.2). This could be a reason for trying an ARFIMA model, which will be done further in the process. From the above results there is no evidence suggesting that I (1) is needed.

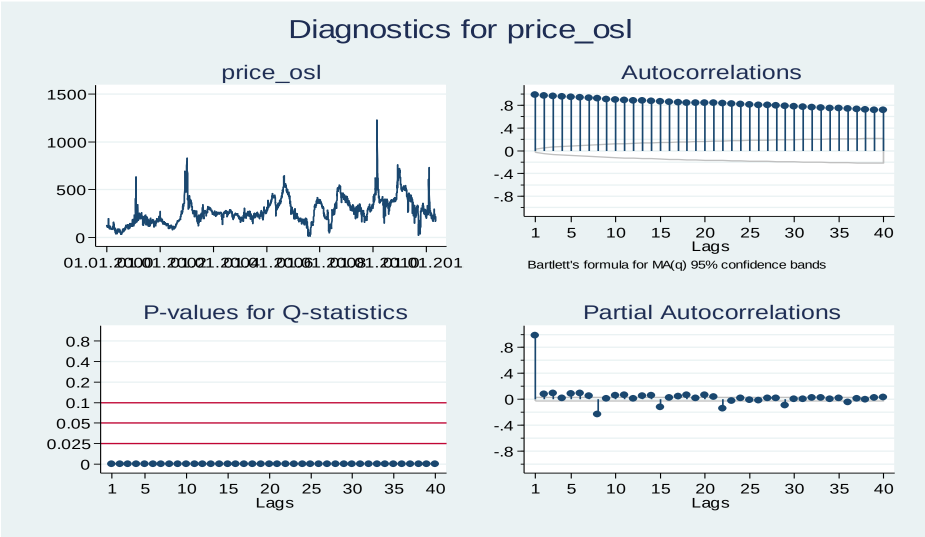

FIGURE 4.2 TIME PLOT, ACF AND PACF FOR NO1 PRICE

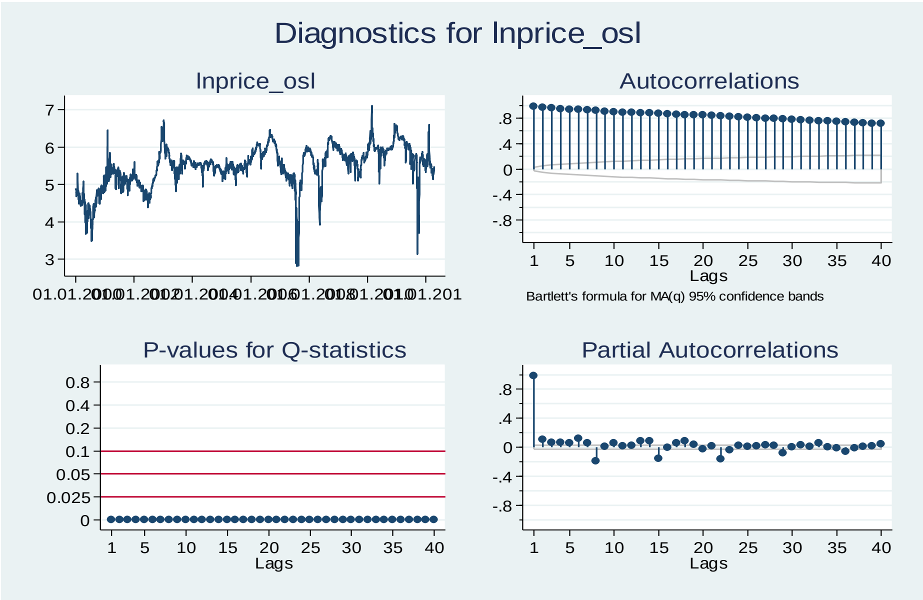

FIGURE 4.3 TIME PLOT, ACF AND PACF FOR LOG OF NO1PRICE

Seasonality

Figures 4.2 and 4.3 display spikes at every 7th lag in the partial autocorrelation. That goes for both regular price and log prices. That is also the case in figure 4.1 for the differenced prices. This is a clear sign of a weekly effect and needs to be controlled for. As described, this might be handled by the SARIMA model, where you include it with seven day seasonality. It might also be possible that the price experiences other seasonalities, such as monthly or yearly. A common prediction of electricity prices is that the winter prices are high (joined with brief spikes), and summer prices lower. However, the plot does not reveal any exact yearly pattern. From the winter of 2003 to 2007, there are not as many clear winter spikes as in the other years. It seems as the only seasonality which needs to be considered is weekly.

AR and MA selection

When it comes to identifying AR and MA terms, figure 4.2 and 4.3 shows the relevant information. The ACF is slowly decaying but the PACF shows a strong lag 1 and negative spikes at every 7th lag. These negative spikes show indication of weekly effects. This would imply at least AR (1) and controlling for the weekly spikes. Another alternative would be to including weekly seasonality in the model. This information has been used as the basis for finding a model and it is the same for log prices.

Many different models have been tried out with different AR, MA and model types, namely the SARIMA and ARFIMA. It is not possible to find the right model only by investigating the ACF and PACF. When a model is identified, analysis of the residuals can show if there is serial correlation left in any of the lags. This can then be corrected by expanding the model for those lags.

The main identified models are the following;

TABLE 4.4 ARMA IDENTIFICATION

#	Model description	Model Type
1	AR (1) MA (7, 14, 21)	ARMA
2	ARFIMA, AR (1, 2) MA (7, 14, 21)	ARFIMA
3	AR (1) MA (1, 7) SARIMA (1,0,1,7)	SARFIMA

The first model follows from the discussion on figure 4.2.When identifying the ARFIMA model it was found that it needed an additional AR-term for the second lag. The SARIMA was identified by the weekly effects clearly shown in previous results. Other SARIMA models were tested, but this gave the best results.

4.2 Parameter Estimation and Testing

Underneath, table 4.5 displays the results of the test scores. For details about the models coefficients etc, see Appendix B.

TABLE 4.5 ARMA RESULTS

#	AIC	BIC	Correlation^2
1	40532	40571	0,971
2	40520	40571	0,971
3	40153	40167	0,973

The pure ARMA model was natural from the looks of the spikes in ACF and PACF. However, the models residuals showed a spike at the 28th lag. When controlling for this lag, the model becomes too complex for standard calculations. Stata´s diffuse option was implemented to simplify the calculation though the score got worse than model 1. Standard ARMA does not seem to be able to control for all the serial correlation because of the spikes left in the residuals. There were also models with more AR-terms, but they gave insignificant parameters or a worse fit.

Model 2 is an ARFIMA model and it uses a fractional integration. The result for this parameter was 0.49 and significant, suggesting that some integration is in order. The AIC and BIC are not substantially different for this model then model 1. Although it seems as ARFIMA is suitable for modeling the prices, there is another model which can perform even better.

The last model is based on the SARIMA estimation. Different models were tested as well, but the one that stood out was model 3 (the others had either insignificant coefficients or worse scores). This one implements weekly seasonal adjustment with a seasonal AR and MA lag. It is clear that number 3 stands out as the best one, in terms of AIC and BIC, out of all the models. In Appendix B it can be verified that the model fulfills the requirements of AR (1) < 1 and the invertibility condition.

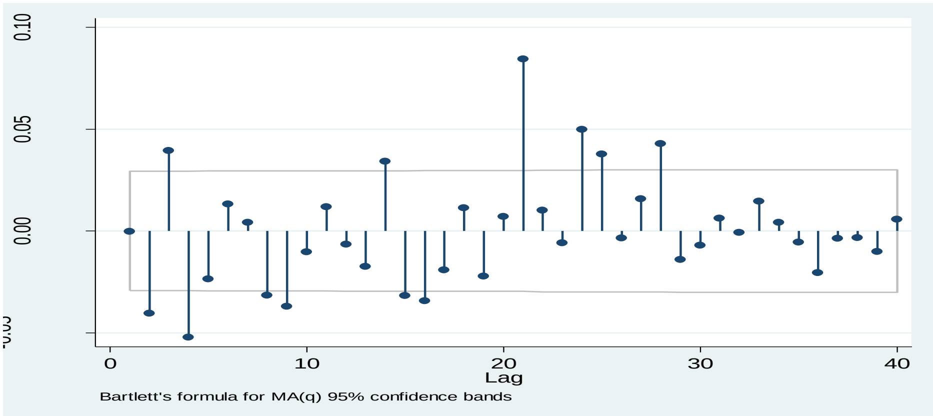

FIGURE 4.4 ACF - PLOT OF THE RESIDUALS

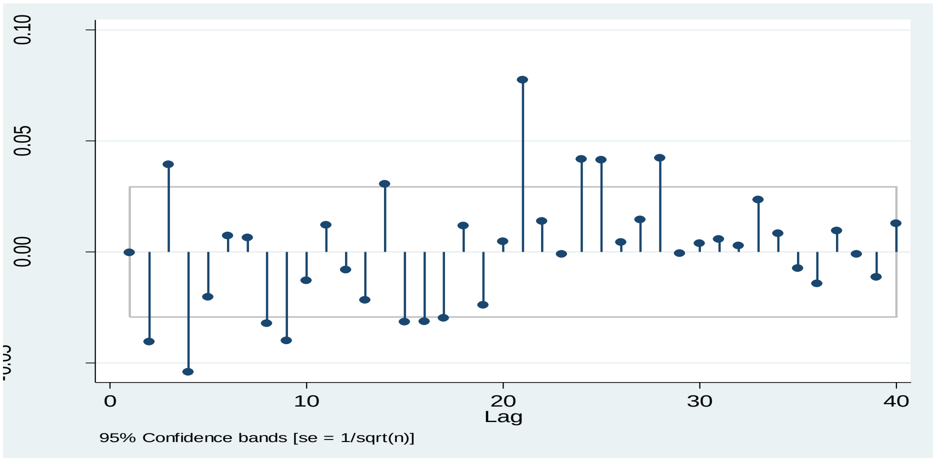

FIGURE 4.5 PACF - PLOT OF THE RESIDUALS

In order to find out if there is any serial correlation left in the model, it is helpful to use ACF and PACF of the residuals. Because of the spikes in figure 4.4 and 4.5, it seems as the model cannot completely describe all the serial correlation. However, it does describe a fair amount of it. Since some of the bigger spikes are located around at the 14^{th}, 21^{st}, and 28^{th} lag it appears as the weekly seasonality is not correctly accounted for. This could be helped by including a day of week variable, DOW. When the 21st lag was included in model 4 it was insignificant and also the model got difficult to calculate.

All the models so far have been estimated from real prices. But it was shown in the descriptive statistics that the logarithm of prices had a distribution slightly more similar to the normal distribution. It is difficult to compare test scores between the two different types, since the logarithmic model returns logarithmic values. This results in different AIC and BIC score levels. To compare the two, figure 4.6 shows the plot of the Oslo prices, predicted SARIMA prices, and the predicted logarithmic SARIMA prices (recalculated to normal prices).

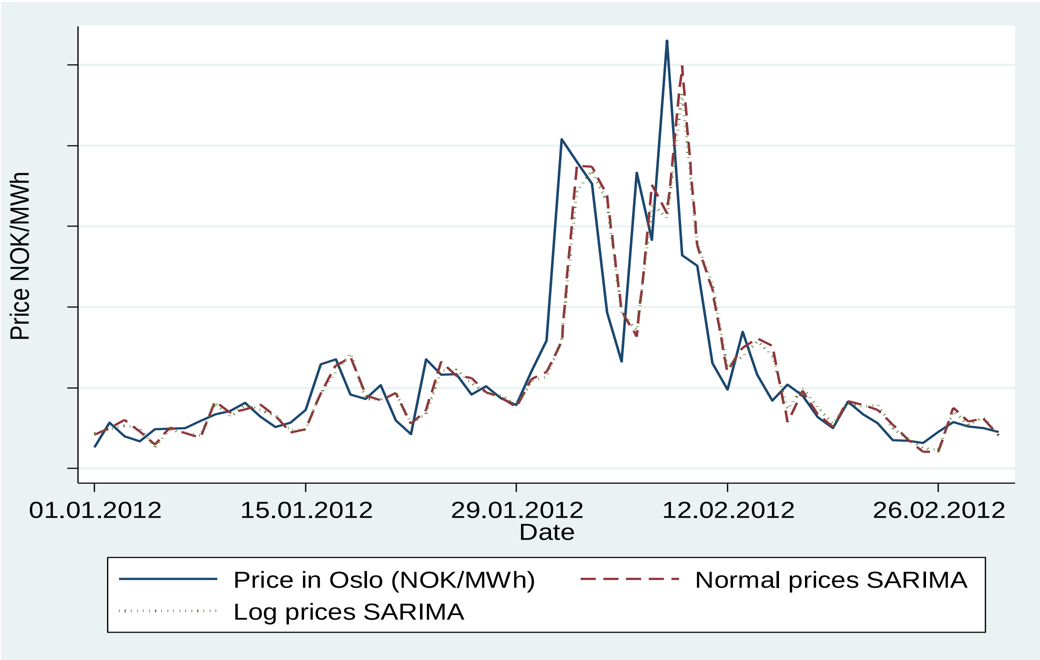

FIGURE 4.6 NORMAL PRICES VS. LOG PRICES, 01.01.2012-01.03.2012

The results are very similar and it is not easy determining which one is the best. Logarithmic regressions can be preferred when the coefficients are of interest since they give the change in percentage relative to dependent variable. In our case that is not the most important factor. Instead, we need the forecast in real prices. And it seems as though the logarithm of prices does not help to make a better forecast, which is why normal prices are going to be used for the proceeding part.

Model 3 is the best identified model, and this is the one which we have chosen to elaborate further in terms of adding explanatory variables. The variables temperature and precipitation are used to investigate if the forecasting performance of the model will improve with this weather information.

TABLE 4.6 SARIMAX RESULTS

Model description	AIC	BIC	Correlation^2
AR (1) MA (1, 7) SARIMAX (1,0,1,7) precipitation	40140	40192	0,973
AR (1) MA (1, 7) SARIMAX (1,0,1,7) temp	40050	40101	0,974
AR (1) MA (1, 7) SARIMAX (1,0,1,7) temp, precipitation	40043	40100	0,974
AR (1) MA (1, 7) SARIMAX (1,0,1,7) DOW	Not signifi		

The scores indicate an improvement from previous models. It appears that temperature and precipitation gives rise to a model with better fit and thus improves the overall model. The intercept in the model was not significant though. This can be a problem if one wants to use the base performance in the absence of the variables. Since we are not interested in only the intercept, this might be acceptable. Most of the decline seems to be originating from including temperature and not so much from precipitation. This is seen in the test scores for the models that incorporates the variables individually. The fact that temperature has a bigger impact on the test score could be expected since temperature has an direct effect on consumers, for example increase the heating in a household. Precipitation, on the other hand, can be stored for later use and therefore have a lagging effect. But since the water is saved, it should affect the water value and therefore have an direct effect of the expectations of tomorrow's prices. Judging from the results, it seems as temperature has a more direct effect on the prices.

When DOW was added to the model, it was insignificant, whereas it was not possible to include it in a final model. Therefore, some of the serial correlation from model 3 is still not explained. Even though precipitation helps the model only marginally, the one with both temperature and precipitation will be used for forecasting. The reason is because it is important to get information into the forecast equation. And since precipitation is a variable that can be forecasted, it is useful to implement it in the model, even though it only improves slightly. The coefficients for temperature and precipitation are -1.38 and -0.20, respectively. If temperature decreases with 1 Celsius the price should increase with 1.38 NOK/MWh, all else being equal. The same interpretation can be done with precipitation. The signs of the coefficients are as expected; if it rains the prices should decrease due to increased supply and if it gets warmer prices should decrease due to decreased demand.

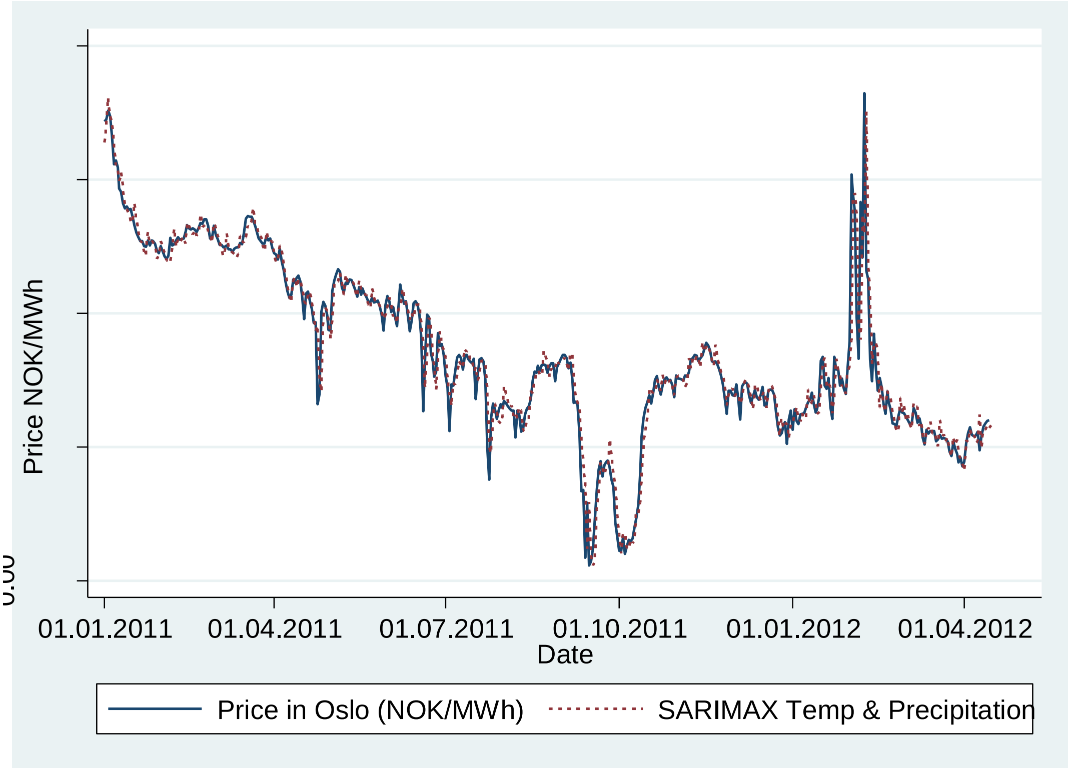

FIGURE 4.7 PLOT OF PRICES AND MODEL ESTIMATION, 01.01.2011-14.04.2012

By inspecting the plot of both the NO1 prices and the models predicted values in figure 4.7, it seems as though the model fits ok. It has difficulties with spikes, both high and low. Overall it follows the prices rather well.

4.3 FORECASTING

Now that the model has been identified and tested, it is used for forecasting the price. It is done by out-of-sample forecasting, as opposed to one-step-ahead predictions, where Stata include all the models parameters, including the explanatory variables for SARIMAX. The one-step-ahead uses all the information to predict the prices, while the out-of-sample only uses the previous price information from before the forecasting date. After this point, it bases further forecasts on the forecast itself. That is why the forecast gets more and more uncertain the further ahead it is done. This is a forecast Stata calls dynamic, where it needs to be specified for which date it will start forecasting. For that it needs observations for the explanatory variables in order to forecast. The forecasting is done for two separate dates, 2nd of April and 25th of January 2012. Those dates are the first forecasted values, the values before them are usual one-step-ahead predictions. The dates are influenced by different volatilities in order to test the models ability to forecast under different situations.

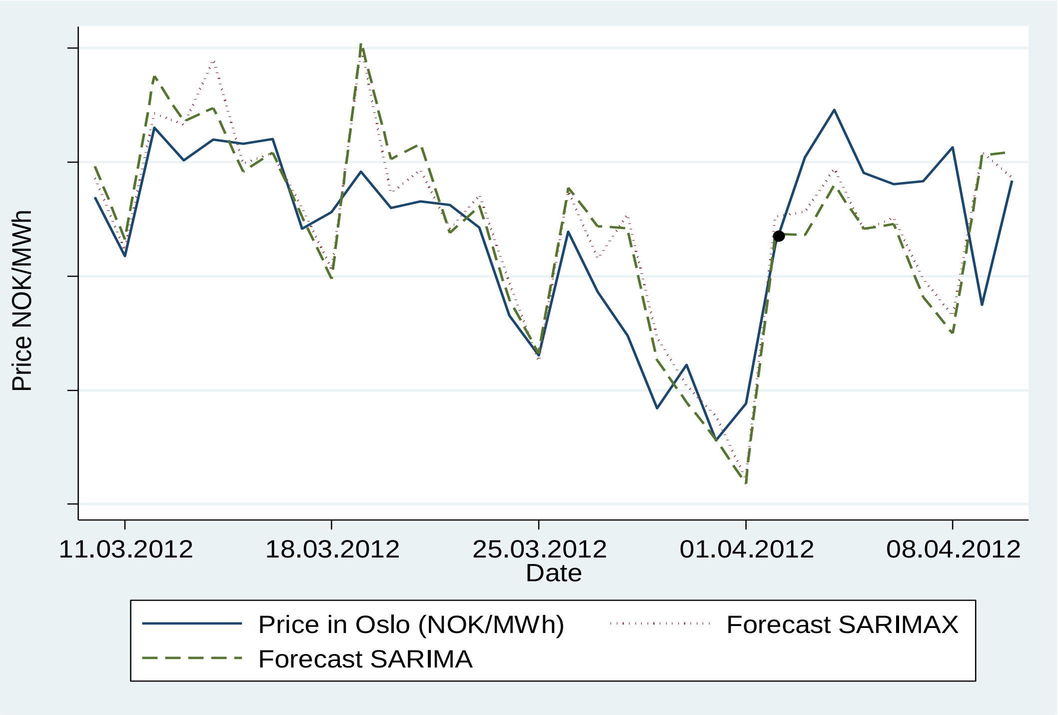

FIGURE 4.8 PLOT OF PRICES AND FORECASTS FOR 2ND OF APRIL 2012

The black dot in the figure above is the date of the first forecast. The forecast is quite accurate in this case. There is slight difference between the SARIMA model and SARIMAX. At the first forecast SARIMA outperforms SARIMAX. However, as the forecast continues the SARIMAX forecasts lies closer to the real prices, which give an indication that temperature and precipitation may increase the accuracy only somewhat. Even though the forecast gets the tendency right, it still under-shoots the real prices until the 6th when the forecast reacts opposite to the prices. The period for which the forecast is done is influenced by some volatility, and it seems as though the model has difficulties forecasting the right level in these circumstances. For the dates after the 2nd of April, the forecast still gets the tendencies somewhat right, though the under-shooting increases. As expected, the forecast worsens the longer it is projected. That is due to the fact that the information the models has gets more uncertain the further the forecast gets.

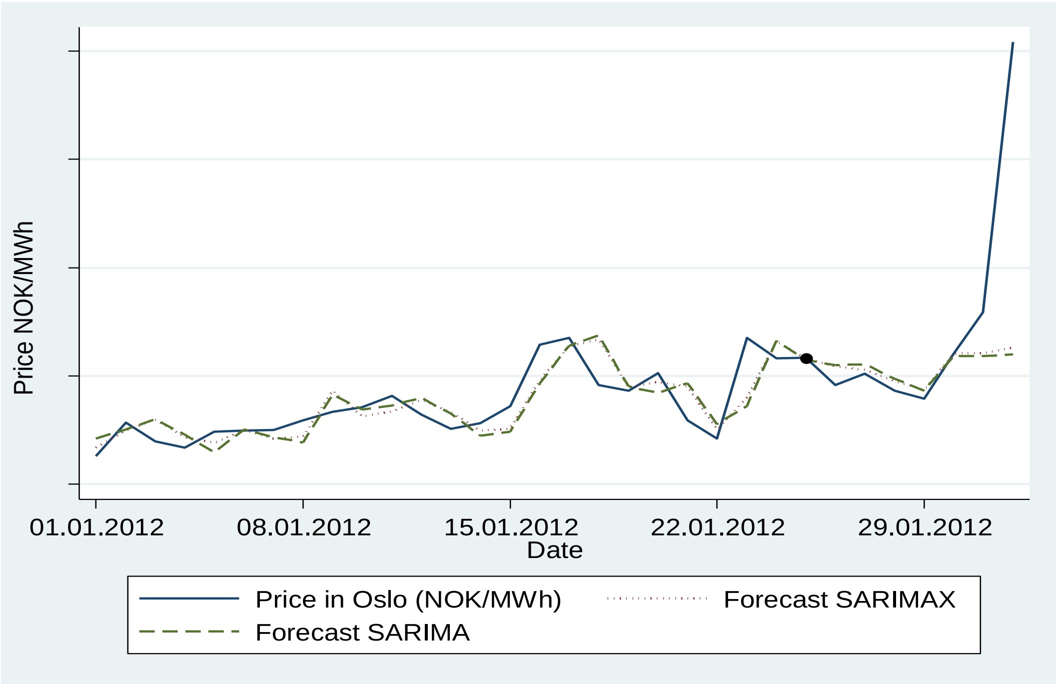

FIGURE 4.9 PLOT OF PRICES AND FORECAST FOR 25TH OF JANUARY

In order to test the model further, a different date was used for forecasting. In this scenario, there is a slight upward trend before the forecast and it ends in a fast increase, see figure 4.9. Both models are spot on for the first forecast. Then it diverge somewhat in the second one and follows rather well until the spike from 30th of January, which is five days after the first forecast. The forecast therefore does a good job for the nearest days, but experiences difficulties further ahead. There is only a small difference, barely visible, between the two different models indicating the lack of improvement from including temperature and precipitation.

In general the models seem to be able to forecast tomorrow's price rather well. It gives an ok indication of the electricity prices the next four to five days. However, they are not as accurate as the first forecast which is expected. The weather variables do not affect the forecasting for the next day, but appear to increases the certainty slightly for forecasts further away.

5. Conclusion

The objective for the thesis has been to identify an ARIMA model, use it to forecast tomorrow's electricity price and test whether explanatory variables such as temperature and precipitation can help improve the forecast. For the identification, it was shown that it was not needed to use differentiation in order to get the prices stationary. However, when using ARFIMA, it appeared as fractional integration was significant. The model did not show improved test scores and consequently it was not used any further. The basic ARMA model was not able to account for all the serial correlation from the weekly seasonality. Instead the SARIMA model was used, which improved the model considerably. However, even this model was not able to account for all the serial correlation. After testing several different models, it was concluded that it was not possible to account for all the serial correlation in the prices.

The SARIMA model was extended to include temperature and precipitation, and it improved the model only slightly. This is highly interesting since weather variables often are regarded as price drivers. One reason for this result is that the model is not able to extract all the weather information from these variables. The end results show that the model could not describe all the variation of the electricity prices, but a fair amount. There is room for improvement of this model by including more variables or by extending it to a more advanced seasonal model, to completely remove the serial correlation left in the residuals.

Another reason for the low improvement could be ascribed to the weather data. There is a big measurement obstacle when using weather as input. It is difficult to get historical data for areas; instead it is available only for single locations. And the data for precipitation often experiences missing values. The data was collected so that there were no missing values, but it is still difficult to get data that reflects the true outcome. Investing more time in weather data might lead to a better SARIMAX model. The one found in this thesis is not a significant improvement from the one without the weather variables. It could be argued that other data would improve the model, e.g. water reservoirs levels. The problem is that it is measured every week, while the prices used here are daily. Other possible variables could be import and export. In the future, if the amount of wind power manages to increase to a significant level, it could be of interest to test wind data.

The forecasts for the next day are relatively good. They give a fine estimation of the price to come. They also give a good indication of the prices for the next few days. When the forecast reaches four days, it starts to lose its ability to indicate movements and gets far off the real prices. From the forecasting results, it seems as the SARIMA (1, 0, 1 & 7) $(1, 0, 1)_7$ return valuable forecasts in short term. One needs to be cautious when using such a model in times with high volatility and possible price spikes. These periods are often during winter or when the electricity market experiences unusual situations, such as extreme low water reservoirs or similar. These situations are difficult to predict with this kind of model. Instead, it is important to know that during these periods, the model might not forecast as well as it should. For better forecasting of spikes, models such as GARCH might be appropriate.

References

Contreras, J., Espinola, R., Nogales, F. & Conejo, A., 2003. ARIMA Models to predict next-day electricity prices. *IEEE Transactions on power systems*, August, Volume 18.

Danish Energy Agency, 2011. *Energy Statistics 2010*, Copenhagen: Danish Energy Agency.

Enders, W., 1994. *Applied Enonometric Time Series.* s.l.:John Wiley & Sons, Inc.

Førsund, F. R., 2007. *Hydropower Economics.* New York: Springer Science+Business Media.

Fridolfsson, S.-O. & Tangerås, T. P., 2008. *Market Power in the Nordic Wholesale Electricity Market: A Survey of the Empirical Evidence*, Stockholm: Research Institute of Industrial Economics.

Haldrup, N. & Nielsen, M., 2004. *A regime switching long memory model for electricity prices*, Aarhus: University of Aarhus.

Hjelset, K. & Monsbakken, L., 2005. *Analyse av områdepriser på Nordpool*, Trondheim: NTNU.

Hughes, W. R. & Parece, A., 2002. The Economics of Price Spikes in Deregulated Power Markets. *The Elctricity Journal*, 15(6), pp. 31-44.

Mauritzen, J., 2010. *What happens when it´s windy in Denmark? An emperical analysis of wind power on price volatility in the nordic electricity market.*, s.l.: Norwegian School of Economics and Business Administration.

Ministry of Petroleum and Energy, 2003. *St.meld. nr 18 - Om forsyningssikkerheten for strøm mv.*, Oslo: Ministry of Petroleum and Energy.

Nau, R. F., 2005. *Decision 411-Identifying the numbers of AR or MA terms.* [Online]
Available at: http://www.duke.edu/~rnau/411arim3.htm#plots

Nord Pool Spot AS, 2012. *Nord Pool Spot.* [Online]
Available at: http://www.nordpoolspot.com/How-does-it-work/

NVE, 2010. *Energistatus*, Oslo: NVE.

Pereira, M., Campodónico, N. & Kelman, R., 1999. *Application of Stochastic Dual DP and extension to hydrothermal scheduling*, s.l.: PSRI.

Schwert, W. G., 1989. Test for unit roots: A monte carlo investigation. *Journal of Business and Ecnomic Statistics*, Volume 7, pp. 147-160.

Stata Corp, 2011. *Stata Release 12.* College Station, TX: Stata Press.

Statnett, 2011. *Statnett Sydvestlinken.* [Online]
Available at: http://www.statnett.no/no/Prosjekter/SydVestlinken/

Woldridge, J. M., 2009. *Introductory Econometrics - A Modern Approach.* 4:e ed. s.l.:South-Western Cengage Learning.

Appendix A

List of precipitation locations

Code	Name
RR_1080	Hvaler
RR_11710	Einavatn
RR_11900	Biri
RR_12680	Lillehammer - Sætherengen
RR_17000	Strømtangen Fyr
RR_17150	Rygge
RR_17850	Ås
RR_18950	Tryvannshøgda
RR_20250	Hole
RR_22730	Hedal i Valdres 2
RR_22790	Grimsrud i Bergnadalen
RR_22840	Reinli
RR_24210	Sokna 2
RR_24600	Grimeli i Krødsherad
RR_25100	Hemsedal
RR_26240	Hiåsen
RR_26380	Eggedal 3
RR_26990	Sande - Galleberg
RR_29350	Uvdal Kraftverk
RR_29600	Tunhovd
RR_30860	Bergeligrend
RR_3780	Igsi i Hobøl
RR_4780	Gardermoen
RR_6620	Elverum-Fagertun

APPENDIX B

Model 1, AR (1) MA (7 14 21)

Sample: 03.01.2000 - 14.04.2012 | Number of obs = 4486
Wald chi2(4) = 386381.90
Log likelihood = -20260.05 | Prob > chi2 = 0.0000

price_osl	Coef.	OPG Std. Err.	z	P>|z|	[95% Conf.	Interval]
price_osl						
_cons	271.7148	20.20228	13.45	0.000	232.119	311.3105
ARMA						
ar						
L1.	.9710457	.0016651	583.17	0.000	.9677822	.9743093
ma						
L7.	.1822326	.0063138	28.86	0.000	.1698577	.1946075
L14.	.1250645	.0088371	14.15	0.000	.1077441	.1423848
L21.	.1673936	.009241	18.11	0.000	.1492815	.1855057
/sigma	22.12841	.0460456	480.58	0.000	22.03816	22.21866

Model 2, ARFIMA AR (1 2) MA (7 14 21)

Sample: 03.01.2000 - 14.04.2012 | Number of obs = 4486
Wald chi2(6) = 12738.47
Log likelihood = -20251.936 | Prob > chi2 = 0.0000

price_osl	Coef.	OIM Std. Err.	z	P>|z|	[95% Conf.	Interval]
price_osl						
_cons	253.3728	398.4846	0.64	0.525	-527.6427	1034.388
ARFIMA						
ar						
L1.	.4159697	.0167726	24.80	0.000	.383096	.4488434
L2.	.1374142	.0149336	9.20	0.000	.1081449	.1666835
ma						
L7.	.2126717	.0152427	13.95	0.000	.1827967	.2425468
L14.	.1357014	.0133707	10.15	0.000	.1094953	.1619074
L21.	.1698073	.0144853	11.72	0.000	.1414166	.198198
d	.4947843	.0073467	67.35	0.000	.4803851	.5091834
/sigma2	487.485	10.2943	47.35	0.000	467.3085	507.6614

Model 3, SARIMA

Sample: 03.01.2000 - 14.04.2012	Number of obs	= 4486
	Wald chi2(5)	= 2.59e+06
Log likelihood = -20069.3	Prob > chi2	= 0.0000

price_osl	Coef.	OPG Std. Err.	z	P>|z|	[95% Conf.	Interval]
price_osl						
_cons	236.6251	266.4311	0.89	0.374	-285.5703	758.8204
ARMA						
ar						
L1.	.9790819	.0014416	679.17	0.000	.9762565	.9819074
ma						
L1.	-.0799693	.0031882	-25.08	0.000	-.0862181	-.0737205
L7.	.0980272	.0068301	14.35	0.000	.0846405	.1114139
ARMA7						
ar						
L1.	.9990247	.0008398	1189.55	0.000	.9973787	1.000671
ma						
L1.	-.9795105	.0036593	-267.68	0.000	-.9866826	-.9723383
/sigma	21.18004	.0417818	506.92	0.000	21.09815	21.26193

SARIMAX

Sample: 03.01.2000 - 14.04.2012	Number of obs	= 4486
	Wald chi2(7)	= 3.11e+06
Log likelihood = -20012.35	Prob > chi2	= 0.0000

price_osl	Coef.	OPG Std. Err.	z	P>|z|	[95% Conf.	Interval]
price_osl						
temp2	-1.377478	.1150789	-11.97	0.000	-1.603028	-1.151928
precip	-.205082	.0787688	-2.60	0.009	-.3594661	-.0506979
_cons	246.6504	263.2051	0.94	0.349	-269.2221	762.523
ARMA						
ar						
L1.	.9790418	.0014525	674.03	0.000	.9761949	.9818887
ma						
L1.	-.0980776	.0032678	-30.01	0.000	-.1044823	-.0916729
L7.	.1050342	.0069418	15.13	0.000	.0914285	.11864
ARMA7						
ar						
L1.	.999186	.0007488	1334.43	0.000	.9977184	1.000654
ma						
L1.	-.9811066	.003573	-274.59	0.000	-.9881096	-.9741035
/sigma	20.91133	.0425319	491.66	0.000	20.82797	20.99469