Project Number 260041

SUPPORTING ACTION

EnRiMa
Energy Efficiency and Risk Management in Public Buildings

D3.1 Model description, calibration, and validation

Start date of the project: 2010-10-01
Duration: 42 months
Organisation name of lead contractor for this deliverable: SINTEF
Lead authors: SINTEF, SU
Revision: 2.0, 2012-01-23

Project funded by the European Commission within the Seventh Framework Programme (2007-2013)		
Dissemination Level		
PU	Public	**x**
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Executive summary

The decision support system (DSS) to be developed in the EnRiMa project is built around two optimization models, one focused on operational models and one on strategic ones. The quality of the solutions suggested by these models depends heavily on the quality of their input parameters. The preceding deliverables D1.1 and D4.1 in work packages 1 and 4, respectively, emphasized that some of these input data are subject to uncertainty and identified the key uncertainties for the selected test sites.

The role of WP3 is to provide data for the stochastic parameters in the form of scenarios, i.e., a discretization of the underlying stochastic distributions and/or processes. The present report, deliverable D3.1, describes the first steps toward this goal: collection and analysis of relevant data. The subsequent deliverable D3.2 will be concerned with the process of actually creating scenarios for the model input parameters.

After a brief introduction, we start by describing the two optimization models in Section 2. We describe their requirements with respect to stochastic data, explaining the differences between the operational and strategic models. Since the strategic model is planned as a two-stage optimization problem, there is no need for a full time-series analysis for these parameters. What is needed instead will depend on the actual time scale of the strategic model, which is still being decided upon. We will therefore come back to these parameters in the forthcoming deliverable D3.2.

Even though the main focus of this report is on time-series analysis of historical data, we explain that some of the stochastic parameters cannot be addressed properly using this approach. These parameters are, therefore, treated separately in Section 3. Without them, and having excluded the strategic parameters in the previous step, the only parameters we need to do the time-series analysis for are hourly electricity prices.

The last step before the time-series analysis itself is data collection. In Section 4, we, thus, present all relevant data sources. Section 5 then presents the main results of this report, namely the time-series analysis of the electricity prices. Since the project's test sites are in Austria and Spain, we had to analyze prices in these two countries. We have tried two different frameworks, one based on mean-reverting stochastic processes and the other on time-series models. For both of them, and for both countries, we tried to model either the whole hourly series at once, or to split it into 24 series, one for each hour of the day. Since this would be too many results to present and the mean-reverting framework turned out to be inferior to the one based on time series, we present only results of the latter one in this report. In addition, we include one of the mean-reverting reports in the appendix, as it might work better for some other data series.

In the presented time-series based framework, we used a combination of seasonal autoregressive moving-average (seasonal ARMA, or SARMA) and generalized autoregressive conditional-heteroskedasticity (GARCH) models for the series of log-returns of the prices, for both countries. In both cases, modelling the whole series at once was, on the whole, better than splitting it into 24 per-hour series. We have also observed that, for both countries, the modelling errors are largest for the early morning hours.

It should be pointed out that the optimization models, for which we analyze the data, are still under development, so it is possible that some additional data analysis will be needed. In such a case, we would include results of this analysis in some suitable later EnRiMa report.

Contents

List of Figures

List of Tables

List of acronyms

ACF autocorrelation function

AIC Akaike information criterion

ARMA autoregressive moving-average (model)

BIC bayesian information criterion

DSS decision support system

EXAA Energy Exchange Austria

GARCH generalized autoregressive conditional-heteroskedasticity (model)

MAE mean absolute error

MAPE mean absolute percentage error

RMSE root mean squared error

PV photovoltaic

PACF partial autocorrelation function

SARMA seasonal autoregressive moving-average (model)

WPx work package x

1. Introduction

The goal of work package 3 (WP3) is specified as operational goal O2 in the description of work:

> We will model energy prices and loads to use as part of the stochastic model within the DSS Engine. Characteristics such as seasonality, short-term mean reversion, and weather fluctuations will be considered in generating short- and long-term price scenarios. By scenarios, we refer to a distribution of random parameters and not just point estimates. Such forecasts are used for the stochastic optimisation in Objective O3, where we want to find decisions that will be robust in all possible futures, not just the most probable one.

In other words, the goal of this package is to provide the stochastic models developed in WP4 with data for their stochastic parameters, in a format suitable for the models. Since we use stochastic programming as our optimization tool, the "suitable format" means a discretization of the underlying stochastic distributions and/or processes, usually referred to as *scenarios* (see, for example, Kall and Wallace, 1994; Birge and Louveaux, 1997; Dupačová et al., 2000; Høyland and Wallace, 2001). The work needed to achieve the goal can be divided in the following steps:

1. Collect all the required data
2. Perform time-series analysis of the data, to extract trends, seasonal effects, etc.
   - this means identifying and fitting a suitable time-series model
   - the result is a set of models, their parameters, and the historical residuals (differences between the predicted and observed values)
3. Generate scenarios for the *residuals*
4. Use the time-series models to convert the scenarios of the residuals into scenarios of the actual values

Deliverable D3.1, described in this report, covers the first two steps of the sequence, while the rest will be done in deliverable D3.2. In addition, there are some stochastic parameters that do not fit the above framework, mostly because we do not need to—or cannot—do the time-series analysis. As an example consider the weather, for which we can use external forecasts instead of doing the analysis ourselves. These parameters are treated in more detail in Section 3 of this report.

1.1. Relation to WP4

The goal of WP3 is to produce input data for the optimization models developed in WP4, in a format suitable for the models developed for the EnRiMa decision-support systems (DSS)s. This report is based on deliverable D4.1, which includes specifications for stochastic parameters for the models. It is, however, important to realize that the specifications can be expected to change as the models get developed and tested; we might find out that some of the parameters are not necessary, or that we have missed some, or that we have to change the time scale of the model, etc. In such a case, we would have to re-run the time-series analysis presented in this report. Nevertheless, the foundation of the methodology will have been laid.

This all implies that this report describes only an initial version of the scenario-generation toolbox, which will have to be further adjusted as the optimization models get developed and tested. Such changes would then be included in a later report.

The rest of the report is organized as follows: in the next section, we explain the differences between the operational and strategic optimization models, as they have different data requirements. Section 3 then describes stochastic parameters that for some reason need special treatment, while Section 4 describes the data sources used needed for the deliverable. Finally, Section 5 presents the main results of this deliverable: for each data series that needs analyzing, we include a report in which we *describe* the time-series model used for that series, *calibrate* (fit) it to the data and then *validate* the results using several goodness-of-fit measures. The paper is then concluded in Section 6.

2. Operational vs. strategic model

The goal of WP4 is to develop a model based DSS for operators of energy facilities. The DSS will be composed of two modules, operational and strategic, each based on a corresponding optimization model. The *operational model* will be used for planning the operation of the installed devices for the next period. It is expected to have a time horizon of between one day and one week, with hourly resolution. Since the operational model does not consider any changes to the installed equipment or building structure, these are all taken as input data. The *strategic model*, on the other hand, will be used for deciding investments into new equipment. This model is planned to have a horizon of several years, possibly up to ten or even twenty years.

2.1. Operational model

The operational model is expected to work with one-hour time steps and a horizon of up to one week. At the moment, all the test sites have long-term contracts for all energy sources (electricity, gas, heat), so the corresponding parameters of the operational model are assumed to be known. This means that the main operational uncertainties are occupancy and weather, with the latter being treated in a special way described in Section 3.1.

As for occupancy, none of the sites seems to have actual data that we could analyze, so the different occupancy patterns will have to be estimated in co-operation with the site managers, by dividing users into groups and establishing user profiles for each of them (expert opinion). The occupancy scenarios will then be converted into energy loads using the LoadCalc tool, which we describe in more detail in Section 3.2. This all implies that as long as we have fixed energy prices, there are no historical data series to be analyzed for the operational model, at least not in context of this deliverable[1].

[1] Deliverable D3.1 is about time-series analysis, i.e. analyzing data for which we have a whole series of historical values. Handling of the other parameters will be described in deliverable D3.2.

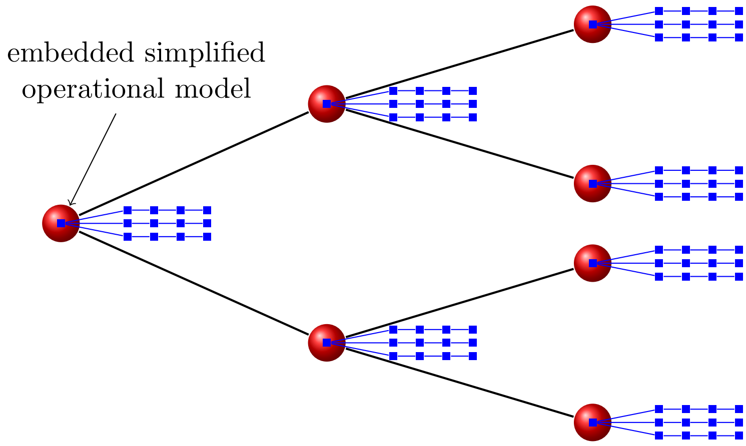

Figure 1: The strategic model, with operational models embedded at each node. The strategic tree, represented by the circular nodes, has period lengths of months, or even years. The period length of the embedded trees, corresponding to the operational model, is in hours; each scenario represents a 'typical' or 'critical' day/week of the corresponding strategic period.

Real-time pricing

Since it is expected that real-time pricing of electricity will become available throughout Europe in the near future, it is important to test the effect that this would have on the operation of the sites. We, thus, plan to test our models under the assumption of real-time pricing, which means that we have to collect and analyze the electricity spot prices, with at least hourly resolution.

2.2. Strategic model

The goal of the strategic model is to support long-term decision making. The model has to deal with two seemingly contradictory requirements: by definition, it has to be long term, with a horizon of at least several years (up to twenty years). On the other hand, the model has to be able to evaluate how the infrastructure it proposes performs on a daily basis. For this, the model should have a time resolution close to the operational model, i.e. on the scale of hours—yet having hourly resolution in a model that spans ten or twenty years is simply not realistic. Fortunately, it is not needed, either: the strategic decisions are clearly not made every hour, not even every day or week. In other words, the hourly resolution is needed only to evaluate the performance of the strategic decisions. Consequently, the scenario tree for the strategic *decisions* can have period lengths of months or years, as long as we can *evaluate* it with the resolution required by the operational model.

Our solution is a kind of "hybrid tree", presented in Figure 1: a tree for the strategic decisions, where every strategic-decision node (circular nodes in the figure) includes the operational model, or more likely a simplified version of it, for evaluation. Ideally, the operational model should be run for every day or week for the strategic period, with several scenarios for each day (week)—which is obviously not possible. Instead, the idea is to select a small set of 'representative' and 'critical' days (weeks), with assigned probabilities. An example of the former would be a "typical" summer day, while an example of the latter would be an "extremely cold" winter day. The number of these operational scenarios will depend on the solution time of the strategic model and will therefore be determined first when the model gets implemented (in deliverable D4.3 of WP4). In Figure 1, there are three such operational scenarios, denoted by the small rectangular nodes, for each strategic node.

Since the strategic model is expected to have a horizon of several decades, we will have to model the way the parameters of these typical and critical days (weeks) evolve during the time scale of the model. In addition, there are stochastic parameters specific to the strategic model itself, i.e. parameters influencing the infrastructure decisions. These include prices and performance characteristics of the available technologies as well as changes in the tariffs' structure (such as availability of time-of-use and real-time tariffs). Most of these parameters are close to impossible to estimate using a time-series analysis of historical data, since they depend on many external factors: new inventions in the case of future technology parameters or political decisions for the electricity tariffs. Fortunately, there are third-party predictions for many of these parameters, which we intend to use wherever possible. This topic will be further discussed in the forthcoming D3.2 report.

Furthermore, the strategic model is planned to be a two-stage stochastic program, which means that we have to estimate the stochastic parameters only *one period ahead*. This implies that the parameters do not need to be modelled as stochastic *processes*, but as *distributions*—which means that we do not need to do time-series analysis for them. This places these parameters outside the scope of this report; we will come back to them in the forthcoming D3.2 report.

This all implies that the only data series for which we have to do time-series analysis are the hourly electricity prices, needed by the operational model. Results of these analyses are presented in Section 5.

3. Stochastic parameters requiring special treatment

Not all stochastic parameters of the optimization models can be treated in the way described above: for some parameters, we do not have data for the values directly, so they have to be inferred from other observable data. Other parameters might have data available, but there might be better ways of estimating their future development than doing data analysis. In this section, we list these parameters and describe how to treat them.

3.1. Weather

By weather, we mean the following set of parameters of the optimization models: outside temperature, humidity, wind speed (and possibly also direction) and solar irradiation (the amount of sunlight, needed to compute the output of solar technologies). While there are historical data available for most of the parameters, we do not need to use data analysis for their prediction for the operational models; instead, we will use weather forecasts, which should be better than any time-series model we can hope to develop. After all, meteorologists use some of the most powerful computers for a reason.

Using weather forecasts is straightforward if we want to trust them completely and use them as a deterministic predictor of the weather. If we, on the other hand, want to have weather as a stochastic parameter in the optimization models, then there is no obviously correct way of doing this. We have identified the following options:

1. Get a weather forecast with some information about uncertainty of at least some of the parameters. Such forecasts are available, for example, from the European Centre for Medium-Range Weather Forecasts[2], but they are quite expensive, so there is little chance that the sites would be willing to subscribe for such service. The only stochastic information in freely available weather forecasts is the precipitation probability, which is not enough for our purpose.

2. Get some information about the error in weather forecasts, for example, by comparing historical forecasts with actually observed weather. The critical issue is the availability of the historical forecasts: none of the freely available forecasts seems to provide this information, so we would probably have to start collecting the forecasts now and do the analysis at some time (perhaps one year) in the future. Another issue with this approach is that the forecasting error is not constant: in some situations, a weather forecast may be relatively certain even several days ahead, while another time even the next day's forecast might be uncertain.

3. Finally, we could take advantage of the fact that there are several freely available weather forecasts and simply use each of them as one possible scenario of the future weather. We would just have to check that they use different sources, to avoid using the same forecasts several times.[3] The advantage of this approach is its simplicity, the downside is that it is likely to underestimate the uncertainty of the weather, as all the forecasts present what they believe is the most likely outcome.

After some deliberation, we have decided to use the last option, since it is the one most readily available. Then, once the system is in place, we should start collecting the weather forecasts and weather data. This will allow us to estimate the distributions of the errors these forecasts make throughout the year. With this data in place, we should be able to use the second way of getting stochastic forecasts, or possibly even to combine the last two approaches.

Forecasting output of solar technologies

In addition to forecasting weather parameters, we also need to be able to convert the weather forecasts into forecasts of the output of solar technologies such as photovoltaic (PV) and solar heating. This requires quite complex calculations, since the output depends on the latitude of the location, direction and inclination of the panels, cloud cover, temperature, date and time, to name at least the main parameters.

[2]See http://www.ecmwf.int/, with an example forecast at http://www.ecmwf.int/products/d/sampler/epsgrams/europe/page.html

[3]Though we cannot expect complete independence of the forecasts, since many European weather offices base their forecasts on the same global forecast, but process it differently.

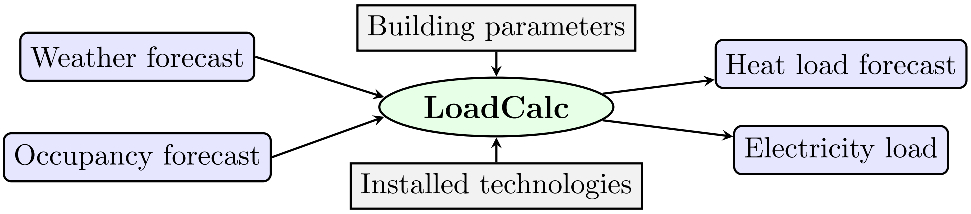

Figure 2: Using LoadCalc to forecast the energy load (electricity and heat).

Fortunately, there is the Photovoltaic Geographic Information System (PVGIS[4]), which includes an European solar radiation database with grid resolution of 1 km × 1 km. PVGIS includes also an online tool[5] to calculate the output of PV technologies at any location, taking into account local geographical conditions (mountains etc). The output can be either cumulative per year or month, or in the form of typical daily profiles. There is also a description of how to calculate the outputs[6], based on Šúri et al. (2007). These formulae can be used to calculate PV outputs for the operational model, taking into account current forecast for temperature and cloud cover.

3.2. Energy loads

Energy loads of the buildings are perhaps the most important stochastic parameters of the optimization models. They are also an example of parameters for which there are no historical data for estimation, as the energy consumption is measured annually, while the operational model has hourly resolution.

Even if we had the data available, it is not sure how useful they would be. The reason is that the energy loads are themselves a function of several other parameters, the most important ones being building occupancy and outside temperature. Using historical data for the loads would, by definition, be based on historical weather and occupancy, so it would not allow us to take into account the current weather forecast or expected occupancy.

The solution is to forecast, and generate scenarios for, the underlying parameters and then calculate the load from them. This calculation then forms an additional step that can be seen either as a post-process of the scenario generation, or as a pre-process of the optimization.

The calculation of energy loads from the parameters is done using LoadCalc, a building-simulation model developed at CET, Austria. LoadCalc creates a simplified model of a building, based on German norms DIN 18599 (draft) and VDI 6020 (Balada et al., 2007; Gebäudeausrüstung, 2001). This implies the following simplifications: room air temperature is constant within the entire operative zone, calculations are made for a stationary situation, the building is regarded as a cuboid or cube, and heat flow is normal to the given surfaces. This makes the energy flows in the building tractable, while keeping the model sufficiently accurate to meet the requirements of the EnRiMa project. For more information about LoadCalc, see Groissböck et al. (2011).

[4]See http://re.jrc.ec.europa.eu/pvgis/.
[5]See http://re.jrc.ec.europa.eu/pvgis/apps4/pvest.php.
[6]See http://re.jrc.ec.europa.eu/pvgis/solres/solmod3.htm.

The goal of LoadCalc is to obtain hourly load data for cooling, heating, and electricity, used primarily for the strategic optimization model. To do this, it needs the following input data: weather (temperature, humidity and wind), heat transfer and heat transfer coefficient, external loads (solar radiation), internal loads (people, working machines), natural infiltration, air ventilation systems (without humidity constraints), and business and non-business hours. Since LoadCalc is a deterministic tool, it will be used for every scenario in the way illustrated in Figure 2: the weather and occupancy forecasts for a given scenario are used to generate forecasts for electricity and heat demands for the scenario.

4. Available data sources

In this section, we describe data sources for estimating the stochastic parameters. We include also sources for the data on which we will not run a time-series analysis, such as weather forecasts. The data are divided with respect to the model type to which they belong.

4.1. Data for the operational models

As we have described in Section 2.1, the main operational uncertainties are occupancy, weather, and, in case of real-time pricing, electricity prices. Unfortunately, none of the sites has historical occupancy data, so these values would have to be estimated in a co-operation with building operators—which is outside of the scope of this deliverable.

Weather

There are several online weather forecast providers, whose forecasts differ not only in actual values but also in the parameters they forecast and the granularity and length of the forecasts. We have selected services that provide weather forecasts of at most 3-hour granularity for all our test sites.

`http://www.weather24.com/` is a German site that provides one-week ahead forecasts with three-hour resolution, including precipitation probability and amount; plus 16-day forecasts without precipitation data.

`http://www.yr.no/` is a Norwegian site providing 48-hour-ahead forecasts with three-hour resolution, including cloud coverage (split between fog, low, middle, and high-clouds); plus a nine-day forecast with precipitation data (only amount).

`http://www.weather.com/weather/` is an US-based weather site with a more limited selection of European locations (it does not, for example, have forecasts directly for Pinkafeld); otherwise, it provides three-day forecasts with one-hour resolution and ten-day forecasts with half-a-day resolution, both including precipitation probability.

`http://www.aemet.es/` is the site of the Spanish meteorological agency (Agencia Estatal de Meteorología, AEMET), limited to forecasts for Spain. It provides forecasts one week ahead, with four values per day for the first two days, two per day for the next two, and one for the last three days; all these include precipitation probability.

Electricity prices

For Austria, we can get historical spot electricity prices from Energy Exchange Austria (EXAA[7]). They provide data from March 22, 2002, with one hour granularity, plus daily averages, peak and off-peak prices, and prices for eleven other intervals. (Those intervals are partly overlapping: for example, the period 01–08 hours can be divided either into 'dream' (01–06) plus 'wakeup' (07–08), or into 'moon' (01–04) and 'sun' (05–08).)

For Spain, we have historical spot electricity prices with hourly resolution, starting from January 1, 2006. These data were provided by HCE (one of the EnRiMa partners).

4.2. Data for the strategic models

As we have already mentioned in Section 2, we have decided that most—if not all—parameters for the strategic models will not be modelled using time-series analysis of the historical data. Nevertheless, we will at some point need data also for these parameters, so we present data sources for all identified strategic parameters below.

Weather data

Weather data will be needed also for the strategic models, namely for the "embedded" operational scenarios. These scenarios should include both representative days, for estimating the average per-year running costs, and extreme days, to ensure that the installed infrastructure can guarantee user comfort even in extreme weather. Assuming stable weather patterns—a reasonable assumption, given the time scale of the model—, we can use the same set of weather scenarios in all the strategic nodes. This will simplify their generation considerably.

Energy prices

Just like the weather, the energy prices will appear in the embedded operational scenarios. However, unlike the weather, the energy prices in the later stages are different in each strategic node, so we cannot simply use the same scenarios throughout the tree.

What we will do instead is the following: we generate operational scenarios for the root node and then record their values *relative to* the average price of the last year or so. Then, we generate scenarios for the strategic level, modelling the development of the average annual energy prices. In other words, each of the strategic nodes of the tree will have a different value for the long-term average of the energy prices. Then we take, for each strategic node, the operational scenarios expressed in relative terms and combine them with the node's prices to produce the scenario values for the node.

How we generate values for the long-term scenarios depends on the time horizon of the model. Should the horizon be ten or even twenty years, we would use expert-created scenarios instead of historical data, as the data cannot be expected to have much predictive power so far ahead. If the time horizon turns out to be only a couple

[7] See http://en.exaa.at/market/historical/austria_germany/.

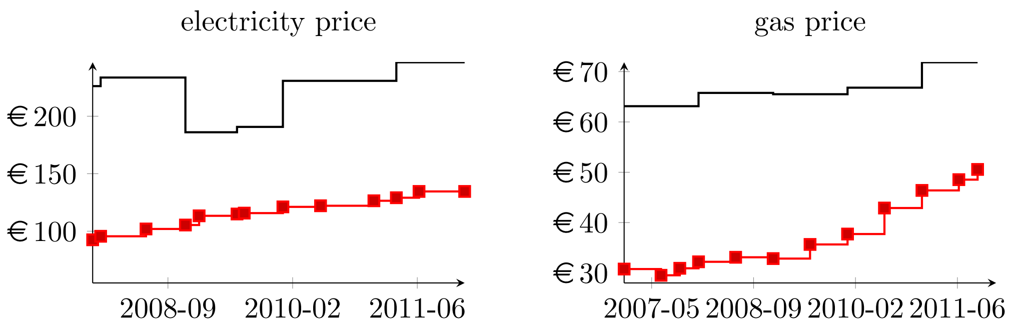

Figure 3: Electricity and gas prices at FASAD. Line "——" depicts the power prices in €/month and line "—■—' energy prices in €/MWh.

of years, which is the current proposal, then we can estimate the distribution of future prices by analyzing the historical distribution of changes over the given time horizon—of course providing we have long-enough time series.

For electricity and gas, we can get half-yearly prices for all European countries from the European Commission's Eurostat web page[8]. The length of the series varies between the countries; most of them start at the end of the 1980s, giving us enough data for our purpose. In addition, if the time horizon turns out to be five years or less, then we can use the *futures* prices as the means of the scenario values to guarantee the generated data agree with the market's expectations. The current futures prices can be obtained from the European Energy Exchange (EEX[9]), for both electricity and gas.

The problem with the above approach is that it generates scenarios for the *market price* of the energy sources. On the other hand, all our test sites have long-term energy contracts, so the price they pay can be significantly different from the market price. This means that the prices will have to be adjusted and for this we have to have data for the tariff prices faced by the sites.

For this purpose, we have obtained the electricity and gas prices for FASAD, both for power (fixed price per month) and for the consumed energy. The electricity data include twelve irregularly spaced prices, starting from December 2007, while the gas data include twelve irregularly spaced prices, starting from January 2007; see Figure 3. We have also the current version of a conversion table used to calculate the gas price from a six-month average of the Brent oil price and three-month average of the Euro/Dollar exchange rate.

We have also received monthly data for several commodity indices, including the IPE natural gas index (gas futures delivered at the National Balancing Point (NBP) in London), which is the base for gas prices faced by HCE. The file also includes monthly EUR/USD and EUR/GBP exchange rates.

Feed-in tariffs

In addition to the prices the sites pay for the energy they buy, we have to have data for the prices they get for the electricity they sell to the grid—the *feed-in* tariffs. The

[8]See http://epp.eurostat.ec.europa.eu/portal/page/portal/energy/data/database.
[9]See http://www.eex.com/en/Market Data/Trading Data.

structure of these tariffs varies between the countries: some countries do not use any special tariffs at all, while other countries have a wide range of feed-in tariffs, depending on the way the electricity was produced.

We have data feed-in tariffs in Spain, with quarterly resolution since July 2007. This data are published by the Ministry of Industry of Spain and are available from https://www.boe.es. In addition to the tariff prices, the data set includes also values for the consumer price index (CPI) and natural gas index, which are used for updating the tariff prices.

5. Results of the analyses

As we have explained at the end of Section 2, the only series on which we actually need to run a time-series analysis are the electricity prices. It is, of course, possible that some developments in the optimization models will require further data to be analyzed; in such a case, we will include these analyses in a later report.

We have used two different frameworks to analyze the electricity prices, one based on mean-reverting stochastic processes and the other on time-series models. For both of them, we tried two ways of modelling the hourly price series: (i) one entire hourly series with all available data and (ii) 24 separate series, one for each hour of the day. Since we have to do all this for two countries (Austria and Spain), we have done in total eight different analyses. On the other hand, we observe that the mean-reverting models perform worse than the time-series ones for all the studied cases, so we report only the latter in this section. In addition, we include one analysis based on the mean-reverting processes in the Appendix,in order to document the approach. After all, it might turn out to be better for some future data sets.

In this section, we therefore present, fit and test time-series models for the hourly electricity prices in Austria and Spain, using both the mentioned approaches. In both cases, we split the data set into an in-sample period for estimating the unknown parameters of the models and an out-of-sample one for evaluating the forecasting performance of the models. All the series exhibit seasonality, as their autocorrelation function shows a peak every twenty-four hours. The linear dependence in the series and the seasonality is captured by a seasonal autoregressive moving-average (seasonal ARMA, or SARMA) model, while the volatility is modelled using a generalized autoregressive conditional-heteroskedasticity (GARCH) model. For more background reading on time-series analysis, see, for example, Hamilton (1994); Tsay (2002); Bueno (2008); Aiube (2007).

5.1. Methodology

Basic concepts

A time-series analysis begins by fitting an ARMA(p, q) model to the series and examining the behavior of the error term. We start with the following ARMA(p, q) model:

$$\Phi_p\,(L)\,r_t = \Theta_q(L)\epsilon_t\,, \tag{1}$$

where L is the lag operator, Φ and Θ are polynomials of degrees p and q, respectively, r_t is the studied time series, and ϵ_t is a N(0,σ_ϵ^2) error term. In our case, we will be analysing log-return series for electricity prices p_t,

$$r_t = \log\left(\frac{p_t}{p_{t-1}}\right) . \tag{2}$$

The basis of time-series analysis is stationarity. In the finance literature, it is common to assume that returns are weakly stationary, and this characteristic can be checked through statistical tests, such as unit root tests.

A linear time-series model can be characterized by its autocorrelation function (ACF), and modelling makes use of the sample ACF to capture the linear dynamics of the data. From the literature on financial time series, in general, an AR(p) is enough to capture the linear dependence. The order (p, q) of an ARMA model in financial applications may depend on the frequency of the return series and can be obtained analyzing the ACF and the Partial ACF. In the case of electricity log-return series considered here, it is necessary to use a Seasonal ARMA (SARMA) model to capture the seasonality as well.

Also, it is well known that there is a strong non-linear dependence in the second moment, which can be seen by the fat tails in the histogram of the log-returns. To verify all these properties in practice, some analysis and statistical tests are conducted through QQ-plots, normality tests and dependence tests on the residuals and on the square of the residuals.

After modelling the possibly existing linear dependence, it is necessary to study the volatility. The variance will be modelled using an ARCH-GARCH model to capture these stylized facts. The conditional variance will be

$$\mathbb{E}\left[\left.{\upsilon_t}^2\right| I_{t-1}\right] = \sigma_t^2 ,$$

where I_{t-1} is the information available until time $t-1$. The residuals in Eq. (1) can be written as

$$\epsilon_t = \sigma_t \eta_t , \tag{3}$$

where $\eta_t \sim \mathrm{N}(0,1)$ and independent on each other. It follows that $\epsilon_t | I_{t-1} \sim \mathrm{N}(0, \sigma_t^2)$.

Using an in-sample period, the unknown parameters for the linear ARMA and GARCH models are estimated. The forecast is obtained and analyzed for the out-of-sample period. For a generic ARMA(p, q) and GARCH(m, s) model, we have

$$r_t = \sum_{i=1}^{p} \phi_i r_{t-i} + \epsilon_t - \sum_{j=1}^{q} \theta_j \epsilon_{t-j} \tag{4}$$

$$\sigma_t^2 = \alpha_o + \sum_{i=1}^{m} \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^{s} \beta_j \sigma_{t-j}^2 \tag{5}$$

Under this framework, the best model will be chosen using the Akaike information criterion (AIC) and Bayesian information criterion (BIC). Also, the comparison among different models relating to estimation in the in-sample period and forecasting for the out-of-sample period can guide us to the best fitting model. In addition, the model will be prepared to capture seasonality and intra-day effects.

To verify the forecast performance, we use the following aggregate error measures:

1. Mean Absolute Error (MAE), given by:

$$MAE = \frac{\sum_{t=1}^{N} |E_t|}{N}\,, \tag{6}$$

where E_t is the difference between the actual and the forecasted value.

2. Mean Absolute Percentage Error (MAPE), given by:

$$MAPE = \frac{\sum_{t=1}^{N} \left|\frac{E_t}{p_t}\right|}{N}\,, \tag{7}$$

where p_t is the actual value.

3. Root Mean Squared Error (RMSE), given by:

$$RMSE = \sqrt{\frac{\sum_{t=1}^{N} E_t^2}{N}}\,. \tag{8}$$

Two different approaches

We test two different approaches for modelling the hourly returns. First, we work with the entire hourly series. The hourly prices observations are taken in sequence, as a whole time series, in the way described above.

In the second approach, we work with 24 time series, one for each hour of the day. Following the methodology of modelling the linear dependence by a SARMA model and the conditional variance by a GARCH model, we will obtain 24 different time series models for each hour series, with different orders and different estimated values for the parameters.

For each hour h, we have

$$r_{ht} = \sum_{i=1}^{p} \phi_{hi} r_{ht-i} + \epsilon_{ht} - \sum_{j=1}^{q} \theta_{hj} \epsilon_{ht-j} \tag{9}$$

$$\sigma_{ht}^2 = \alpha_{ho} + \sum_{i=1}^{m} \alpha_{hi} \epsilon_{ht-i}^2 + \sum_{j=1}^{s} \beta_{hj} \sigma_{ht-j}^2\,. \tag{10}$$

Each of these series will then be analyzed in the same way as the complete series in the first approach.

5.2. Austrian data

Descriptive analysis

A total of 43 080 hourly observations over five years of electricity spot prices in €/MWh from Austria/Germany markets, provided by Energy Exchange Austria (EXAA), are

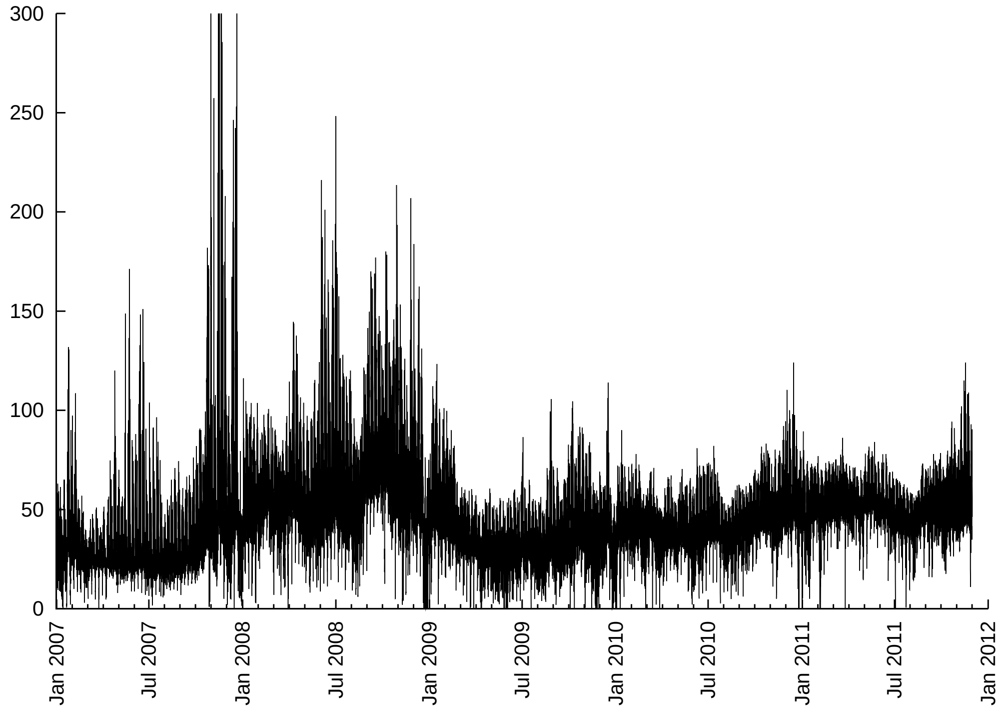

Figure 4: Electricity spot price in Austria, in €/MWh – whole series. Note that the figure has been capped, the highest value is €519.93 from 2007-11-15 18:00.

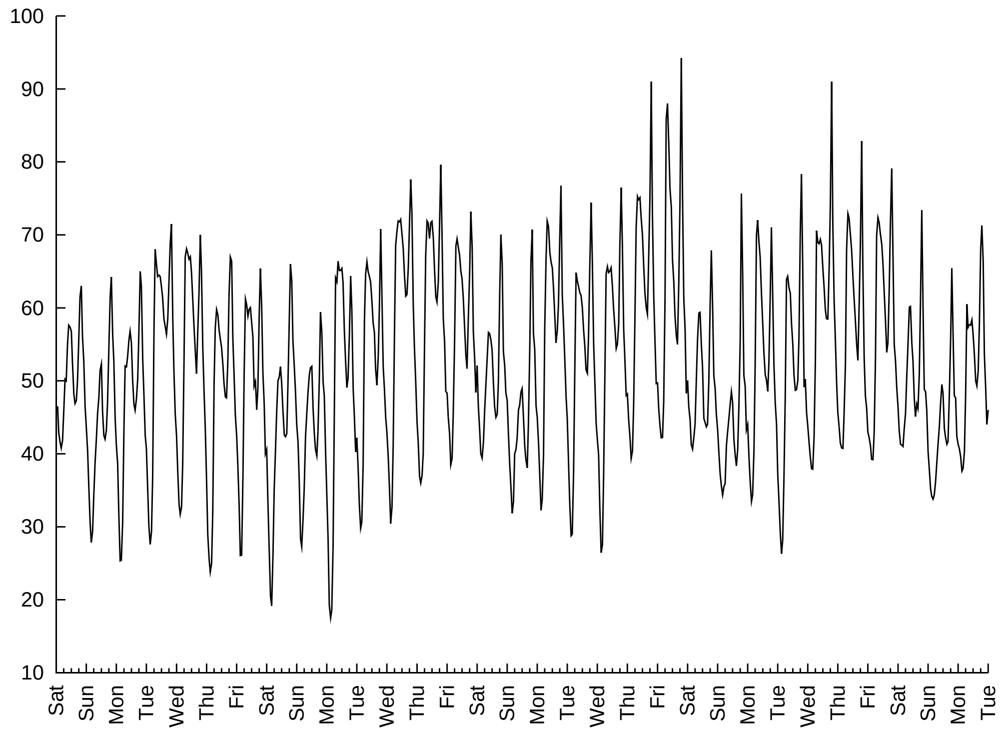

Figure 5: Electricity spot price in Austria, in €/MWh – October 2011

Table 1: Descriptive statistics for the entire hourly series, Austrian data

Statistic	Price ($\frac{€}{\text{MWh}}$)	Log-return
Mean	48.20	2.06e-05
Std. dev.	23.92	0.36
Variance	572.32	0.13
Skewness	2.40	0.48
Kurtosis	25.67	316.42
# of obs.	43080	43079

Table 2: Descriptive statistics for per-hour series, Austrian data (1795 observations)

	Electricity price in €/MWh, per hour											
Statistic	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**
Mean	37.23	32.72	29.32	26.90	27.24	31.68	38.51	50.55	55.44	58.81	61.17	64.79
Std. dev.	12.37	12.34	12.41	12.31	12.47	13.40	17.50	24.16	23.75	23.74	24.29	27.30
Variance	153.0	152.3	154.1	151.6	155.6	179.6	306.4	583.8	564.2	563.5	590.2	745.1
Skewness	0.17	0.02	-0.03	0.08	0.07	-0.08	-0.01	0.57	0.88	1.31	1.60	1.97
Kurtosis	2.86	2.67	2.56	2.43	2.48	2.72	2.78	3.71	4.74	5.95	6.80	8.76
Statistic	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**	**24**
Mean	60.41	56.93	53.72	51.17	51.39	57.86	61.72	59.02	54.09	48.56	47.18	40.43
Std. Dev.	22.45	21.86	21.38	20.47	21.63	34.35	33.82	23.64	18.25	14.72	13.52	12.10
Variance	503.9	478.0	456.9	418.9	468.0	1179.7	1143.7	558.9	333.1	216.5	182.9	146.4
Skewness	1.55	1.39	1.32	1.24	1.54	5.51	5.03	2.05	1.02	0.75	0.45	0.29
Kurtosis	6.60	5.92	5.66	5.32	7.18	55.83	49.25	13.57	4.75	3.87	3.20	2.99

available. The sample period begins on January 1$^{\text{st}}$, 2007 and ends on November 30$^{\text{th}}$, 2011. There are five missing values, all of them for hour 3. These were replaced by 0.001, in order to obtain the return series. The resulting series is presented in Figure 4. In addition, Figure 5 presents only prices for one month, so we can see the daily and weekly cycles. The data set is then split into two periods:

- an in-sample period, from January 1$^{\text{st}}$, 2007 to December 31$^{\text{st}}$, 2010 (35 064 observations), which is used to estimate the unknown parameters, and
- an out-of-sample period, from January 1$^{\text{st}}$, 2011 to November 30$^{\text{th}}$, 2011 (8016 observations), which is used to assess the forecast of the model proposed.

The descriptive statistics of the entire hourly series of prices, as well as the log-returns, are presented in Table 1. Statistics for each of the 24 hours of the day separately are presented in Tables 2 for the prices and Table 3 for the log-returns.

First approach: one entire hourly series in sequence

In this approach, we work with one entire hourly series, with the observations in sequence. The autocorrelogram for the series is presented in Figure 6. There, column “AC” includes

Table 3: Descriptive statistics for log-returns of 24 per-hour series of Austrian electricity prices (1794 observations)

	Log-returns of electricity price, per hour											
Statistic	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**
Mean	4e-04	7e-04	1e-03	2e-03	5e-03	5e-03	5e-03	5e-03	1e-03	1e-03	1e-03	4e-03
Std. Dev.	0.51	0.75	1.18	0.98	1.07	1.07	1.10	1.06	0.66	0.35	0.28	0.31
Variance	0.3	0.6	1.4	1.0	1.2	1.2	1.2	1.1	0.4	0.1	0.1	0.1
Skewness	-0.08	-0.29	-0.03	0.14	0.11	0.19	0.63	0.74	-0.46	0.90	1.05	6.19
Kurtosis	168.4	77.68	49.41	40.03	34.36	35.50	28.79	38.18	89.10	16.99	5.85	124.7

Statistic	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**	**24**
Mean	5e-04	6e-04	7e-04	7e-04	7e-04	7e-04	7e-04	5e-04	6e-04	6e-04	4e-04	5.e-04
Std. Dev.	0.24	0.27	0.31	0.31	0.30	0.26	0.22	0.18	0.15	0.14	0.21	0.18
Variance	0.1	0.1	0.1	0.1	0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0
Skewness	0.96	1.04	0.90	1.08	1.11	1.13	1.18	0.57	0.41	0.46	-0.06	-1.39
Kurtosis	6.22	5.67	15.97	5.25	5.93	9.79	11.97	8.33	5.62	48.01	254.6	190.2

autocorrelations and "PAC" the partial autocorrelations, for each lag. The "Q-Stat" and "Prob" columns then show values of the Ljung-Box statistic and their p-values, respectively. The figure was created using the EViews statistical package.

Following the methodology of modelling the linear dependence by a Seasonal ARMA model and the conditional variance by a GARCH model, the best fitted model for the series was a $\text{SARMA}(2,2) \times (2,2)_{24}$ + $\text{GARCH}(1,1)$. The estimation results are presented in Table 4. The estimated SARMA and $\text{GARCH}(1,1)$ equations are

$$\begin{aligned}\left(1-0.68L-0.15L^2\right)\left(1-1.50L^{24}+0.55L^{48}\right) r_t \\ = \left(1+0.60L+0.40L^2\right)\left(1+1.24L^{24}-0.38L^{48}\right)\epsilon_t \,. \end{aligned} \tag{11}$$

$$\sigma_t^2 = 0.0019 + 5.2663\epsilon_{t-1}^2 + 0.1360\sigma_{t-1}^2 \,. \tag{12}$$

The autocorrelogram of the residuals is presented in Figure 7. We can see that the seasonality was mostly captured by the SARMA model. Although statistically significant, the autocorrelation absolute values are not relevant except for lags 23, 24 and 25, related to seasonality. However, we decided to work with a parsimonious model, which presents order 2 for both seasonal autoregressive and moving average parts, and provides an acceptable forecast error (forecast performance is shown at the end of this section).

In Figure 8, we present the autocorrelogram for the squared residuals. The results show that the non-linear dependence was well captured by GARCH. Even though lag 24 presents a relevant value, we decided to work with a GARCH(1,1) model, which provides an acceptable forecast error as already mentioned.

From the fitted model, the forecasts are obtained for the out-of-sample period, from January 1st, 2011, hour 1 to November 30th, 2011, hour 24. The aggregate error measures (RMSE, MAPE and MAE) are presented in columns 2–4 of Table 5. We present now the overall results for the forecasted series. Besides, in order to compare the accuracy of this model—which works with an entire hourly series—, with the model in the next section,

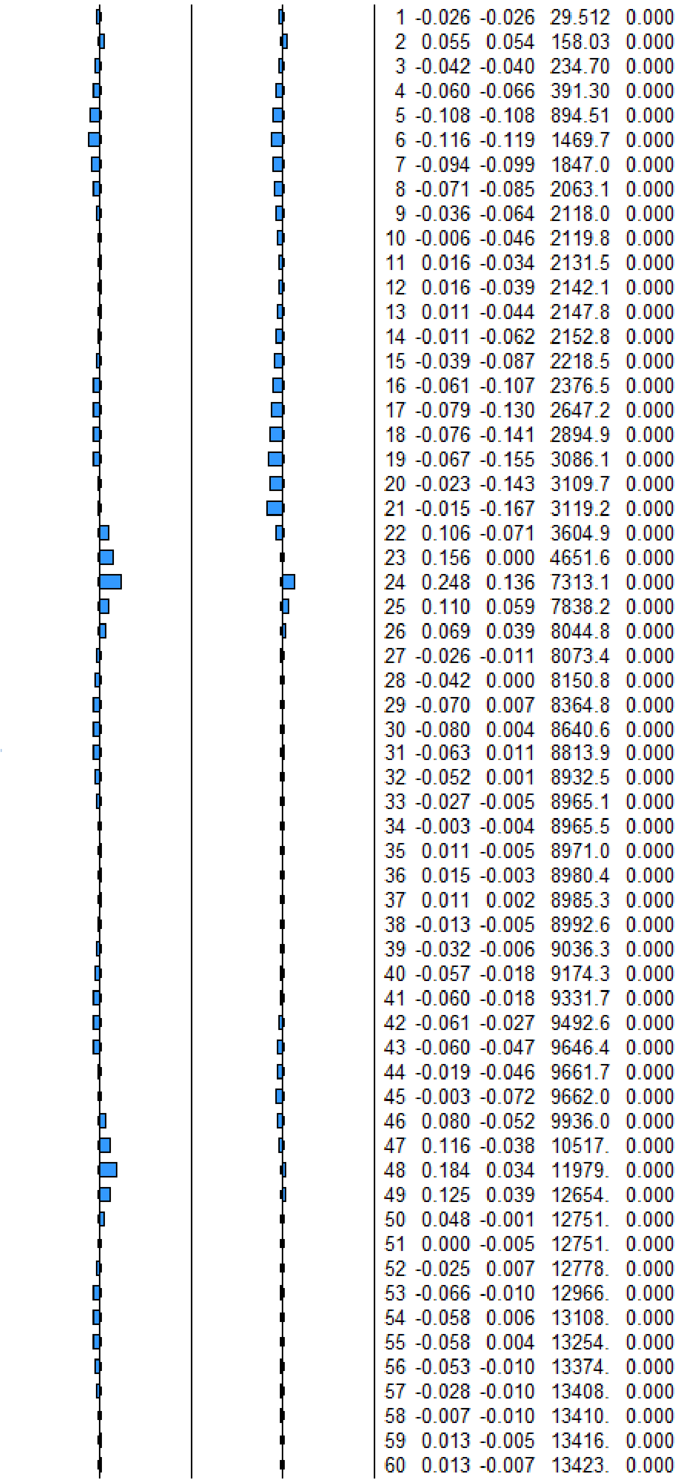

Date: 01/24/12 Time: 16:31
Sample: 1 43080
Included observations: 43079

	AC	PAC	Q-Stat	Prob
1	-0.026	-0.026	29.512	0.000
2	0.055	0.054	158.03	0.000
3	-0.042	-0.040	234.70	0.000
4	-0.060	-0.066	391.30	0.000
5	-0.108	-0.108	894.51	0.000
6	-0.116	-0.119	1469.7	0.000
7	-0.094	-0.099	1847.0	0.000
8	-0.071	-0.085	2063.1	0.000
9	-0.036	-0.064	2118.0	0.000
10	-0.006	-0.046	2119.8	0.000
11	0.016	-0.034	2131.5	0.000
12	0.016	-0.039	2142.1	0.000
13	0.011	-0.044	2147.8	0.000
14	-0.011	-0.062	2152.8	0.000
15	-0.039	-0.087	2218.5	0.000
16	-0.061	-0.107	2376.5	0.000
17	-0.079	-0.130	2647.2	0.000
18	-0.076	-0.141	2894.9	0.000
19	-0.067	-0.155	3086.1	0.000
20	-0.023	-0.143	3109.7	0.000
21	-0.015	-0.167	3119.2	0.000
22	0.106	-0.071	3604.9	0.000
23	0.156	0.000	4651.6	0.000
24	0.248	0.136	7313.1	0.000
25	0.110	0.059	7838.2	0.000
26	0.069	0.039	8044.8	0.000
27	-0.026	-0.011	8073.4	0.000
28	-0.042	0.000	8150.8	0.000
29	-0.070	0.007	8364.8	0.000
30	-0.080	0.004	8640.6	0.000
31	-0.063	0.011	8813.9	0.000
32	-0.052	0.001	8932.5	0.000
33	-0.027	-0.005	8965.1	0.000
34	-0.003	-0.004	8965.5	0.000
35	0.011	-0.005	8971.0	0.000
36	0.015	-0.003	8980.4	0.000
37	0.011	0.002	8985.3	0.000
38	-0.013	-0.005	8992.6	0.000
39	-0.032	-0.006	9036.3	0.000
40	-0.057	-0.018	9174.3	0.000
41	-0.060	-0.018	9331.7	0.000
42	-0.061	-0.027	9492.6	0.000
43	-0.060	-0.047	9646.4	0.000
44	-0.019	-0.046	9661.7	0.000
45	-0.003	-0.072	9662.0	0.000
46	0.080	-0.052	9936.0	0.000
47	0.116	-0.038	10517.	0.000
48	0.184	0.034	11979.	0.000
49	0.125	0.039	12654.	0.000
50	0.048	-0.001	12751.	0.000
51	0.000	-0.005	12751.	0.000
52	-0.025	0.007	12778.	0.000
53	-0.066	-0.010	12966.	0.000
54	-0.058	0.006	13108.	0.000
55	-0.058	0.004	13254.	0.000
56	-0.053	-0.010	13374.	0.000
57	-0.028	-0.010	13408.	0.000
58	-0.007	-0.010	13410.	0.000
59	0.013	-0.005	13416.	0.000
60	0.013	-0.007	13423.	0.000

Figure 6: Autocorrelations of the entire hourly series, Austria

Date: 01/24/12 Time: 16:54
Sample: 52 43080
Included observations: 43029
Q-statistic probabilities adjusted for 8 ARMA term(s)

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.042	0.042	76.182	
		2	0.067	0.066	271.91	
		3	0.011	0.006	277.28	
		4	0.019	0.014	292.15	
		5	-0.007	-0.009	294.14	
		6	0.001	0.000	294.20	
		7	0.008	0.008	296.69	
		8	0.008	0.007	299.14	
		9	0.020	0.019	316.02	0.000
		10	0.029	0.027	352.89	0.000
		11	0.045	0.041	441.67	0.000
		12	0.024	0.017	466.20	0.000
		13	0.026	0.019	496.22	0.000
		14	0.011	0.006	501.71	0.000
		15	-0.007	-0.012	503.70	0.000
		16	-0.009	-0.010	507.43	0.000
		17	-0.002	-0.002	507.68	0.000
		18	0.006	0.006	509.12	0.000
		19	0.005	0.004	510.29	0.000
		20	0.023	0.020	533.36	0.000
		21	0.025	0.019	559.52	0.000
		22	0.040	0.032	628.34	0.000
		23	0.082	0.074	918.46	0.000
		24	0.050	0.038	1026.1	0.000
		25	0.076	0.063	1273.2	0.000
		26	0.030	0.019	1311.1	0.000
		27	0.006	-0.005	1312.8	0.000
		28	-0.003	-0.007	1313.3	0.000
		29	-0.017	-0.019	1325.0	0.000
		30	-0.015	-0.017	1334.6	0.000
		31	-0.004	-0.006	1335.2	0.000
		32	0.002	-0.004	1335.3	0.000
		33	0.008	-0.001	1338.1	0.000
		34	0.002	-0.010	1338.3	0.000
		35	0.005	-0.006	1339.4	0.000
		36	0.005	-0.005	1340.3	0.000
		37	-0.002	-0.007	1340.5	0.000
		38	-0.012	-0.013	1347.0	0.000
		39	-0.014	-0.011	1356.0	0.000
		40	-0.025	-0.019	1382.8	0.000
		41	-0.016	-0.011	1393.8	0.000
		42	-0.011	-0.008	1399.1	0.000
		43	-0.004	-0.006	1399.9	0.000
		44	-0.008	-0.012	1402.5	0.000
		45	-0.001	-0.009	1402.5	0.000
		46	0.005	-0.005	1403.6	0.000
		47	0.015	0.005	1412.7	0.000
		48	0.007	-0.004	1415.0	0.000
		49	0.033	0.027	1461.2	0.000
		50	0.015	0.011	1470.4	0.000
		51	0.004	0.003	1471.2	0.000
		52	0.001	0.006	1471.2	0.000
		53	-0.007	-0.001	1473.2	0.000
		54	-0.014	-0.008	1481.7	0.000
		55	-0.021	-0.016	1501.1	0.000
		56	-0.016	-0.013	1512.2	0.000
		57	-0.010	-0.007	1516.8	0.000
		58	-0.007	-0.006	1519.1	0.000
		59	-0.002	-0.001	1519.3	0.000
		60	0.007	0.007	1521.2	0.000

Figure 7: Autocorrelations of residuals of the entire hourly series, Austria

Date: 01/24/12 Time: 17:01
Sample: 52 43080
Included observations: 43029
Q-statistic probabilities adjusted for 8 ARMA term(s)

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.000	0.000	0.0107	
		2	-0.001	-0.001	0.0361	
		3	0.000	0.000	0.0458	
		4	0.000	0.000	0.0458	
		5	-0.001	-0.001	0.0612	
		6	0.001	0.001	0.1230	
		7	-0.001	-0.001	0.1394	
		8	0.000	0.000	0.1394	
		9	-0.001	-0.001	0.1599	0.689
		10	-0.001	-0.001	0.1760	0.916
		11	0.000	0.000	0.1772	0.981
		12	-0.001	-0.001	0.1933	0.996
		13	0.000	0.000	0.2020	0.999
		14	-0.001	-0.001	0.2212	1.000
		15	-0.001	-0.001	0.2398	1.000
		16	0.001	0.001	0.2970	1.000
		17	0.000	0.000	0.3016	1.000
		18	0.000	0.000	0.3032	1.000
		19	-0.001	-0.001	0.3223	1.000
		20	0.000	0.000	0.3250	1.000
		21	0.000	0.000	0.3295	1.000
		22	0.000	0.000	0.3366	1.000
		23	0.001	0.001	0.3793	1.000
		24	0.103	0.103	459.27	0.000
		25	0.000	0.000	459.27	0.000
		26	-0.001	0.000	459.28	0.000
		27	0.000	0.000	459.29	0.000
		28	0.000	0.000	459.29	0.000
		29	-0.001	0.000	459.30	0.000
		30	0.000	-0.001	459.30	0.000
		31	0.000	0.000	459.31	0.000
		32	0.000	0.000	459.32	0.000
		33	-0.001	-0.001	459.34	0.000
		34	-0.001	0.000	459.35	0.000
		35	-0.001	0.000	459.36	0.000
		36	-0.001	-0.001	459.38	0.000
		37	-0.001	-0.001	459.40	0.000
		38	-0.001	0.000	459.42	0.000
		39	0.000	0.000	459.42	0.000
		40	0.000	0.000	459.42	0.000
		41	0.000	0.000	459.42	0.000
		42	0.000	0.000	459.43	0.000
		43	-0.001	0.000	459.44	0.000
		44	0.000	0.000	459.44	0.000
		45	0.000	0.000	459.44	0.000
		46	0.000	0.000	459.45	0.000
		47	0.000	-0.001	459.46	0.000
		48	0.032	0.022	504.07	0.000
		49	0.000	0.000	504.08	0.000
		50	0.000	0.000	504.08	0.000
		51	0.000	0.000	504.09	0.000
		52	-0.001	-0.001	504.10	0.000
		53	0.000	0.000	504.10	0.000
		54	0.000	0.000	504.11	0.000
		55	0.000	0.000	504.11	0.000
		56	-0.001	0.000	504.13	0.000
		57	-0.001	0.000	504.14	0.000
		58	-0.001	-0.001	504.16	0.000
		59	-0.001	-0.001	504.18	0.000
		60	-0.001	-0.001	504.20	0.000

Figure 8: Autocorrelations of squared residuals of the entire hourly series, Austria

Table 4: Estimated parameters for SARMA-GARCH in entire hourly series, Austria

Model	Coef.	Std. err.	z-stat.	Prob.
AR (1)	0.677	0.002362	286.527	0.000
AR (2)	0.153	0.002154	70.919	0.000
SAR(24)	1.498	0.000823	1820.920	0.000
SAR(48)	−0.547	0.000681	−803.382	0.000
MA (1)	−0.599	0.001899	−315.661	0.000
MA (2)	−0.400	0.001898	−210.751	0.000
SMA(24)	−1.244	0.000808	−1539.009	0.000
SMA(48)	0.378	0.000465	811.090	0.000
C	0.002	0.000011	165.035	0.000
ARCH(1)	5.266	0.006853	768.490	0.000
GARCH(1)	0.136	0.000312	435.133	0.000

we also present measures for each hour separately. To obtain this result, we considered the entire hourly forecasted series and separated the forecasts in 24 series, one for each hour of the day. For example, the aggregate error for hour 1 series is calculated from the forecasted series composed by $\mathbb{E}[p_1]$, $\mathbb{E}[p_{25}], \mathbb{E}[p_{49}]$, etc, where the expectation is conditional on the in-sample values.

For hours 3, 5 and 8, the MAPE values are very high. This is due to very low data values that influence the overall measure. Including in-sample and out-of-sample periods, besides the missing values, there is a total of 184 observations with values equal or less than €1/MWh. In the out-of-sample period, there is one missing value in the original data and 14 very low values (equal or less than €1/MWh). Excluding those values for the additional analysis of the forecasting performance, the aggregate error measures are presented in the last three columns of Table 5. These results are much better than the previous ones.

Second approach: 24 separate series for each hour of the day

In this approach, we work with 24 time series, one for each hour of the day, each treated the same way as the whole series in the previous section. For each series, we adjusted several models and analyzed their residuals in order to find the best one, observing the AIC and BIC criteria. Table 6 shows the best fitted model for each one of the 24 series.

From the fitted models, we forecast the prices in the out-of-sample period, from January 1st, 2011 to 30th. Table 7 presents the aggregate error measures (MAE, MAPE and RMSE) for the price series forecast. For the hours between 3 and 9, the MAPE measures are very high. Just as in the previous case, this is due to very low values in the data series, which influence the overall measures. Excluding those values for the additional analysis of the forecasting performance, the aggregate error measures improve significantly, as shown in Table 8. In the same table, we also include the corresponding results from the previous section, so we can compare the two approaches.

From MAPE and MAE measures, the model using one entire series gives better results than the 24 separate series model. Regarding the RMSE measure, for some hours in the

Table 5: Aggregate measures for the SARMA-GARCH model for Austria

	Complete series			Excluding small values		
hour	**MAE**	**MAPE**	**RMSE**	**MAE**	**MAPE**	**RMSE**
1	2.74	7%	3.93	2.74	7%	3.93
2	1.82	6%	2.50	1.82	6%	2.50
3	2.39	13187%	5.05	2.26	8%	4.45
4	4.90	14%	42.97	4.91	14%	43.03
5	4.00	209%	40.64	4.00	11%	40.88
6	2.62	9%	3.92	2.65	7%	3.95
7	6.04	17%	8.83	6.05	16%	8.85
8	4.97	449%	7.51	4.96	10%	7.49
9	4.49	9%	28.63	4.51	8%	28.68
10	2.70	5%	9.15	2.70	5%	9.14
11	1.51	3%	3.09	1.50	3%	3.09
12	1.42	2%	1.98	1.42	2%	1.98
13	1.17	2%	1.66	1.18	2%	1.66
14	2.46	5%	3.06	2.46	5%	3.06
15	1.61	3%	2.09	1.61	3%	2.08
16	1.28	3%	1.66	1.28	3%	1.66
17	1.46	3%	2.05	1.47	3%	2.06
18	2.49	4%	3.72	2.51	4%	3.74
19	2.41	4%	3.45	2.41	4%	3.45
20	2.33	3%	3.74	2.34	3%	3.74
21	2.17	4%	2.93	2.17	4%	2.92
22	1.71	3%	2.30	1.71	3%	2.30
23	1.44	3%	1.89	1.44	3%	1.89
24	2.04	4%	2.55	2.04	4%	2.55
Overall	**2.59**	**582%**	**13.99**	**2.59**	**6%**	**14.03**

morning, the result using 24 separate series is better. Considering the comparison as a whole, and since in the entire series approach we work with only one model instead of 24 different models, we consider the first one the best model for the considered data. It is worth mentioning that both models perform worst in the early morning hours.

5.3. Spanish data

Descriptive Analysis

A total of 42 720 hourly observations (close to five years) of Spanish electricity spot prices in €/MWh are available. The sample period begins on January 1$^{\text{st}}$, 2007 and ends on November 15$^{\text{th}}$, 2011. There are 5 missing values and 379 zeros; these values were, again, replaced by 0.001 in order to be able to obtain the log-return series. The resulting series can be seen in Figure 9, while Figure 10 again presents prices from a one month sub-interval. The data set is split into two periods: an in-sample period

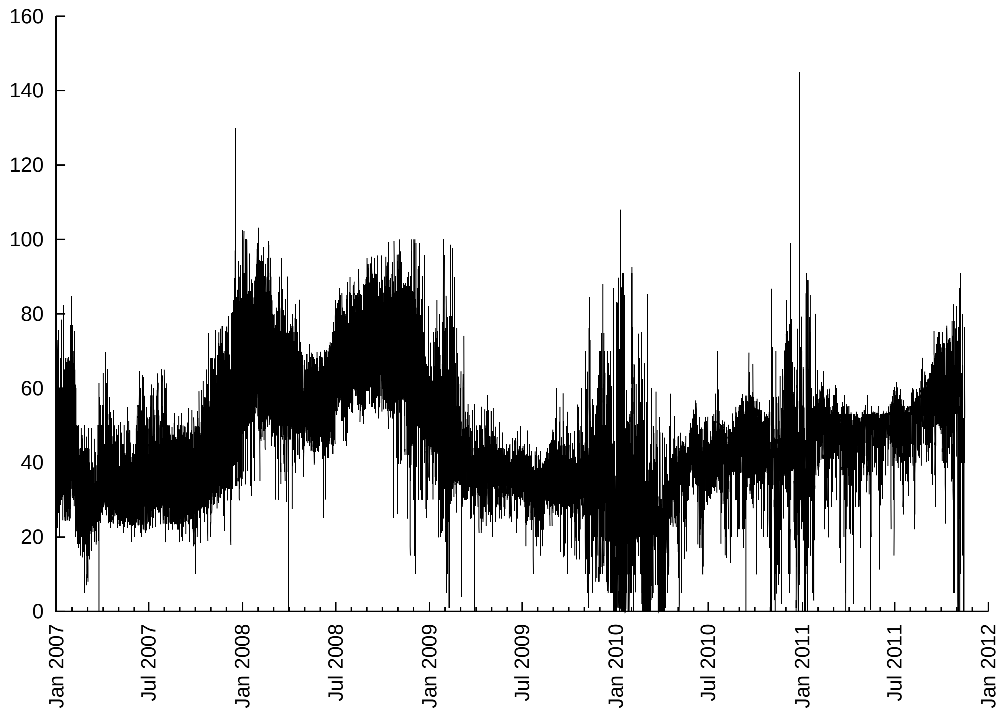

Figure 9: Electricity spot price in Spain, in €/MWh – whole series

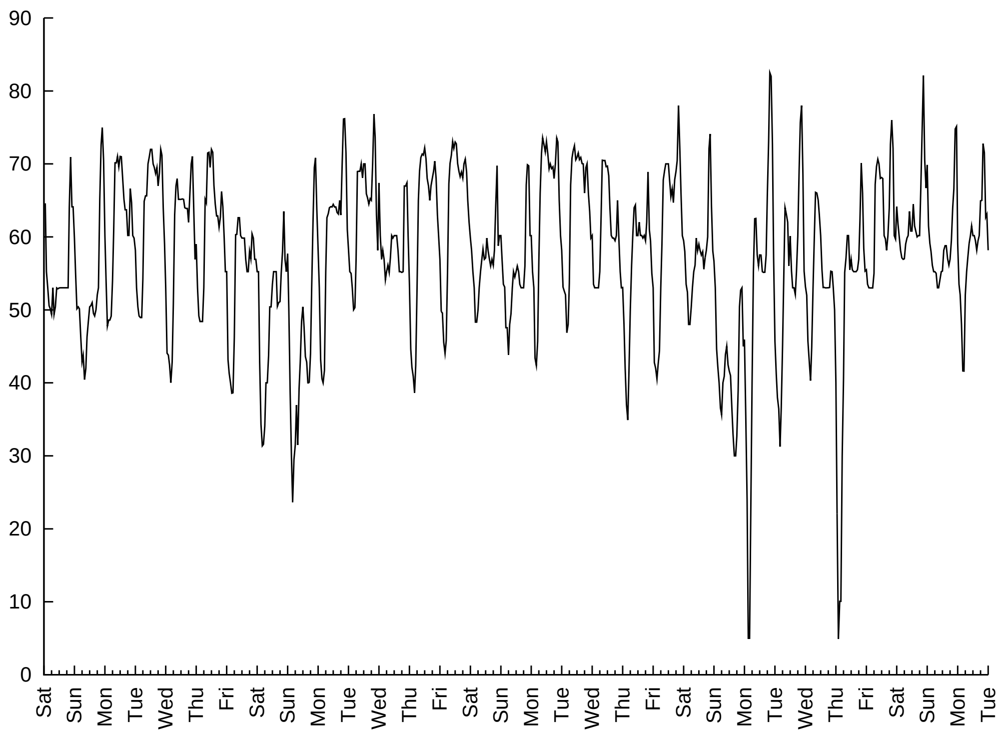

Figure 10: Electricity spot price in Spain, in €/MWh – October 2011

Table 6: Best fitted models for 24 separate series, Austria

hour	**Best fitted model**
1	$\text{SARMA}(0,1) \times (1,1)_7 + \text{GARCH}(2,0)$
2	$\text{SARMA}(1,1) \times (1,1)_7 + \text{GARCH}(1,1)$
3	$\text{SARMA}(0,2) \times (1,1)_7 + \text{GARCH}(2,1)$
4	$\text{SARMA}(0,1) \times (1,1)_7 + \text{GARCH}(2,1)$
5	$\text{SARMA}(1,1) \times (1,1)_7 + \text{GARCH}(1,1)$
6	$\text{SARMA}(3,1) \times (1,1)_7 + \text{GARCH}(2,1)$
7	$\text{SARMA}(1,1) \times (2,1)_7 + \text{GARCH}(2,1)$
8	$\text{SARMA}(1,1) \times (1,1)_7 + \text{GARCH}(2,0)$
9	$\text{SARMA}(1,1) \times (1,2)_7 + \text{GARCH}(3,1)$
10	$\text{SARMA}(1,1) \times (1,2)_7 + \text{GARCH}(1,1)$
11	$\text{SARMA}(2,1) \times (1,2)_7 + \text{GARCH}(2,1)$
12	$\text{SARMA}(1,2) \times (2,1)_7 + \text{GARCH}(2,2)$
13	$\text{SARMA}(1,2) \times (1,2)_7 + \text{GARCH}(2,2)$
14	$\text{SARMA}(1,2) \times (1,2)_7 + \text{GARCH}(2,1)$
15	$\text{SARMA}(1,2) \times (1,2)_7 + \text{GARCH}(1,1)$
16	$\text{SARMA}(1,1) \times (2,2)_7 + \text{GARCH}(2,1)$
17	$\text{SARMA}(2,1) \times (2,1)_7 + \text{GARCH}(2,2)$
18	$\text{SARMA}(1,1) \times (1,2)_7 + \text{GARCH}(2,2)$
19	$\text{SARMA}(2,2) \times (2,2)_7 + \text{GARCH}(1,1)$
20	$\text{SARMA}(1,1) \times (2,1)_7 + \text{GARCH}(2,1)$
21	$\text{SARMA}(1,1) \times (2,1)_7 + \text{GARCH}(2,1)$
22	$\text{SARMA}(1,2) \times (1,2)_7 + \text{GARCH}(2,1)$
23	$\text{SARMA}(1,1) \times (1,1)_7 + \text{GARCH}(2,1)$
24	$\text{SARMA}(1,1) \times (1,1)_7 + \text{GARCH}(1,1)$

(January 1$^{\text{st}}$, 2007 to December 31$^{\text{st}}$, 2010) with 35 064 observations, used to estimate the unknown parameters; and an out-of-sample period (January 1$^{\text{st}}$, 2011 to November 15$^{\text{th}}$, 2011) with 7656 observations, used to assess the model's forecast.

A summary of the descriptive statistics of the entire hourly series, for both the prices and log-returns, is presented in Table 9. Considering the 24 series for each hour of the day, their statistics are presented in Table 10 for prices and Table 11 for log-returns.

First approach: one entire hourly series in sequence

Just like for the Austrian data, we start by analyzing the hourly prices as one data series. The autocorrelogram for the series is presented in Figure 11.

Following the methodology of modelling the linear dependence by a Seasonal ARMA model and the conditional variance by a GARCH model, the best fitted model for the series was a $\text{SARMA}(2,2) \times (2,2)_{24} + \text{GARCH}(2,1)$. The estimation results are presented

Date: 01/24/12 Time: 17:24
Sample: 1 42720
Included observations: 42719

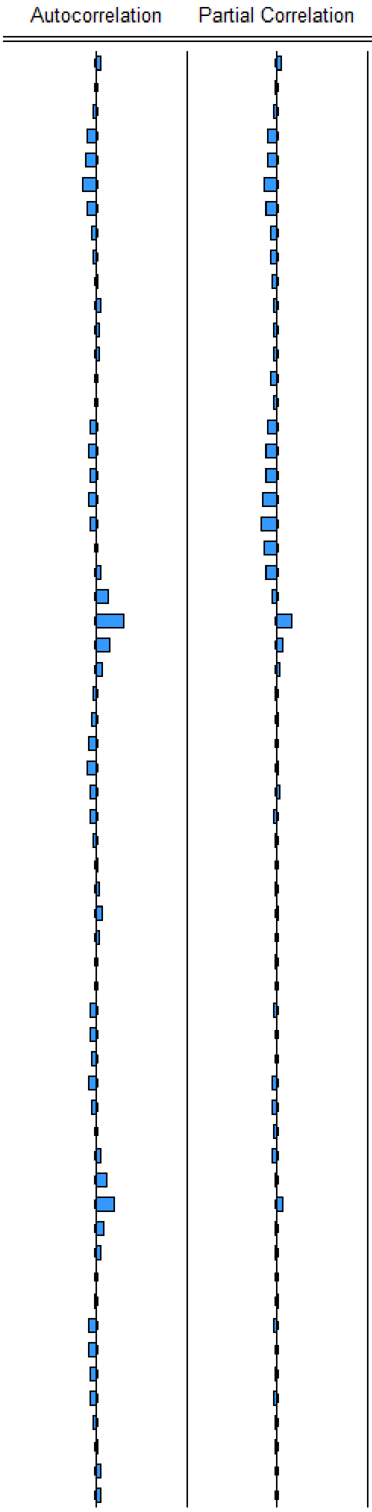

	AC	PAC	Q-Stat	Prob
1	0.056	0.056	135.45	0.000
2	-0.005	-0.008	136.60	0.000
3	-0.027	-0.026	167.77	0.000
4	-0.114	-0.111	719.20	0.000
5	-0.121	-0.110	1341.1	0.000
6	-0.151	-0.146	2317.3	0.000
7	-0.112	-0.115	2855.8	0.000
8	-0.046	-0.069	2945.1	0.000
9	-0.034	-0.078	2993.7	0.000
10	0.019	-0.042	3009.2	0.000
11	0.047	-0.025	3102.9	0.000
12	0.036	-0.035	3159.5	0.000
13	0.029	-0.036	3195.4	0.000
14	-0.007	-0.058	3197.7	0.000
15	-0.006	-0.040	3199.0	0.000
16	-0.069	-0.099	3403.0	0.000
17	-0.093	-0.119	3775.3	0.000
18	-0.079	-0.118	4038.7	0.000
19	-0.087	-0.149	4360.7	0.000
20	-0.068	-0.167	4559.8	0.000
21	-0.007	-0.147	4562.0	0.000
22	0.054	-0.120	4685.6	0.000
23	0.135	-0.053	5461.7	0.000
24	0.315	0.179	9697.6	0.000
25	0.151	0.067	10670.	0.000
26	0.066	0.038	10858.	0.000
27	-0.034	-0.025	10906.	0.000
28	-0.046	0.024	10998.	0.000
29	-0.092	0.003	11361.	0.000
30	-0.104	0.012	11826.	0.000
31	-0.069	0.026	12030.	0.000
32	-0.078	-0.030	12291.	0.000
33	-0.029	-0.019	12327.	0.000
34	0.020	-0.006	12344.	0.000
35	0.032	-0.011	12389.	0.000
36	0.061	0.017	12546.	0.000
37	0.028	0.001	12581.	0.000
38	-0.004	-0.010	12581.	0.000
39	-0.006	-0.002	12583.	0.000
40	-0.072	-0.032	12808.	0.000
41	-0.060	-0.007	12964.	0.000
42	-0.057	-0.004	13100.	0.000
43	-0.088	-0.048	13429.	0.000
44	-0.050	-0.044	13538.	0.000
45	0.002	-0.032	13538.	0.000
46	0.043	-0.047	13617.	0.000
47	0.120	-0.013	14231.	0.000
48	0.221	0.069	16323.	0.000
49	0.098	-0.021	16734.	0.000
50	0.046	-0.022	16824.	0.000
51	0.007	0.003	16825.	0.000
52	-0.016	0.024	16836.	0.000
53	-0.093	-0.027	17203.	0.000
54	-0.083	-0.005	17496.	0.000
55	-0.071	-0.013	17714.	0.000
56	-0.069	-0.029	17915.	0.000
57	-0.026	-0.018	17942.	0.000
58	0.013	-0.011	17949.	0.000
59	0.052	0.004	18067.	0.000
60	0.054	-0.005	18190.	0.000

Figure 11: Autocorrelations for Spanish log-return series

Table 7: Aggregate error measures for the per-hour series, Austrian data

hour	MAE	MAPE	RMSE	hour	MAE	MAPE	RMSE
1	3.44	9%	4.93	13	3.54	6%	4.71
2	4.65	16%	6.58	14	3.76	7%	5.15
3	7.42	8959%	9.95	15	3.62	7%	4.95
4	5.33	690%	7.16	16	3.61	7%	5.01
5	7.69	1411%	10.28	17	3.66	7%	5.01
6	6.20	1589%	9.08	18	3.65	6%	4.99
7	5.30	901%	7.88	19	3.79	6%	5.45
8	7.94	968%	11.47	20	3.47	6%	4.54
9	5.60	658%	8.50	21	3.21	5%	4.20
10	4.04	8%	5.72	22	2.86	5%	3.61
11	3.58	6%	4.94	23	2.62	5%	3.43
12	3.56	6%	4.86	24	2.36	5%	3.21

in Table 12. The resulting SARMA and GARCH$(2,1)$ models are

$$\begin{aligned}\left(1-1.55L+0.79L^2\right)\left(1-0.08L^{24}-0.26L^{48}\right)r_t& \\ =\left(1+1.39L-0.70L^2\right)\left(1-0.13L^{24}+0.08L^{48}\right)\epsilon_t& \quad (13)\end{aligned}$$

$$\sigma_t^2=0.00012+2.6585\epsilon_{t-1}^2-2.2343\epsilon_{t-2}^2+0.8680\sigma_{t-1}^2\,. \quad (14)$$

The autocorrelogram of the residuals is presented in Figure 12. The seasonality was partially captured by a SARMA model, as for the Austrian data. The absolute autocorrelation values are negligible, even if statistically significant, with the exception of lags 24, 25 and 48, related to seasonality. Nevertheless, we decided to work with a parsimonious model, which presents order 2 for both seasonal autoregressive and moving-average parts, and provides an acceptable forecast error.

In Figure 13, we present the autocorrelogram for the squared residuals. The results show that the non-linear dependence was successfully captured by GARCH. The aggregate error measures (RMSE, MAPE and MAE) are presented in columns 2–4 of Table 13. In addition to the overall results, this table presents measures for each hour separately, so we can compare the results with the Austrian case. These values were obtained in the same manner as for the Austrian data.

We see that for hours 2, 3, 4, 9 and 24 the MAPE values are very high. This is again due to very low values in the database. Including in-sample and out-of-sample periods, there are 5 missing values, 379 zeros, and 84 values that are very low (equal or less than €1/MWh), 468 observations in total. In the out-of-sample period, there are 32 missing values / zeros in the original data and 15 very low values. Excluding those values for the additional analysis of the forecasting performance, the aggregate error measures improve significantly, as we can see in the last three columns of Table 13. It is interesting to note that, just like in the Austrian case, the largest errors occur in the early morning hours.

Table 8: Aggregate measures for series with excluded small values, both approaches for the Austrian data

	24 separate series			1 entire series		
hour	**MAE**	**MAPE**	**RMSE**	**MAE**	**MAPE**	**RMSE**
1	3.44	9%	4.93	2.74	7%	3.93
2	4.65	16%	6.58	1.82	6%	2.5
3	7.28	22%	9.68	2.26	8%	4.45
4	5.28	20%	7.06	4.91	14%	43.03
5	7.54	24%	10.08	4.00	11%	40.88
6	6.01	17%	8.77	2.65	7%	3.95
7	5.23	16%	7.72	6.05	16%	8.85
8	7.89	17%	11.44	4.96	10%	7.49
9	5.55	12%	8.43	4.51	8%	28.68
10	4.04	8%	5.72	2.70	5%	9.14
11	3.58	6%	4.94	1.50	3%	3.09
12	3.56	6%	4.86	1.42	2%	1.98
13	3.54	6%	4.71	1.18	2%	1.66
14	3.76	7%	5.15	2.46	5%	3.06
15	3.62	7%	4.95	1.61	3%	2.08
16	3.61	7%	5.01	1.28	3%	1.66
17	3.66	7%	5.01	1.47	3%	2.06
18	3.65	6%	4.99	2.51	4%	3.74
19	3.79	6%	5.45	2.41	4%	3.45
20	3.47	6%	4.54	2.34	3%	3.74
21	3.21	5%	4.20	2.17	4%	2.92
22	2.86	5%	3.61	1.71	3%	2.3
23	2.62	5%	3.43	1.44	3%	1.89
24	2.36	5%	3.21	2.04	4%	2.55
Overall				**2.59**	**6%**	**14.03**

Table 9: Descriptive statistics of the Spanish data

Statistic	**Price** ($\frac{€}{\text{MWh}}$)	**Log return**
Mean	45.38	1.59e-06
Std. dev.	16.37	0.52
Variance	268.1	0.27
Skewness	0.24	0.75
Kurtosis	3.55	246.77
# of obs.	42720	42719

Table 10: Descriptive statistics for 24 series of electricity prices, Spain (1780 observations)

	Electricity price in €/MWh, per hour											
Statistic	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**
Mean	44.79	40.25	35.88	34.03	32.59	33.69	37.80	42.70	44.76	47.22	49.98	50.18
Std. Dev.	13.51	13.63	14.05	14.18	14.08	13.81	13.98	14.69	15.91	15.80	15.78	15.59
Variance	182.4	185.8	197.3	201.0	198.2	190.6	195.5	215.8	253.1	249.7	249.1	243.2
Skewness	0.14	-0.04	-0.19	-0.20	-0.18	-0.23	-0.25	0.01	0.04	0.19	0.31	0.33
Kurtosis	3.43	3.54	3.26	3.03	2.89	3.00	3.31	3.40	3.39	3.38	3.42	3.41
Statistic	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**	**24**
Mean	50.90	49.76	46.74	45.61	45.37	46.85	49.66	52.75	54.16	55.70	51.11	46.71
Std. Dev.	15.38	15.37	14.49	14.58	15.00	15.27	16.50	17.53	16.77	15.65	14.34	14.42
Variance	236.5	236.2	209.9	212.7	225.1	233.3	272.3	307.1	281.2	244.8	205.7	208.0
Skewness	0.33	0.29	0.24	0.10	0.10	0.20	0.56	0.57	0.58	0.68	0.52	0.27
Kurtosis	3.45	3.43	3.45	3.45	3.42	3.31	3.41	3.01	2.78	3.32	3.03	3.27

Table 11: Descriptive statistics for log-returns of the 24 series of Spanish electricity prices (1779 observations)

	Log-returns of electricity price, per hour											
Statistic	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**
Mean	4e-05	-4e-06	1e-04	2e-04	3e-04	3e-04	4e-04	7e-04	9e-04	9e-04	5e-04	4e-04
Std. Dev.	0.72	1.02	1.43	1.43	1.37	1.32	1.12	1.04	1.17	0.94	0.90	0.75
Variance	0.52	1.03	2.04	2.05	1.88	1.74	1.26	1.08	1.38	0.88	0.81	0.56
Skewness	-0.38	0.20	0.23	-0.45	-0.31	0.04	-0.41	0.69	0.38	-0.16	-0.60	0.39
Kurtosis	158.4	66.99	39.11	38.10	40.34	43.65	64.35	79.50	64.30	95.19	112.7	148.3
Statistic	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**	**24**
Mean	4.e-04	4e-04	4e-04	5e-04	5e-04	4e-04	2e-04	2e-04	4e-05	-4e-05	-1e-04	-2e-04
Std. Dev.	0.75	0.73	0.78	0.81	0.80	0.74	0.50	0.35	0.17	0.17	0.28	0.94
Variance	0.56	0.53	0.61	0.65	0.65	0.54	0.25	0.13	0.03	0.03	0.08	0.89
Skewness	1.31	1.47	0.94	-0.21	0.27	1.21	1.05	1.51	-1.68	-0.63	-0.89	-0.41
Kurtosis	153.5	154.3	135.7	126.5	119.7	140.7	327.6	480.1	33.86	20.27	257.6	119.4

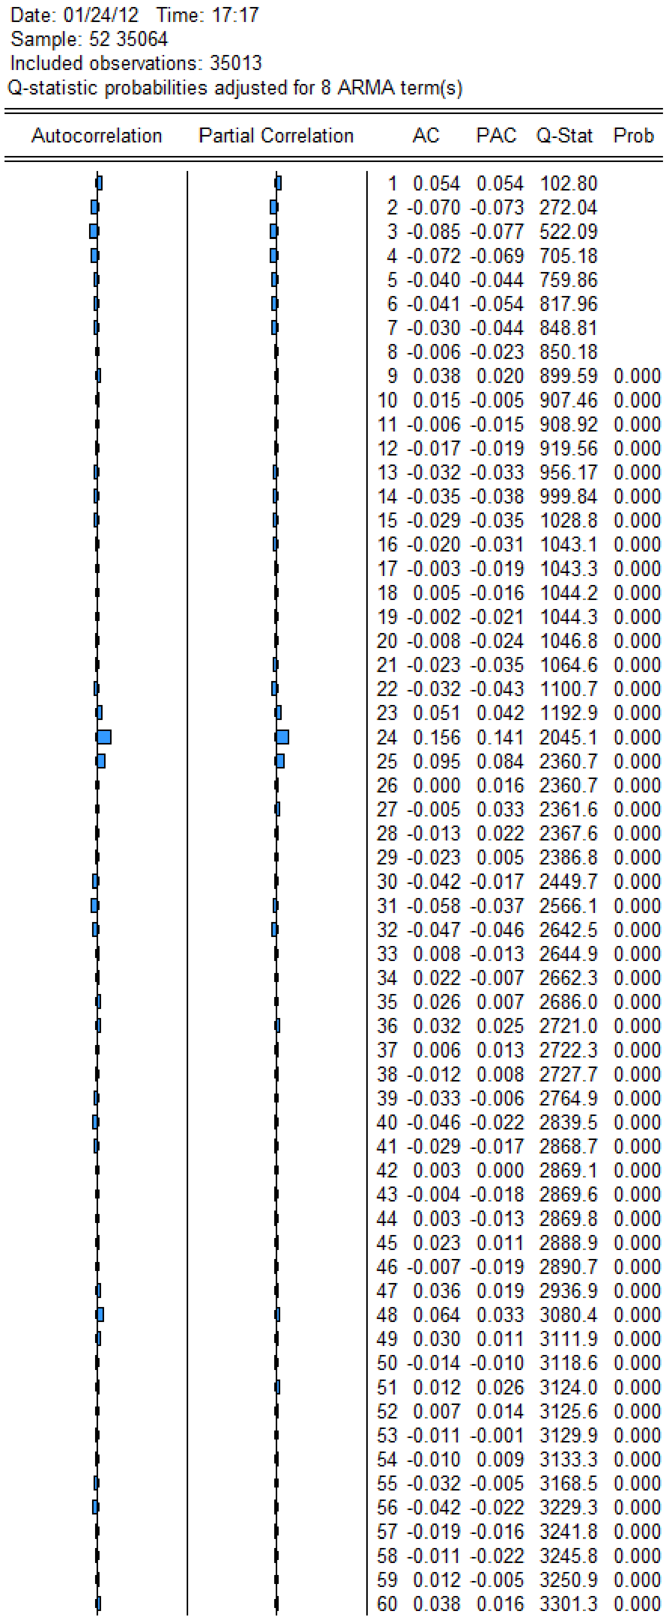

Date: 01/24/12 Time: 17:17
Sample: 52 35064
Included observations: 35013
Q-statistic probabilities adjusted for 8 ARMA term(s)

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.054	0.054	102.80	
		2	-0.070	-0.073	272.04	
		3	-0.085	-0.077	522.09	
		4	-0.072	-0.069	705.18	
		5	-0.040	-0.044	759.86	
		6	-0.041	-0.054	817.96	
		7	-0.030	-0.044	848.81	
		8	-0.006	-0.023	850.18	
		9	0.038	0.020	899.59	0.000
		10	0.015	-0.005	907.46	0.000
		11	-0.006	-0.015	908.92	0.000
		12	-0.017	-0.019	919.56	0.000
		13	-0.032	-0.033	956.17	0.000
		14	-0.035	-0.038	999.84	0.000
		15	-0.029	-0.035	1028.8	0.000
		16	-0.020	-0.031	1043.1	0.000
		17	-0.003	-0.019	1043.3	0.000
		18	0.005	-0.016	1044.2	0.000
		19	-0.002	-0.021	1044.3	0.000
		20	-0.008	-0.024	1046.8	0.000
		21	-0.023	-0.035	1064.6	0.000
		22	-0.032	-0.043	1100.7	0.000
		23	0.051	0.042	1192.9	0.000
		24	0.156	0.141	2045.1	0.000
		25	0.095	0.084	2360.7	0.000
		26	0.000	0.016	2360.7	0.000
		27	-0.005	0.033	2361.6	0.000
		28	-0.013	0.022	2367.6	0.000
		29	-0.023	0.005	2386.8	0.000
		30	-0.042	-0.017	2449.7	0.000
		31	-0.058	-0.037	2566.1	0.000
		32	-0.047	-0.046	2642.5	0.000
		33	0.008	-0.013	2644.9	0.000
		34	0.022	-0.007	2662.3	0.000
		35	0.026	0.007	2686.0	0.000
		36	0.032	0.025	2721.0	0.000
		37	0.006	0.013	2722.3	0.000
		38	-0.012	0.008	2727.7	0.000
		39	-0.033	-0.006	2764.9	0.000
		40	-0.046	-0.022	2839.5	0.000
		41	-0.029	-0.017	2868.7	0.000
		42	0.003	0.000	2869.1	0.000
		43	-0.004	-0.018	2869.6	0.000
		44	0.003	-0.013	2869.8	0.000
		45	0.023	0.011	2888.9	0.000
		46	-0.007	-0.019	2890.7	0.000
		47	0.036	0.019	2936.9	0.000
		48	0.064	0.033	3080.4	0.000
		49	0.030	0.011	3111.9	0.000
		50	-0.014	-0.010	3118.6	0.000
		51	0.012	0.026	3124.0	0.000
		52	0.007	0.014	3125.6	0.000
		53	-0.011	-0.001	3129.9	0.000
		54	-0.010	0.009	3133.3	0.000
		55	-0.032	-0.005	3168.5	0.000
		56	-0.042	-0.022	3229.3	0.000
		57	-0.019	-0.016	3241.8	0.000
		58	-0.011	-0.022	3245.8	0.000
		59	0.012	-0.005	3250.9	0.000
		60	0.038	0.016	3301.3	0.000

Figure 12: Autocorrelations for the residuals, Spain

Date: 01/24/12 Time: 17:21
Sample: 52 35064
Included observations: 35013
Q-statistic probabilities adjusted for 8 ARMA term(s)

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	-0.002	-0.002	0.1938	
		2	-0.001	-0.001	0.2374	
		3	0.002	0.002	0.3237	
		4	0.004	0.004	0.8206	
		5	-0.001	-0.001	0.8397	
		6	0.000	0.000	0.8464	
		7	-0.001	-0.001	0.8753	
		8	0.000	0.000	0.8811	
		9	-0.002	-0.002	0.9617	0.327
		10	-0.001	-0.001	1.0203	0.600
		11	-0.002	-0.002	1.1948	0.754
		12	-0.002	-0.002	1.3405	0.854
		13	-0.001	-0.001	1.3666	0.928
		14	0.001	0.001	1.3773	0.967
		15	-0.001	-0.001	1.4264	0.985
		16	0.001	0.001	1.4707	0.993
		17	-0.001	-0.001	1.5155	0.997
		18	-0.001	-0.001	1.5360	0.999
		19	-0.001	-0.001	1.6130	0.999
		20	-0.001	-0.001	1.6256	1.000
		21	-0.001	-0.001	1.6510	1.000
		22	0.000	0.000	1.6519	1.000
		23	0.001	0.001	1.7173	1.000
		24	0.006	0.006	3.1696	1.000
		25	0.007	0.007	4.6781	0.999
		26	-0.001	-0.001	4.7072	0.999
		27	0.000	-0.001	4.7158	1.000
		28	0.002	0.001	4.7996	1.000
		29	-0.001	-0.001	4.8458	1.000
		30	0.001	0.001	4.8677	1.000
		31	-0.002	-0.002	4.9623	1.000
		32	-0.001	-0.001	5.0405	1.000
		33	0.000	0.000	5.0418	1.000
		34	-0.002	-0.002	5.1654	1.000
		35	-0.001	-0.001	5.2430	1.000
		36	-0.001	-0.001	5.3165	1.000
		37	-0.002	-0.002	5.4128	1.000
		38	-0.001	-0.001	5.4577	1.000
		39	0.000	0.000	5.4624	1.000
		40	-0.002	-0.002	5.5700	1.000
		41	-0.001	-0.001	5.6210	1.000
		42	-0.001	-0.001	5.6315	1.000
		43	0.000	0.000	5.6351	1.000
		44	-0.001	-0.001	5.6812	1.000
		45	0.003	0.003	6.0574	1.000
		46	0.000	0.000	6.0575	1.000
		47	0.003	0.003	6.3745	1.000
		48	0.021	0.021	22.557	0.988
		49	0.000	0.000	22.565	0.991
		50	0.000	0.000	22.571	0.994
		51	0.000	0.000	22.576	0.996
		52	0.000	0.000	22.576	0.997
		53	-0.001	-0.001	22.614	0.998
		54	-0.001	-0.001	22.668	0.998
		55	-0.001	-0.001	22.732	0.999
		56	-0.001	-0.001	22.781	0.999
		57	-0.002	-0.002	22.871	0.999
		58	-0.002	-0.002	22.964	1.000
		59	-0.001	-0.001	23.040	1.000
		60	-0.002	-0.002	23.138	1.000

Figure 13: Autocorrelations for the squared residuals, Spain

Table 12: Estimated parameters for SARMA-GARCH models, Spain

Model	Coef.	Std. err.	z-stat.	Prob.
AR (1)	1.553	0.003385	458.915	0.000
AR (2)	−0.794	0.002591	−306.311	0.000
SAR(24)	0.082	0.005103	16.138	0.000
SAR(48)	0.255	0.002996	85.265	0.000
MA (1)	−1.388	0.003305	−420.035	0.000
MA (2)	0.702	0.002848	246.328	0.000
SMA(24)	0.132	0.005197	25.399	0.000
SMA(48)	−0.084	0.001969	−42.668	0.000
C	0.000	0.000003	43.570	0.000
ARCH(1)	2.659	0.011179	237.812	0.000
ARCH(2)	−2.234	0.010239	−218.226	0.000
GARCH(1)	0.868	0.000731	1187.177	0.000

Second approach: 24 separate series for each hour of the day

In this approach, we work with 24 time series, one for each hour of the day, as we presented for Austrian data. Just as before, we have tested many different models for each of the series, to find the one that gives the best fit, measured by the AIC and BIC criteria. The identified models are presented in Table 14.

We then use the fitted models to forecast prices for the out-of-sample period and measure the forecast errors; results are in Table 15. For hours between 2 and 10, and hour 24, MAPE measures are really high. This is again caused by several very low prices, which have a big impact on the measure. Moreover, the proposed ARMA+GARCH models for early hours in the morning do not fit the data very well.

Just like in the previous cases, excluding the very low prices for the data set improves the forecasting perfomance considerably, and shown in columns 2–4 of Table 16. The table includes also the error measures from the previous section for comparison. We can see that the model using all hours as one series is better in terms of the MAE and MAPE measures, for 23 out of 24 hours. The results are less clear when looking at the RMSE measure, where the per-hour approach performs better in the morning, and the one-series approach in the afternoon. On the whole, the one-series model performs slightly better than the per-hour models; it is also significantly less labour-intensive, as we need to fit one series instead of 24.

5.4. Summary of the analysis

In this section, we have presented, fitted and tested models under a time-series framework in order to analyze hourly electricity prices in Austria and Spain. For both countries, we tried to model the hourly prices either as one entire series, or as 24 separate series, one for each hour. Hence, we did, in total, 25 analyses for each country.

For both countries, the best-fitted models for all the series were SARMA + GARCH, with different orders. Analysing the error terms of the chosen models and the forecasting

Table 13: Aggregate measures for the SARMA-GARCH model for Spain

	Complete series			Excluding small values		
hour	**MAE**	**MAPE**	**RMSE**	**MAE**	**MAPE**	**RMSE**
1	6.08	13%	40.76	6.08	13%	40.76
2	3.18	53%	5.11	3.13	10%	5.02
3	3.76	1554%	6.30	3.74	21%	6.25
4	2.40	1357%	4.27	2.44	10%	4.32
5	2.15	9%	3.85	2.20	8%	3.90
6	2.84	10%	7.01	2.92	8%	7.11
7	6.19	16%	26.26	6.28	16%	26.47
8	4.90	16%	14.90	4.94	11%	14.99
9	4.84	1786%	19.03	4.90	12%	19.17
10	6.98	17%	29.22	7.00	17%	29.26
11	3.27	9%	9.19	3.27	9%	9.19
12	1.87	5%	7.11	1.87	5%	7.11
13	2.44	5%	5.05	2.44	5%	5.05
14	1.99	4%	4.10	1.99	4%	4.10
15	1.77	4%	3.15	1.77	4%	3.15
16	1.56	4%	2.63	1.56	4%	2.63
17	1.18	3%	2.04	1.18	3%	2.04
18	1.66	4%	3.75	1.66	4%	3.75
19	2.31	5%	4.35	2.31	5%	4.35
20	2.26	4%	3.89	2.26	4%	3.89
21	2.59	5%	5.05	2.59	5%	5.05
22	2.33	4%	4.98	2.33	4%	4.98
23	2.60	5%	4.85	2.60	5%	4.85
24	2.73	13382%	5.42	2.60	5%	4.87
Overall	**3.08**	**761%**	**13.36**	**3.09**	**8%**	**13.39**

accuracy measures, the entire-series approach provided the best results, compared to the 24 separate series—in addition to requiring significantly less work. In both cases, the forecasting power of the fitted model was best in the afternoon and worst in the early morning hours and the MAPE error measures improved significantly when we excluded the lowest prices from the test. Comparing results for the two countries, the Austrian forecast results were better than the Spanish.

Overall, we consider the quality of the predictions to be sufficient for our purposes, that is for generating scenarios for the optimization models. Should we find out that this is not the case, however, we might consider the following directions in refining the analysis:

- Since SARMA partially captured the seasonality, it is worth testing other alternatives to model it, using dummies or trigonometric functions. However, the gain in forecast accuracy might not be relevant.

Table 14: Best fitted models for 24 separate series, Spain

hour	Best fitted model
1	$\mathrm{SARMA}(1,1) \times (1,1)_7 + \mathrm{GARCH}(1,1)$
2	$\mathrm{SARMA}(1,1) \times (1,1)_7 + \mathrm{GARCH}(2,1)$
3	$\mathrm{SARMA}(1,2) \times (1,1)_7 + \mathrm{GARCH}(1,1)$
4	$\mathrm{SARMA}(2,1) \times (1,1)_7 + \mathrm{GARCH}(1,1)$
5	$\mathrm{SARMA}(1,1) \times (1,1)_7 + \mathrm{GARCH}(2,1)$
6	$\mathrm{SARMA}(2,3) \times (2,2)_7 + \mathrm{GARCH}(2,2)$
7	$\mathrm{ARMA}(2,2) + \mathrm{GARCH}(2,1)$
8	$\mathrm{SARMA}(2,1) \times (2,2)_7 + \mathrm{GARCH}(2,1)$
9	$\mathrm{SARMA}(2,2) \times (1,1)_7 + \mathrm{GARCH}(2,2)$
10	$\mathrm{SARMA}(2,2) \times (1,1)_7 + \mathrm{GARCH}(2,1)$
11	$\mathrm{SARMA}(2,2) \times (1,1)_7 + \mathrm{GARCH}(1,1)$
12	$\mathrm{SARMA}(1,1) \times (1,1)_7 + \mathrm{GARCH}(1,1)$
13	$\mathrm{SARMA}(2,1) \times (1,1)_7 + \mathrm{GARCH}(2,1)$
14	$\mathrm{SARMA}(1,2) \times (1,1)_7 + \mathrm{GARCH}(2,1)$
15	$\mathrm{SARMA}(1,1) \times (1,0)_7 + \mathrm{GARCH}(1,1)$
16	$\mathrm{SARMA}(1,2) \times (1,1)_7 + \mathrm{GARCH}(2,1)$
17	$\mathrm{SARMA}(2,1) \times (1,1)_7 + \mathrm{GARCH}(2,1)$
18	$\mathrm{SARMA}(1,1) \times (2,2)_7 + \mathrm{GARCH}(1,0)$
19	$\mathrm{SARMA}(1,2) \times (1,1)_7 + \mathrm{GARCH}(1,1)$
20	$\mathrm{SARMA}(1,1) \times (1,1)_7 + \mathrm{GARCH}(2,1)$
21	$\mathrm{SARMA}(1,1) \times (1,1)_7 + \mathrm{GARCH}(2,1)$
22	$\mathrm{SARMA}(1,1) \times (1,1)_7 + \mathrm{GARCH}(2,1)$
23	$\mathrm{SARMA}(2,1) \times (1,1)_7 + \mathrm{GARCH}(1,2)$
24	$\mathrm{SARMA}(2,1) \times (1,1)_7 + \mathrm{GARCH}(1,0)$

Table 15: Aggregate error measures for the per-hour series, Spanish data

hour	MAE	MAPE	RMSE
1	3.33	7%	4.83
2	5.56	109%	8.44
3	10.23	3612%	21.95
4	8.52	26955%	11.57
5	9.26	21452%	12.97
6	8.15	9032%	11.52
7	13.21	18142%	40.47
8	6.65	23925%	10.83
9	11.19	33176%	20.17
10	10.30	117%	12.78
11	5.90	15%	8.91
12	6.31	14%	7.98

hour	MAE	MAPE	RMSE
13	4.55	11%	7.37
14	3.83	9%	6.10
15	4.54	12%	7.02
16	4.69	18%	7.92
17	4.88	17%	7.46
18	5.46	14%	8.67
19	4.40	10%	6.51
20	3.97	8%	6.26
21	3.15	6%	4.96
22	3.15	5%	5.47
23	2.82	5%	4.35
24	3.75	15600%	6.86

Table 16: Aggregate measures for series with excluded small values, both approaches for the Spanish data

	24 separate series			1 entire series		
hour	**MAE**	**MAPE**	**RMSE**	**MAE**	**MAPE**	**RMSE**
1	3.33	7%	4.83	6.08	13%	40.76
2	5.46	17%	8.18	3.13	10%	5.02
3	10.28	49%	22.07	3.74	21%	6.25
4	8.44	30%	11.21	2.44	10%	4.32
5	9.27	43%	12.79	2.20	8%	3.90
6	8.08	28%	11.23	2.92	8%	7.11
7	13.13	38%	40.31	6.28	16%	26.47
8	6.36	20%	10.12	4.94	11%	14.99
9	10.84	23%	19.66	4.90	12%	19.17
10	10.30	117%	12.79	7.00	17%	29.26
11	5.90	15%	8.91	3.27	9%	9.19
12	6.31	14%	7.98	1.87	5%	7.11
13	4.55	11%	7.37	2.44	5%	5.05
14	3.83	9%	6.10	1.99	4%	4.10
15	4.54	12%	7.02	1.77	4%	3.15
16	4.69	18%	7.92	1.56	4%	2.63
17	4.88	17%	7.46	1.18	3%	2.04
18	5.46	14%	8.67	1.66	4%	3.75
19	4.40	10%	6.51	2.31	5%	4.35
20	3.97	8%	6.26	2.26	4%	3.89
21	3.15	6%	4.96	2.59	5%	5.05
22	3.15	5%	5.47	2.33	4%	4.98
23	2.82	5%	4.35	2.60	5%	4.85
24	3.60	7%	6.28	2.60	5%	4.87
Overall				**3.09**	**8%**	**13.39**

- We recommend outlier treatment to obtain better forecasts, regarding the quantity of missing, zero or low-value observations.

6. Conclusions

This report presents the first deliverable of WP3 of the EnRiMa project, whose goal is to generate scenarios for stochastic optimization models being developed as part of WP4. It is, in particular, based on findings presented in deliverable D4.1, which identified the stochastic parameters needed by the optimization models.

In this report, we describe the first steps needed for the scenario-generation process: data sources, data series, and their analyses. In particular, we present detailed time-series analyses of hourly electricity prices from Austria and Spain. The time-series models identified there will be needed for the scenario-generation procedure that is being developed in WP3. In addition, we describe how to deal with parameters where straightforward data analysis is not possible or optimal to do, such as the weather and the energy loads of the buildings. The next step to fulfilling the goals of WP3 is to use the results from this report to actually generate the scenarios; this task will result in deliverable D3.2.

We would like to point out that these results should not be considered final: it can be expected that the optimization models will change during their development and testing and that those changes might trigger the need to change some of the input data, or even to add new data series. Such changes would then be described in a future report of work package 4.

Acknowledgments

This report was jointly written by Michal Kaut, Adrian T. Werner, Fernando Aiube, Tara Baidya, Frances Blank, Ana B. Mattos, Wagner Saboia, and Afzal Siddiqui. The data discussed in this deliverable were provided by CET and HCE, while SU was responsible for the time-series analyses. All other partners also contributed with valuable input about available data sources and suggestions how to handle the various types of uncertainty in the context of the EnRiMa DSS.

The manuscript has been subject to an internal project quality control by IIASA, URJC, and UCL.

References

F. A. L. Aiube. Econometria para séries financeiras. Technical report, Pontifícia Universidade Católica do Rio de Janeiro, 2007.

S. Balada, H. Erhorn, T. Dr. Hartmann, C. Händel, K. Dr. Jagnow, A. Dr. Maas, B. Prof. Dr. Oschatz, W. Reiners, H. Schiller, J. Schilling, and J. Dr. de Boer. Energy efficiency of buildings – Calculation of the net, final and primary energy demand for heating, cooling, ventilation, domestic hot water and lighting. Technical Report DIN V 18599, 2007.

J. R. Birge and F. Louveaux. *Introduction to stochastic programming.* Springer-Verlag, New York, 1997. ISBN 0-387-98217-5.

R. D. L. d. S. Bueno. *Econometria de Séries Temporais.* Cengage, São Paulo, first edition, 2008.

A. K. Dixit and R. S. Pindyck. *Investment under uncertainty.* Princeton University Press, Princeton, NJ, 1994.

J. Dupačová, G. Consigli, and S. W. Wallace. Scenarios for multistage stochastic programs. *Annals of Operations Research*, 100(1–4):25–53, 2000. ISSN 0254-5330. doi: 10.1023/A:1019206915174.

S. Ekern and P. Bjerksund. Contingent claims evaluation of mean-reverting cash flows in shipping. chapter 12, pages 207–220. Praeger, 1995.

V.-F. T. Gebäudeausrüstung. Requirements on methods of calculation to thermal and energy simulation of buildings and plants. buildings. Technical Report VDI 6020-1, VDI–Gesellschaft Bauen und Gebäudetechnik, 2001.

M. Groissböck, M. Stadler, T. Edlinger, A. Siddiqui, S. Heydari, and E. Perea. The first step for implementing a stochastic based energy management system at campus Pinkafeld. Technical Report C-2011-1, Center for Energy and innovative Technologies, Hofamt Priel, Austria, 2011.

J. D. Hamilton. *Time Series Analysis*. Princeton University Press, first edition, 1994.

S. Heydari and A. Siddiqui. Valuing a gas-fired power plant: A comparison of ordinary linear models, regime-switching approaches, and models with stochastic volatility. *Energy Economics*, 32(3):709–725, 2010. doi: 10.1016/j.eneco.2009.10.001.

K. Høyland and S. W. Wallace. Generating scenario trees for multistage decision problems. *Management Science*, 47(2):295–307, 2001. doi: 10.1287/mnsc.47.2.295.9834.

R. Huisman, C. Huurman, and R. Mahieu. Hourly electricity prices in day-ahead markets. *Energy Economics*, 29(2):240–248, 2007. doi: 10.1016/j.eneco.2006.08.005.

P. Kall and S. W. Wallace. *Stochastic Programming*. John Wiley & Sons, Chichester, 1994.

E. Schwartz. The stochastic behaviour of commodity prices: Implications for valuation and hedging. *Journal of Finance*, 52(3):923–73, 1997.

R. S. Tsay. *Analysis of Financial Time Series*. John Wiley, New York, first edition, 2002.

M. Šúri, T. A. Huld, E. D. Dunlop, and H. A. Ossenbrink. Potential of solar electricity generation in the european union member states and candidate countries. *Solar Energy*, 81(10):1295–1305, 2007. doi: 10.1016/j.solener.2006.12.007.

Appendix A. Austrian prices using a mean-reverting process

In this Appendix, we analyse historical data for hourly electricity prices in Austria. We propose a model based on mean-reverting stochastic process for the logarithms of hourly prices. The tested data set consists of spot prices in Austria/Germany markets, provided by EXAA, for the period of January 1st, 2007 (hour 1) to November 30th, 2011 (hour 24).

We split the data set into two periods: from January 1st, 2007, hour 1, to December 31st, 2010, hour 24, which is used as in-sample period, and from January 1st, 2011, hour 1, to November 30th, 2011, hour 24, which is used as out-of-sample period.

The logarithm of electricity price is decomposed in a sum of two factors: a deterministic seasonal function and a stochastic part. The deterministic part includes intraday effects, and weekly and yearly seasonality. To model the stochastic part, a mean-reverting process is used.

The resulting model will be applied for the hourly series, to obtain a forecast model for electricity prices. The unknown parameters are estimated using the in-sample period data. The out-of-sample period data is used to evaluate the forecasting performance of the models.

In section 2, we present a description of the data used and some descriptive analysis. In section 3, the seasonal function is modelled and some results related to seasonality are discussed. In section 4, the stochastic process is modelled and the parameters are estimated. In section 5, we provide the forecasting results, which are compared to the results obtained from working with 24 separate series..

1. Data and descriptive statistics

In total 43,080 hourly observations over five years of electricity spot prices in €/MWh from Austria/Germany markets, provided by Energy Exchange Austria (EXAA) are available. The sample period begins on January 1st, 2007, hour 1, and ends on November 30th, 2011, hour 24. In order to compare the results presented here with other analyses, we removed the observations with missing values or values equal or less than 1 €/MWh, in a total of 189 observations in the sample (less than 0.5% of the number of observations).

The data set is split into two periods:

- an in-sample period, from 1 January 1st, 2007 to December 31st, 2010 (34,890 valid hourly observations, in a total of 35,064 observations), which is used to estimate the unknown parameters, and
- an out-of-sample period, from January 1st, 2011 to November 30th, 2011 (8001 valid hourly observations, in a total of 8016 observations), which is used to assess the forecast of the model proposed.

According to Heydari and Siddiqui (2010), the natural logarithms of spot prices are decomposed into two factors:

$$\ln(S_t) = X_t + f_t \,, \tag{1}$$

where

S_t – is the energy spot price observed;

X_t – is the stochastic part of log prices; and

f_t – is a deterministic seasonal function.

Considering one series with all the hourly observations, a summary of descriptive statistics of electricity prices and log electricity prices is presented in Table 1.

Table 1 – Summary of descriptive statistics for electricity prices (EUR/MWh) and ln electricity prices

Statistic	Electricity	Ln electricity
Mean	48.41	3.76
Standard Deviation	23.76	0.51
Variance	564.66	0.26
Skewness	2.47	-1.00
Kurtosis	26.32	6.66
# of observations	42,891	42,891

According to EXAA, the hours 9 to 20 are considered peak hours. Off-peak hours include hours 1 to 8 and from hours 21 to 24.

2. Seasonality

Since we have one aggregated hourly series, now the seasonality treatment should include intraday effects, as well as weekly and yearly seasonality. Now, we work with 23 hourly dummies to explain intraday effect, 2 weekdays dummies (one for Monday to Friday and one to Saturday) to explain weekly seasonality and 11 monthly dummies to explain yearly seasonality. The model is represented by the following equation:

$$\begin{aligned} \ln(S_t) = (\alpha + \varepsilon_t) + \{(I_W D_W + I_S D_S) + \\ + \{(I_1 D_1 + I_2 D_2 + I_4 D_4 + I_5 D_5 + \cdots + I_{12} D_{12}) \\ + (A_1 H_1 + A_2 H_2 + A_3 H_3 + A_5 H_5 + \cdots + A_{24} H_{24})\} \end{aligned} \quad (2)$$

where:

$\ln(S_t)$ – is the natural logarithm of the daily energy price;

α – is the independent coefficient of the regression;

D_i – are the binary variables (dummies) related to the weekdays. We define three types of days: type 1: Sundays, type 2: Monday to Friday and type 3: Saturdays.

$$D_i = \begin{cases} 1, if \ \ is\ of\ type\ i \\ 0, for\ other\ types \end{cases} \quad i = 2(Monday\ to\ Friday), 3(Saturday)$$

I_i – are the linear coefficients of the regression for weekdays, $i = 2{,}3$

Sundays are used as reference, so there are no dummies or linear coefficients related to this weekday.

D_j – are the binary variables (Dummies) related to the months of the year.

$$D_j = \begin{cases} 1, for\ month\ j \\ 0, for\ other\ months \end{cases} \quad j = 1(January), 2(Feburary), \dots, 12(December)$$

I_j – are the linear coefficients of the regression for the months of the year, $j = 1,2,4,5, \ldots ,12$

The month of March is used as reference, so there are no dummies or linear coefficients related to this month.

A_l – are the binary variables (Dummies) related to the hours of the day.

$$A_l = \begin{cases} 1, for\ hour\ l \\ 0, for\ other\ hours \end{cases} \quad j = 1(hour\ 1), 2(hour\ 2), \ldots ,24\ (hour\ 24)$$

H_l – are the linear coefficients of the regression for the hours of the day, $l = 1,2,3,5, \ldots ,24$

In this regression, we did not use the coefficient related to slope since it was not statistically significant for all analysis.

We execute the regression of Eq.(3) to estimate the coefficients. Afterwards, the seasonal effects are eliminated as follows:

$$\begin{aligned} X_t = \ln(S_t) - (\widehat{I_1}D_1 + \widehat{I_2}D_2 + \widehat{I_4}D_4 + \widehat{I_5}D_5 + \cdots + \widehat{I_{12}}D_{12} + \widehat{I_W}D_W + \widehat{I_S}D_S \\ + \widehat{A_1}H_1 + \widehat{A_2}H_2 + \widehat{A_3}H_4 + \widehat{A_5}H_5 + \cdots + \widehat{A_{24}}H_{24}) \end{aligned} \quad (3)$$

The estimated parameters are shown below in Table 2.

Table 2 - Results for entire hourly series regression

Parameter	*Estimated Value*	*Parameter*	*Estimated Value*
α	2.580	A_7	0.353
I_W	0.475	A_8	0.646
I_S	0.275	A_9	0.787
I_1	0.147	A_{10}	0.879
I_2	0.150	A_{11}	0.932
I_4	0.053	A_{12}	0.997
I_5	-0.002	A_{13}	0.929
I_6	0.107	A_{14}	0.861
I_7	0.114	A_{15}	0.792
I_8	0.037	A_{16}	0.737
I_9	0.239	A_{17}	0.735
I_{10}	0.407	A_{18}	0.832
I_{11}	0.285	A_{19}	0.905
I_{12}	0.180	A_{20}	0.877

Parameter	*Estimated Value*	*Parameter*	*Estimated Value*
A_1	0.399	A_{21}	0.799
A_2	0.243	A_{22}	0.697
A_3	0.119	A_{23}	0.673
A_5	0.018	A_{24}	0.500
A_6	0.184		

Fig. (1) shows the estimated monthly seasonal parameters and Fig.(2) shows the estimated hourly parameters.

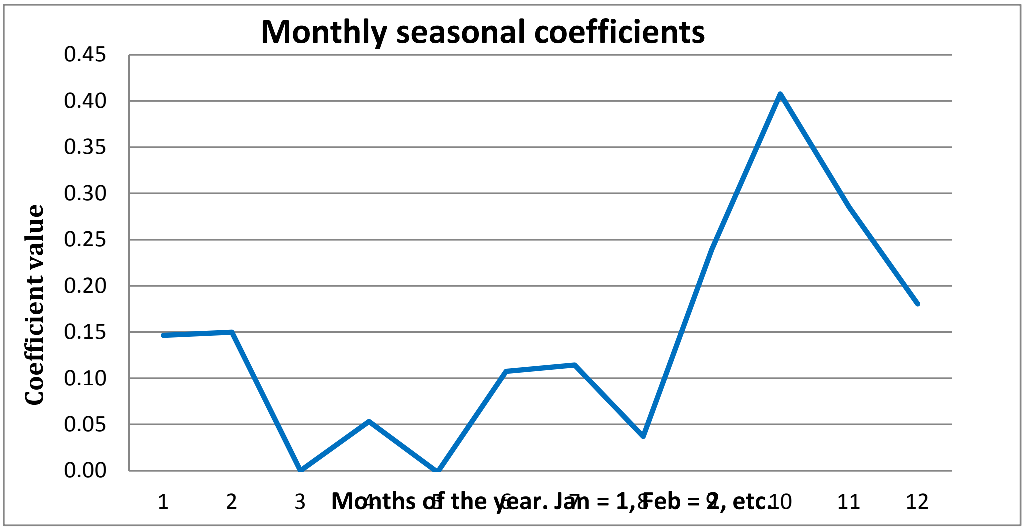

Fig. 1 – Monthly seasonal coefficients

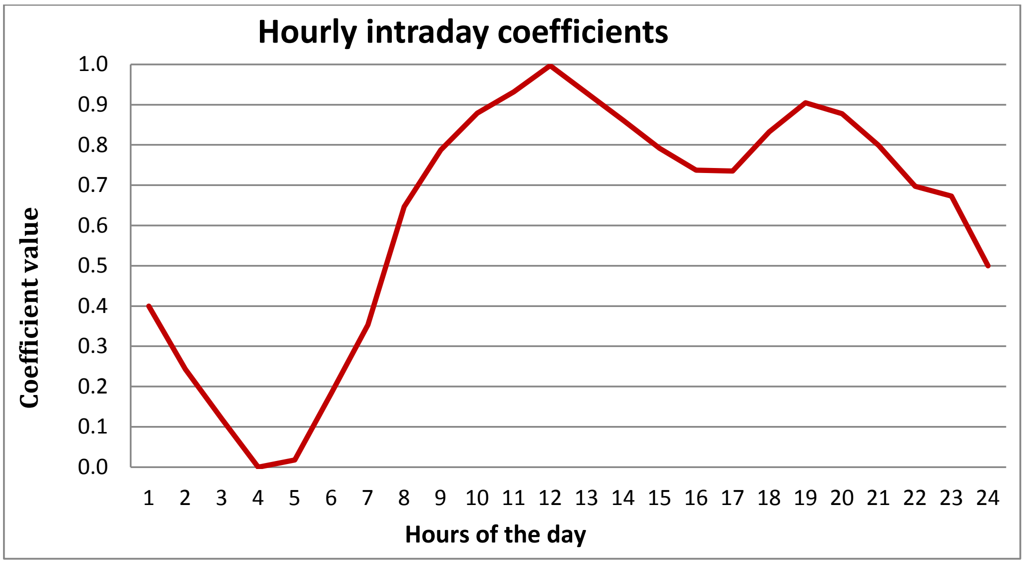

Fig. 2 – Hourly intraday coefficients

The monthly coefficients curve presents an expected shape. We see higher values for the months by the end of the year (September, October, November and December) and lower values for the months of March, April and May. The hourly coefficients curve also presents an expected shape based on the prices levels for peak hours (hours 9 to 20), which are higher, and off-peak hours.

3. Stochastic Linear Model

As explained before, after the elimination of seasonal effects we use an Ornstein-Uhlenbeck stochastic process to model the natural logarithm of energy prices for each hour of the day:

$$dX_t = \kappa(\alpha - X_t)dt + \sigma dz_t \,, \tag{4}$$

where:

> X_t – is the natural logarithm of the hourly electricity price, without seasonal effects;
>
> κ – is the magnitude of the speed of adjustment, which measures the degree of mean reversion to the long-run log price. $\kappa > 0$;
>
> α – is the long-run mean natural logarithm of the price;
>
> σ – is the term of volatility of the process;
>
> dz_t – is the increment of a standard Brownian motion.

In this model, as we are working with one entire hourly series, we will have only one parameter of long-run mean and only one of speed of adjustment.

According to Dixit and Pindyck (1994), Eq.(4) is the continuous-time version of the first-order autoregressive process in discrete time. The limiting case as $\Delta t \rightarrow 0$ is:

$$x_t - x_{t-1} = \alpha(1 - e^{-\kappa}) + (e^{-\kappa} - 1)x_{t-1} + \epsilon_t \quad (5)$$

where ϵ_{h_t} is normally distributed with mean zero and standard deviation $\sigma_{\epsilon h}$, and

$$\sigma_\epsilon^2 = \frac{\sigma^2}{2\kappa}(1 - e^{-2\kappa}) \quad (6)$$

Again, we estimate the parameters of Eq.(4) using the discrete time data available to run the regression of Eq.(7),

$$x_t - x_{t-1} = a + bx_{t-1} + \epsilon_t \quad (7)$$

and then calculate for each hour h:

$$\hat{\alpha} = -\hat{a}/\hat{b} \quad (8)$$

$$\hat{\kappa} = -\ln(1 + \hat{b}) \quad (9)$$

$$\hat{\sigma} = \hat{\sigma}_\varepsilon \sqrt{\frac{2\ln(1 + \hat{b})}{(1 + \hat{b})^2 - 1}} \quad (10)$$

Where $\hat{\sigma}_\varepsilon$ is the standard error of the regression of Eq. (7).

In Table 3, we have the results of the regression. Both parameters are significant at a 5% confidence level. As we reject the null hypothesis of $b = 0$, we can also reject the hypothesis of unit root (random walk), what reinforces the idea of mean-reversion in log hourly prices.

Table 3 - Results of the regression of Eq. (7)

	Coefficients	*Std Error*	*t Stat*	*P Value*
$\hat{a}$	0.221	0.006	39.066	0.000
$\hat{b}$	-0.086	0.002	-39.534	0.000

The application of Eqs. (8) to (10) give the estimated values for mean-reversion parameters as follows in Table (4).

Table 4 - Estimated mean-reversion parameters for hourly price series

$\alpha = 2.580$	$\kappa = 0.090$	$\sigma = 0.120$

4. Forecasting

The value of the stochastic process in Eq. (5) for a future date T, $X(T)$, conditional on the initial value $X(0)$, may be written as the stochastic integral (Bjerksund and Ekern, 1995):

$$X(T) = e^{-\kappa T}X(0) + (1 - e^{-\kappa T})\alpha + \sigma e^{-\kappa T}\int_0^T e^{-\kappa u}dZ(u) \tag{11}$$

$X(T)$ is normally distributed, and its expected value and its variance, are given by:

$$\mathbb{E}_0[X(T)|X(0)] = e^{-\kappa T}X(0) + (1 - e^{-\kappa T})\alpha \tag{12}$$

$$\mathbb{V}ar_0[X(T)|X(0)] = \frac{\sigma^2}{2\kappa}(1 - e^{-2\kappa T}) \tag{13}$$

Fig. (3) presents the real hourly log prices without seasonal effects (dotted line), the expected value of log daily average prices and the confidence interval of 95% (continuous lines).

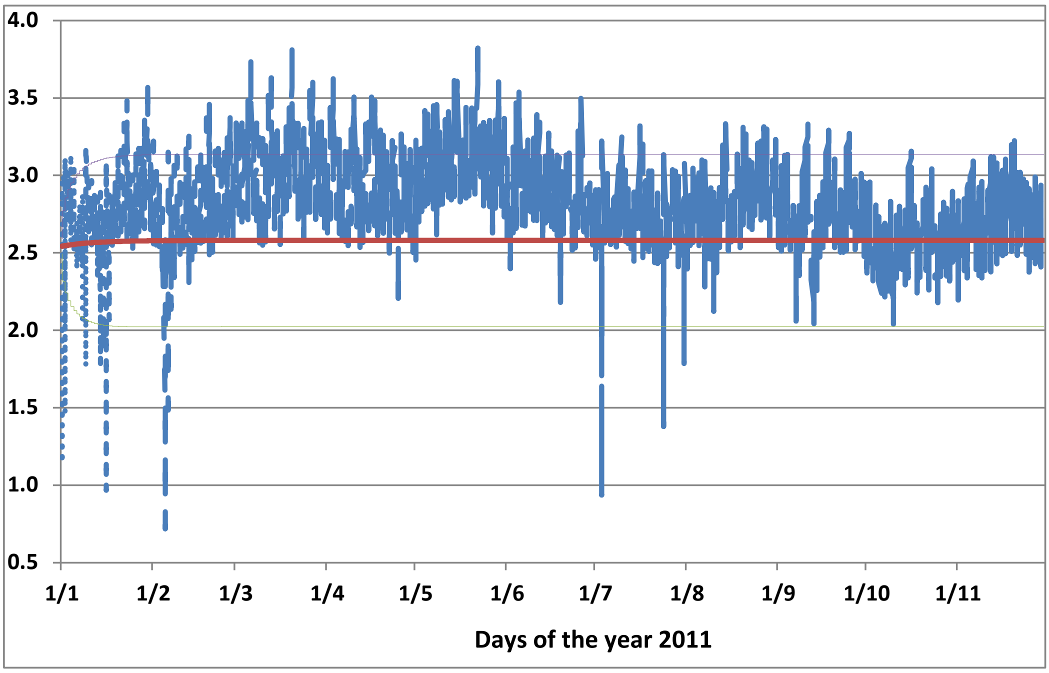

Fig. 3 – Results of forecast for the year 2011 given by Eq. (11)

As discussed before, $X(T)$ is normally distributed, with mean and variance given by Eqs. (12) and (13), respectively. So we can find the expected value of daily prices for each hour of the day:

$$\mathbb{E}[S_t|X(0)] = \mathbb{E}[e^{f_t}\, e^{X_t}|X(0)] = e^{\{f_t+\mathbb{E}_0[X(T)]+\frac{1}{2}\mathbb{V}ar_0[X(T)]\}}, \tag{14}$$

where e^{f_t} is deterministic and e^{X_t} is stochastic. $\mathbb{E}[\boldsymbol{S_t}|X(0)]$ is the conditional forecast of prices of electricity and can be compared to the actual prices $\boldsymbol{S_t}$.

These results provide forecasts for each observation during the out-of-sample period. This means that we have hourly forecasts in sequence starting on January 1st, 2011 hour 1 until November 31st, 2011 hour 24.

We calculate measures of aggregate error and present the results for the forecasted series. Besides, in order to compare the accuracy of the model to the one we worked with 24 separate series, we also present here measures for each hour separately from this model which works with an entire hourly series. To obtain this result, we took the entire hourly forecasted series and separated the forecasts in 24 series, one for each hour of the day. For example, the aggregate error for hour 1 series is calculated from the forecasted series composed by $\mathbb{E}[S_1]$, $\mathbb{E}[S_{25}]$, $\mathbb{E}[S_{49}]$ etc. The results are presented in Table 5.

Considering the measures of aggregate error, there is no significant gain to work with one entire hourly series compared to 24 separate series. The results are very similar. We observe again that the accuracy of the model is better for peak hour series, although still not satisfactory. The accuracy is worst for the off-peak hours in the morning. The real gain is that, instead of working with 24 different models, it can be easier to work and adjust only one model.

Table 5 – Aggregate error measures

	24 series						1 series				
	Ln Prices			Prices			Prices				
Hour	MAE	MAPE	RMSE	MAE	MAPE	RMSE	MAE	MAPE	RMSE	Dif MAE	Dif MAPE
H01	0.31	8.2%	0.35	11.00	24.5%	12.45	11.43	25.3%	13.54	-0.43	-0.7%
H02	0.47	12.8%	0.55	14.13	35.1%	16.28	12.70	31.5%	14.62	1.43	3.6%
H03	0.46	12.6%	0.51	13.13	34.8%	14.82	13.36	34.8%	15.15	-0.23	0.1%
H04	0.52	15.0%	0.59	13.82	39.9%	15.59	14.47	40.4%	16.22	-0.65	-0.6%
H05	0.51	14.2%	0.57	13.83	37.2%	15.60	14.30	37.9%	16.07	-0.48	-0.7%
H06	0.42	11.3%	0.49	12.30	29.9%	14.11	13.01	31.2%	14.76	-0.71	-1.3%
H07	0.39	10.7%	0.48	12.61	29.0%	14.55	14.61	33.2%	16.28	-1.99	-4.2%
H08	0.29	7.7%	0.38	11.33	23.0%	13.48	12.61	24.8%	14.73	-1.27	-1.8%
H09	0.22	5.6%	0.28	10.03	18.7%	12.20	9.97	18.3%	12.32	0.06	0.3%
H10	0.17	4.2%	0.22	8.48	15.1%	10.66	7.90	14.1%	10.06	0.58	1.0%
H11	0.14	3.6%	0.18	8.12	13.7%	10.13	7.62	12.8%	9.68	0.50	0.9%
H12	0.20	4.9%	0.35	10.58	17.7%	13.87	8.74	14.5%	11.07	1.84	3.2%
H13	0.14	3.4%	0.18	7.81	13.3%	9.61	7.99	13.6%	10.21	-0.19	-0.3%
H14	0.15	3.7%	0.19	7.70	13.8%	9.59	7.52	13.4%	9.62	0.18	0.4%
H15	0.16	4.1%	0.21	7.89	14.7%	9.78	7.64	14.0%	9.66	0.26	0.7%
H16	0.18	4.6%	0.23	8.19	15.7%	10.10	7.79	14.7%	9.65	0.40	1.0%
H17	0.20	5.0%	0.25	8.59	16.6%	10.56	7.96	15.0%	9.77	0.62	1.6%
H18	0.21	5.2%	0.26	9.82	17.4%	11.97	8.66	14.7%	11.08	1.16	2.7%
H19	0.20	4.9%	0.24	10.29	16.8%	12.32	9.80	15.1%	12.63	0.49	1.7%
H20	0.21	5.1%	0.25	10.98	17.4%	13.03	9.32	14.7%	11.90	1.65	2.7%
H21	0.22	5.4%	0.26	11.03	18.1%	12.79	9.93	16.3%	12.68	1.10	1.8%
H22	0.22	5.5%	0.25	10.22	17.9%	11.86	10.69	18.9%	13.35	-0.47	-1.0%
H23	0.22	5.4%	0.25	9.82	17.6%	11.33	10.90	19.8%	13.53	-1.08	-2.2%
H24	0.27	7.0%	0.30	10.76	21.5%	12.15	11.09	22.4%	13.31	-0.33	-0.9%
Overall	0.27	7%	0.33	10.52	21.6%	12.45	10.41	21.3%	12.77	0.11	0.4%

5. Conclusions

We implemented a similar mean-reversion methodology we have been using to model electricity prices, but now considering an hourly series, where the observations are hourly prices in sequence.

We decomposed the logarithms of electricity price series in two parts. The first part was related to a deterministic function to model daily, weekly and yearly seasonality. We observed an expected shape in both monthly and hourly coefficients curves. In monthly coefficients curve, we obtained higher values for the months of September, October, November and December and lower values for the months of March, April and May. Regarding the hourly coefficients curve, we observed higher price levels in peak hours (hours 9 to 20).

After the elimination of seasonal effects we used a mean-reverting stochastic process to model the natural logarithm of electricity prices. Using the results of parameters' estimation we calculated the forecast for the stochastic process $\boldsymbol{X_t}$, and then the forecasts for the hourly electricity prices. Some measures of aggregate error are presented, in order to compare the results with the results obtained from working with 24 separate series.

Considering the measures of aggregate error, there is no significant gain to work with one entire hourly series compared to 24 separate series. The results are very similar. We observe again that the accuracy of the model is better for peak hour series, although still not satisfactory. The accuracy is worst for the off-peak hours in the morning. The real gain is that, instead of working with 24 different models, it can be easier to work and adjust only one model.