Politecnico di Torino

Master's Degree in Mathematical Engineering

Academic Year 2025/2026

Statistical models for the electricity market

Focus on load and market prices

Supervisor:
Mauro Gasparini

Candidate:
Luca Innocenti

Contents

List of Figures

List of Tables

Abstract

The ongoing transition of the European energy sector and the increasing complexity of grid management have made accurate forecasting within electricity markets a critical necessity. Short-term electricity load forecasting is fundamentally important for Transmission System Operators to efficiently schedule secondary energy resources, balance supply and demand, and maintain grid frequency at its nominal value. This thesis addresses these operational challenges through a dual-focus approach: modeling short-term electricity load dynamics in Italy and conducting a comparative econometric analysis of electricity prices across distinct European markets.

Methodologically, this research employs robust statistical time series forecasting techniques, specifically utilizing Exponential Smoothing State Space Models, ETS, Autoregressive Integrated Moving Average, ARIMA, and Generalized Autoregressive Conditional Heteroskedasticity, GARCH, models.

The first part of the study focuses on short-term demand forecasting for the Italian electricity grid. To capture spatial and structural complexities, the analysis evaluates the grid both at the aggregate national level and with a specific focus on the major islands. Because these insular regions rely on isolated, independent secondary energy resources, precise short-term load prediction is crucial for timely resource dispatching and ensuring local grid stability.

The second part of the thesis broadens the analytical scope to a cross-national comparison of electricity price time series. The study examines Italy alongside France, Germany, and Denmark. These specific nations were strategically selected to represent fundamentally different energy generation portfolios: Denmark takes advantage of abundant, although intermittent renewable energy sources; France is characterized by a massive reliance on stable nuclear baseload generation; Germany is characterized by a large usage of fossil fuels; And Italy relies mostly on imported natural gas.

By applying the selected statistical models, particularly leveraging GARCH to capture price volatility clustering, the study highlights how varying degrees of renewable integration and differing base load technologies directly impact market price dynamics, seasonality, and overall volatility. Ultimately, this research demonstrates the continued efficacy of traditional time series models in modern grid management, providing valuable insights into both the technical requirements for maintaining stability in isolated networks and the economic implications of diverse national energy policies across Europe.

Chapter 1

Introduction

1.1 Electricity Supply Chain

Over the last few decades, the electricity sector has moved away from the traditional vertically integrated model toward a more segmented and liberalized structure. Before the major reform introduced in Italy in 1999, the national electricity system was managed by a single entity that controlled the entire value chain, from power generation to final sales to consumers.

The liberalization process progressively separated this structure into four main phases, each characterized by specific economic mechanisms, technical constraints, and operational responsibilities:

- **Generation**: Generation refers to the conversion of primary energy sources, either fossil-based or renewable, into electricity. In the current market structure, this activity takes place in a competitive environment, where producers compete to inject electricity into the grid according to their marginal costs, technical availability, and access to primary resources. Since generation depends on plant constraints, fuel availability, and, increasingly, weather-dependent renewable sources, it represents one of the main sources of supply-side variability in the electricity market.

- **Transmission**: Transmission concerns the transport of large quantities of electricity over long distances through the high-voltage national grid. Due to the extremely high investment costs and the impracticality of building parallel grid infrastructures, transmission is naturally organized as a monopoly. In Italy, this activity is managed by Terna S.p.A., whose role is to ensure the security, balance, and reliability of the national electricity system.

- **Distribution**: Distribution is the local delivery of electricity through medium- and low-voltage networks. As with transmission, distribution is also operated as a natural monopoly, generally under state concessions and within defined geographical areas. Distribution system operators are responsible for the management, maintenance, and development of local networks, ensuring that electricity reaches final withdrawal points in a reliable and continuous way.

- **Sale**: The sale of electricity represents the final stage of the supply chain and concerns the commercialization of energy to households, firms, and public institutions. Unlike network activities, electricity retailing operates in a competitive market. For suppliers, accurate demand forecasting is therefore not only a technical requirement, but also a strategic tool.

Forecasting errors may lead to costly imbalances between the energy purchased on the wholesale market and the actual consumption of their customers.

This segmented structure provides the institutional and operational background for the data analyzed in this thesis. While transmission and distribution define the physical infrastructure and the geographical constraints of the power system, generation and sales introduce the main sources of economic and temporal variability. Together, these components shape electricity prices, demand fluctuations, and the need for balancing actions over time.

1.2 Power Reserve System and Frequency Regulation

The operation of an electricity system requires a continuous balance between generation and demand. Since electricity is difficult to store at large scale and consumption changes constantly throughout the day, even small mismatches between production and load may cause deviations in the system frequency. For this reason, the grid must be continuously monitored and controlled in real time in order to keep the frequency close to its nominal value.

To preserve system stability, the Grid Code defines different types of power reserves. These reserves consist of generation capacity, or controllable demand resources, that can be activated at different time scales to compensate for imbalances and restore secure operating conditions.

From an operational perspective, frequency regulation is organized through three main levels of intervention: primary, secondary, and tertiary reserves. Each of them acts with a different response time, degree of automation, and control objective. From a statistical and forecasting point of view, estimating the future need for reserve capacity is not only relevant for system security, but also for reducing balancing costs in the Dispatching Services Market.

1.2.1 Primary Reserve

The primary reserve represents the first automatic response of the power system to an imbalance between total generation and total demand. Its purpose is to contain frequency deviations immediately after a disturbance occurs, preventing the system from moving toward unsafe operating conditions.

This response is activated automatically by the speed governors of the generating units that provide the service. When a frequency deviation is detected, these devices adjust the power output of the participating units within a few seconds. To make this possible, generators must keep part of their available capacity as a reserve margin, so that their production can be increased or decreased when needed.

An important feature of the primary reserve is that it is not managed only at national level. Since the Italian electricity system is interconnected with the wider European power system, primary regulation is provided simultaneously by all enabled units operating in parallel across the interconnected area. This distributed response allows the system to react rapidly, regardless of where the imbalance physically occurs.

However, the primary reserve does not fully restore the frequency to its nominal value. Its main function is to stop the initial frequency deviation and stabilize the system after the disturbance. The task of bringing the frequency back to its reference value is then assigned to slower and more coordinated control mechanisms.

1.2.2 Secondary Reserve

The secondary reserve acts after the primary response has stabilized the initial disturbance. Its role is to correct the remaining imbalance between generation and demand within the National Power System and to restore the scheduled operating conditions.

Unlike the primary reserve, which is based on decentralized automatic responses from individual generators, the secondary reserve is coordinated through a centralized automatic control system managed by the transmission system operator. This controller sends regulation signals to the enabled units, adjusting their output in order to compensate for the imbalance.

The secondary reserve has two main objectives. First, it restores the planned electricity exchanges with neighboring systems, ensuring that cross-border flows return to their scheduled commercial values. Second, by correcting these deviations, it contributes to restoring the frequency of the interconnected European system to its nominal value of 50 Hz.

Due to their specific geographical and electrical characteristics, the major Italian islands may follow different regulation schemes. Sardinia, under normal operating conditions, and Sicily, when electrically disconnected from the mainland grid, manage secondary regulation locally rather than relying entirely on the national centralized control system.

1.2.3 Tertiary Reserve

The tertiary reserve provides an additional level of flexibility for the system operator. While primary and secondary reserves are mainly automatic control mechanisms, the tertiary reserve is activated through explicit dispatching instructions issued by the transmission system operator.

This reserve consists of upward or downward power margins made available by qualified production units, and in some cases by controllable demand resources, through the Dispatching Services Market. These margins may be procured during the scheduling phase or closer to real time, depending on system conditions and expected balancing needs.

The tertiary reserve can operate in two directions. Upward reserve is used when the system requires additional power, which can be obtained either by increasing generation or by reducing electricity withdrawals from the grid. Downward reserve is used when there is excess generation, requiring either a reduction in production or an increase in consumption. Each unit can provide only the type of reserve for which it has been technically qualified.

Overall, tertiary regulation represents the operator's main tool for managing more persistent or significant imbalances, restoring reserve margins after previous activations, and maintaining the system in a secure operating state over time.

1.3 Day-Ahead Electricity Price Formation

In other chapters of this thesis, electricity prices are analyzed across several European countries, inspired by [19]. In order to interpret these data correctly, it is important to clarify what the day-ahead electricity price represents and how it is computed in each market.

The day ahead market is the main short-term wholesale market for electricity. In this market, participants submit supply offers and demand bids for electricity delivery on the following day. For each market time interval, the market clearing mechanism matches supply and demand while accounting for the technical limits of the transmission network and the available cross-border interconnection capacity. The resulting price can therefore be interpreted as the marginal wholesale value of electricity in a given bidding zone and time interval.

At the European level, most day-ahead markets are integrated through the Single Day-Ahead Coupling mechanism. This framework uses a common price-coupling algorithm, known as EUPHEMIA, which jointly clears the coupled markets by maximizing social welfare, while respecting network constraints and cross-border capacities. As a consequence, prices in neighboring bidding zones may coincide when interconnection constraints are not binding, whereas they may diverge when congestion occurs.

A key point for cross country comparisons is that the day ahead price is not always defined in the same way at national level. Some countries are represented by a single bidding zone, while others are divided into multiple zones. In the latter case, either a specific bidding zone price or an aggregate national reference price must be selected.

1.3.1 Italy

In Italy, electricity is traded on the *Mercato del Giorno Prima* (MGP), the Italian day-ahead market operated by the Gestore dei Mercati Energetici (GME). The Italian market is organized into several geographical market zones, reflecting the presence of transmission constraints within the national grid. For each market time interval, the MGP produces a zonal marginal price, determined by the intersection between demand and supply curves within the market-clearing process.

The main national reference price for Italy is the *Prezzo Unico Nazionale* (PUN), currently published as the *PUN Index GME*. Unlike the prices of France or Germany, the PUN is not the clearing price of a single bidding zone. Instead, it is computed as a weighted average of the Italian zonal prices, where the weights are given by the accepted purchase quantities in the corresponding geographical market zones.

Therefore, the Italian series used in this thesis represents a national synthetic price, obtained from the zonal structure of the Italian market. This makes the PUN particularly useful for national-level analyses, but it should be interpreted differently from a pure single-zone market-clearing price.

1.3.2 France

France is generally represented in the day ahead market by a single national bidding zone. The French day ahead price is determined through the European coupled auction mechanism and corresponds to the market clearing price for electricity delivery in the French bidding zone.

This price is obtained from the matching of buy and sell orders submitted for the French delivery area, taking into account the interaction with neighboring countries through cross border transmission capacities.

1.3.3 Germany

Germany is represented in the day-ahead market by the German–Luxembourg bidding zone, commonly referred to as the DE-LU bidding zone. The German day-ahead price used in this thesis is therefore the market-clearing price associated with this bidding zone.

As in the French case, the price is obtained through the coupled European day-ahead auction. It reflects the intersection of supply and demand for the German-Luxembourg bidding zone, while also accounting for cross border exchanges with neighboring bidding zones. When transmission constraints are not binding, the German price may align with nearby markets; when congestion occurs, it may separate from them.

1.3.4 Denmark

Denmark differs from France and Germany because it is divided into two main bidding zones: DK1, corresponding to Western Denmark, and DK2, corresponding to Eastern Denmark. These two areas may display different day-ahead prices because they are connected to different neighboring systems and may be affected by different transmission constraints.

In this thesis, the Danish electricity price is represented by the price of the largest Danish bidding zone, namely DK1.

The DK1 price is determined in the day-ahead auction through the same European market-coupling logic used for the other considered countries. For each time interval, supply and demand bids are matched while considering the available interconnection capacity between DK1 and the neighboring bidding zones.

Chapter 2

Methodology

2.1 Box-Cox Transformation and Guerrero Algorithm

Many time series forecasting models rely, either explicitly or implicitly, on the assumption that the variability of the series remains reasonably stable over time. This property is commonly referred to as homoscedasticity. In practice, however, real world time series often violate this assumption. For example, the magnitude of seasonal fluctuations may increase as the level of the series grows, and periods with higher values may also be associated with larger absolute variations.

This issue is particularly relevant in electricity market data, where prices and demand may exhibit strong seasonality, sudden peaks, and changes in volatility over time. If this changing variance is not properly addressed, the forecasting model may assign excessive importance to high variance periods, produce unstable parameter estimates, or generate prediction intervals that are difficult to interpret.

For this reason, before fitting the forecasting models, a variance stabilizing transformation is applied to the original series. In this thesis, this role is performed by the Box-Cox transformation, whose transformation parameter is selected using the Guerrero method.

2.1.1 The Box-Cox Transformation

The Box-Cox transformation is a family of power transformations introduced by Box and Cox [2] to reduce heteroscedasticity and make a positive valued variable more suitable for statistical modelling. as explained in [2], given a strictly positive time series y_t, the transformed series w_t is defined as

$$w_t = \begin{cases} \log(y_t), & \text{if } \lambda = 0, \\ \dfrac{y_t^\lambda - 1}{\lambda}, & \text{if } \lambda \neq 0, \end{cases}$$

where λ is the transformation parameter.

Different values of λ correspond to different degrees of transformation. When $\lambda = 1$, the transformation is essentially linear and leaves the structure of the data almost unchanged, apart from a shift. When $\lambda = 0.5$, the transformation is equivalent to a square-root transformation. When $\lambda = 0$, the Box-Cox transformation becomes the natural logarithm, which is often appropriate when the variability of the series increases approximately in proportion to its level.

The main purpose of the transformation is not to change the underlying temporal structure of the data, but rather to make its variability more regular over time. This can improve the behaviour of forecasting models, especially when the original series displays level dependent volatility.

2.1.2 Selection of the Transformation Parameter: The Guerrero Algorithm

A central practical aspect of the Box-Cox transformation is the choice of the parameter λ. Selecting this value manually would introduce an arbitrary modelling decision and could lead to different results depending on the scale and variability of the data. To avoid this, the value of λ is selected through the Guerrero algorithm, explained in [8].

The Guerrero method provides a data-driven criterion for choosing the transformation parameter. Instead of focusing only on the global distribution of the series, it evaluates how the variability behaves across different portions of the data. This is particularly useful for seasonal time series, where the relationship between the level of the series and its variance may change across time.

Conceptually, the procedure can be described as follows. The original time series is divided into several non overlapping sub series, typically according to the seasonal structure of the data. For each sub series, the local mean and standard deviation are computed. The algorithm then considers different candidate values of λ and evaluates how stable the transformed variability is across the sub series.

More specifically, the selected value of λ is the one that minimizes the coefficient of variation of the transformed sub-series standard deviations, after accounting for the corresponding local means. In this way, the algorithm identifies the transformation that makes the variance as homogeneous as possible across different periods of the sample.

In the context of this thesis, the Guerrero algorithm is therefore used as an automatic and reproducible procedure for selecting the Box-Cox parameter. This allows the transformed time series to be better aligned with the assumptions of the forecasting models, while preserving the main temporal patterns present in the original data.

2.2 Model Evaluation: K-Fold and Time Series Cross Validation

A reliable assessment of forecasting performance requires evaluating the model on observations that were not used during the estimation phase. This principle is essential because a model may describe the training data very well without necessarily being able to generalize to future observations.

In traditional machine learning, a common approach for this purpose is K-Fold Cross Validation. This method provides a more stable estimate of predictive performance than a single division of the dataset into training and test samples, since the evaluation is repeated over multiple partitions of the data.

2.2.1 Standard K-Fold Cross Validation

In standard K-Fold Cross Validation, the dataset is divided into K subsets of approximately equal size, called folds. The model is trained using all folds except one, while the remaining fold is used as the validation set. This procedure is repeated K times, so that each fold is used once for validation.

At the end of the procedure, the errors obtained across the K validation steps are averaged to produce a single performance measure. This approach reduces the dependence of the evaluation on a specific train and test split and provides a more robust estimate of the model's predictive ability.

2.2.2 Data Leakage in Time Series

Although standard K-Fold Cross Validation is widely used in many machine learning applications, it is not directly suitable for time series forecasting. The reason is that time series observations are ordered chronologically, and this order contains essential information about the evolution of the process.

If the data are randomly split, observations from the future may be included in the training sample used to predict observations from the past. For example, electricity prices observed in December could be used to estimate a model that is then evaluated on prices from October of the same year. Such a procedure would not reflect a realistic forecasting setting, since future information is not available at the time the prediction is made.

This problem is known as data leakage. It leads to an overly optimistic evaluation of the model, because the validation procedure allows information from future observations to influence the estimation process. As a result, the model may appear more accurate than it would be in a genuine forecasting application.

2.2.3 Time Series Cross Validation

To avoid data leakage, the evaluation procedure must respect the chronological structure of the data. For this reason, this thesis adopts Time Series Cross Validation, a validation strategy specifically designed for forecasting problems.

The main idea is to reproduce the way in which forecasts would be generated in practice. The model is first estimated using an initial portion of the historical data and is then evaluated on the immediately following observations. After this first validation step, the training sample is expanded to include the newly observed data, and the model is estimated again. The procedure is repeated sequentially over the entire sample.

The process can be summarized as follows:

1. **Initial estimation**: the model is trained on an initial chronological portion of the available data.

2. **First validation step**: the fitted model is used to forecast the next block of observations, which has not been used during training.

3. **Expansion of the training window**: the training sample is enlarged by including the observations from the previous validation block.

4. **Sequential evaluation**: the model is estimated again on the expanded sample and used to forecast the next chronological block.

5. **Repetition of the procedure**: the same process continues until all validation blocks have been evaluated.

This expanding window structure ensures that, at each validation step, the model is trained only on information that would have been available at that point in time. Therefore, the evaluation more closely reflects the real forecasting environment.

2.2.4 Robust Evaluation

By averaging the forecast errors across the different temporal validation blocks, Time Series Cross Validation provides a robust measure of predictive accuracy while preserving the chronological order of the observations. This makes it particularly appropriate for electricity market data, where prices and demand are strongly dependent on past dynamics, seasonal patterns, and evolving market conditions.

In this thesis, the use of Time Series Cross Validation allows the forecasting models to be compared under a realistic evaluation framework. The resulting performance measures are therefore more informative than those obtained from a random validation scheme, since they are based on genuine out of sample forecasts generated in chronological order.

2.3 STL Decomposition

STL, acronym for *Seasonal and Trend decomposition using Loess*, was introduced by [4] as a flexible procedure for decomposing a time series into interpretable structural components. The method is particularly useful when the data exhibit complex seasonal patterns and gradual changes in their long term behaviour.

Unlike more restrictive decomposition techniques, STL does not impose a rigid parametric structure on the series. Instead, it relies on locally weighted regression to estimate the components of the time series in an iterative way. This makes the method able to adapt to local variations in the data and to capture changes in both trend and seasonal behaviour over time.

2.3.1 Decomposition Structure

In its standard additive formulation, STL represents each observation of the time series as the sum of three components:

$$y_t = T_t + S_t + R_t,$$

where y_t denotes the observed value at time t. The term T_t represents the trend component, which describes the long term evolution of the series and the broader movements of the market. The term S_t denotes the seasonal component, which captures systematic and recurrent patterns. Finally, R_t is the remainder component, which contains the part of the series that is not explained by either the trend or the seasonal structure.

This decomposition is especially useful in electricity market applications, where the observed series often combines persistent changes in the average level, regular seasonal effects, and irregular fluctuations caused by market shocks, weather conditions, or unexpected changes in demand and supply.

2.3.2 Loess Smoothing

The central smoothing tool used by STL is Loess, or locally estimated scatterplot smoothing. Loess is a flexible nonparametric regression method that estimates the value of a smooth curve at each point by fitting a local polynomial regression in a neighbourhood of that point.

Rather than estimating a single global function for the entire sample, Loess gives more importance to observations that are closer to the point being estimated. The weights assigned to neighbouring observations are commonly defined through the tri cubic kernel:

$$w(u) = \left(1 - |u|^3\right)^3 \mathbf{1}(|u| < 1),$$

where u is the normalized distance from the target point. According to this weighting scheme, observations close to the target point receive larger weights, while observations farther away have progressively smaller influence. As a result, the estimated curve can follow local changes in the data without requiring a fixed global functional form.

This local nature of Loess is one of the main reasons why STL is well suited for time series with evolving patterns. In the context of electricity data, it allows the decomposition to account for changes in seasonal intensity, market dynamics, and gradual structural movements.

2.3.3 Iterative Estimation Procedure

STL estimates the trend and seasonal components through an iterative procedure. The algorithm alternates between the estimation of the seasonal component and the estimation of the trend component, gradually refining both until a stable decomposition is obtained.

At a conceptual level, the procedure can be described as follows:

1. **Detrending**: an initial estimate of the trend is removed from the original series, so that the seasonal structure can be estimated more clearly.
2. **Seasonal smoothing**: Loess smoothing is applied to seasonal subseries. For example, in a series with daily or weekly seasonality, observations corresponding to the same position within the seasonal cycle are considered together in order to estimate the seasonal pattern.
3. **Seasonal adjustment**: the estimated seasonal component is removed from the original series, producing a seasonally adjusted series.
4. **Trend smoothing**: Loess smoothing is then applied to the seasonally adjusted series in order to update the estimate of the trend component.
5. **Iteration**: the previous steps are repeated until the estimates of the components become sufficiently stable.

In addition to this inner iterative procedure, STL can include an outer robustness loop. This step reduces the influence of observations that are associated with unusually large residuals. Let R_t denote the remainder obtained after removing the estimated trend and seasonal components. Robustness weights can be assigned according to

$$\rho_t = B\left(\frac{|R_t|}{h \cdot \text{median}(|R_t|)}\right),$$

where $B(\cdot)$ is the bisquare weighting function and h is a scaling constant. Observations with large residuals receive smaller weights in the following iterations. This makes the decomposition less sensitive to isolated outliers, which are common in electricity prices due to sudden spikes or abnormal market conditions.

2.3.4 Multiple Seasonality

Electricity time series often display more than one seasonal pattern. For example, electricity demand and prices may be affected by daily, weekly, and yearly cycles at the same time. A single seasonal component may therefore be insufficient to describe the full structure of the data.

When multiple seasonal components are present, the additive decomposition can be extended as

$$y_t = T_t + \sum_{k=1}^{K} S_t^{(k)} + R_t,$$

where $S_t^{(k)}$ represents the seasonal component associated with the corresponding period. In this framework, each seasonal component captures a different recurrent pattern. The components can be estimated sequentially, usually starting from the highest frequency seasonality and then moving toward slower seasonal cycles. After each seasonal effect is estimated, it is removed from the series, allowing the following step to focus on the remaining structure.

This extension is particularly relevant for electricity market data. Prices and consumption are influenced by human activity, working schedules, weather conditions, and institutional market routines. As a consequence, multiple seasonal effects may overlap and interact. Decomposing these effects separately helps to isolate the main regular patterns before applying forecasting models.

2.3.5 STL as a Preprocessing Step for Forecasting

In this thesis, STL is not used only as a descriptive tool, but also as a preprocessing step for forecasting. Many standard forecasting models, such as ETS and ARIMA, are designed to handle relatively simple seasonal structures. When the original time series contains complex or multiple seasonal patterns, directly applying these models may be inadequate.

By combining STL with a forecasting model, the seasonal structure can be separated before model estimation. The predictive model is then applied to the seasonally adjusted series, which is generally easier to model because the main recurrent patterns have already been removed.

Using the `decomposition_model()` framework from the `fpp3` package, the forecasting procedure can be summarized in three main steps:

1. **Decomposition of the original series**: the series y_t is decomposed into its trend component, its seasonal components, and its remainder component.

2. **Forecasting of the adjusted series**: the seasonal components are removed, and a forecasting model is fitted to the seasonally adjusted series

$$y_t^{\text{adj}} = y_t - \sum_{k=1}^{K} S_t^{(k)}.$$

   In this step, the model focuses mainly on the trend dynamics and on the remaining stochastic variation.

3. **Reconstruction of the final forecast**: the forecast of the seasonally adjusted series is combined with the extrapolated seasonal components to obtain the final forecast

$$\hat{y}_{t+h|t} = \hat{y}_{t+h|t}^{\text{adj}} + \sum_{k=1}^{K} S_{t+h}^{(k)}.$$

This approach allows the forecasting model to operate on a simplified representation of the data while still preserving the seasonal structure of the original series in the final predictions. For this reason, STL represents an important component of the modelling strategy adopted in this thesis.

2.4 ETS Models: Error, Trend, and Seasonality

Decomposition methods such as STL are useful for separating the historical components of a time series into trend, seasonal patterns, and residual variation. However, in order to generate forecasts, a formal statistical model is required. In this context, ETS models represent an alternative to ARIMA models. While ARIMA models describe the dependence structure of a series through autoregressive and moving average components, ETS models [12] are based on exponential smoothing and focus directly on the evolution of error, trend, and seasonal components.

An ETS model is identified by a triplet, where each component specifies how the model represents the error term, the trend component, and the seasonal component. This notation makes it possible to describe a wide range of models within a unified framework.

The first component, E, describes the structure of the error term. The error can be additive, denoted by A, or multiplicative, denoted by M. The second component, T, specifies the form of the trend. The model may have no trend, denoted by N, an additive trend, denoted by A, an additive damped trend, denoted by A_d, or a multiplicative trend, denoted by M. The third component, S, describes the seasonal structure. Seasonality may be absent, denoted by N, additive, denoted by A, or multiplicative, denoted by M.

This structure is particularly useful in empirical applications because it allows the forecasting model to adapt to different types of time series behaviour. In electricity market data, for example, the series may present a changing level, persistent movements, recurring seasonal effects, and volatility patterns that vary with the magnitude of the series.

2.4.1 General Formulation in the Additive Case

A useful starting point for understanding the ETS framework is the model ETS(A, A, A), which assumes additive errors, additive trend, and additive seasonality. In this case, the model can be written through the following state space equations:

$$y_t = \ell_{t-1} + b_{t-1} + s_{t-m} + \varepsilon_t, \tag{2.1}$$

$$\ell_t = \ell_{t-1} + b_{t-1} + \alpha\varepsilon_t, \tag{2.2}$$

$$b_t = b_{t-1} + \beta\varepsilon_t, \tag{2.3}$$

$$s_t = s_{t-m} + \gamma\varepsilon_t. \tag{2.4}$$

In these equations, y_t denotes the observed value of the series at time t. The term ℓ_t represents the level of the series, which can be interpreted as its local baseline. The term b_t denotes the trend component, which describes the local direction of movement. The term s_t represents the seasonal component, where m is the length of the seasonal period. Finally, ε_t is the innovation term, usually assumed to be normally distributed with mean zero and variance σ^2 in the additive Gaussian case.

The parameters α, β, and γ are smoothing parameters and take values in the interval $[0, 1]$. They determine how strongly the model reacts to new information. A larger value gives more weight to recent observations and makes the corresponding state component more responsive to new shocks. A smaller value produces a smoother evolution and gives more importance to the historical structure of the series.

This formulation shows the recursive nature of ETS models. At each time point, the forecast error is used to update the level, trend, and seasonal states. The model therefore adapts over time as new observations become available.

2.4.2 Multiplicative Components

In some time series, the amplitude of the fluctuations increases with the level of the series. In such cases, an additive formulation may be too restrictive. Multiplicative ETS models address this issue by allowing errors or seasonal effects to act proportionally to the level of the series.

An example is the model ETS(M, A, M), which has multiplicative errors, additive trend, and multiplicative seasonality:

$$y_t = (\ell_{t-1} + b_{t-1})s_{t-m}(1 + \varepsilon_t), \tag{2.5}$$

$$\ell_t = (\ell_{t-1} + b_{t-1})(1 + \alpha\varepsilon_t), \tag{2.6}$$

$$b_t = b_{t-1} + \beta(\ell_{t-1} + b_{t-1})\varepsilon_t, \tag{2.7}$$

$$s_t = s_{t-m}(1 + \gamma\varepsilon_t). \tag{2.8}$$

In this specification, the seasonal effect and the innovation affect the series through scaling rather than through simple addition. This allows the model to represent situations in which the magnitude of seasonal movements or random shocks changes according to the current level of the series.

For electricity prices and demand, this feature can be relevant because periods with higher average values may also exhibit larger absolute fluctuations. A multiplicative structure can therefore provide a more appropriate representation when variability is level dependent.

2.4.3 Damped Trend

A standard additive trend assumes that the current direction of movement will continue indefinitely into the future. This assumption may be unrealistic, especially over long forecasting horizons. To obtain more conservative forecasts, the ETS framework allows the use of a damped trend.

In a damped trend model, the trend component is multiplied by a damping parameter ϕ, with $0 < \phi < 1$. The level and trend updating equations can be written as

$$\ell_t = \ell_{t-1} + \phi b_{t-1} + \alpha\varepsilon_t, \tag{2.9}$$

$$b_t = \phi b_{t-1} + \beta\varepsilon_t. \tag{2.10}$$

The parameter ϕ controls the persistence of the trend. When ϕ is close to one, the model behaves similarly to a standard linear trend model. When ϕ is smaller, the influence of the trend decreases more rapidly as the forecast horizon increases. As a result, the forecasts tend to approach a more stable long run trajectory.

This feature is useful when the historical data show a clear trend, but it is not reasonable to assume that such a trend will continue with the same intensity indefinitely.

2.4.4 Parameter Estimation and Model Selection

Before forecasts can be produced, the model parameters and the initial states must be estimated. The parameters include the smoothing coefficients α, β, and γ, and, when present, the damping parameter ϕ. The initial states include the initial level, the initial trend, and the initial seasonal factors.

In the ETS framework, these quantities are commonly estimated by maximum likelihood. Under the assumption of additive normally distributed errors, with $\varepsilon_t \sim \text{NID}(0, \sigma^2)$, the log likelihood can be expressed as

$$\mathcal{L} = -\frac{n}{2}\log(2\pi\hat{\sigma}^2) - \frac{1}{2\hat{\sigma}^2}\sum_{t=1}^{n}\varepsilon_t^2, \tag{2.11}$$

where

$$\hat{\sigma}^2 = \frac{1}{n}\sum_{t=1}^{n}\varepsilon_t^2$$

is the estimated residual variance.

Since the ETS framework contains many possible model specifications, a criterion is needed to select the most appropriate one. In this thesis, model selection is based on the corrected Akaike Information Criterion, denoted by AIC_c:

$$\text{AIC}_c = -2\mathcal{L} + 2k + \frac{2k(k+1)}{n-k-1}, \tag{2.12}$$

where k is the number of estimated parameters and n is the sample size.

The AIC_c balances goodness of fit and model complexity. Models with a higher likelihood are preferred, but specifications with too many parameters are penalized. This reduces the risk of selecting a model that fits the historical sample very closely while performing poorly on future observations.

2.4.5 Point Forecasts

Once the model has been estimated and the best specification has been selected, forecasts can be generated recursively from the estimated states. For the additive model ETS(A, A, A), the h step ahead point forecast is given by

$$\hat{y}_{t+h|t} = \ell_t + hb_t + s_{t+h-m(k+1)}, \tag{2.13}$$

where

$$k = \left\lfloor \frac{h-1}{m} \right\rfloor.$$

The seasonal index ensures that the appropriate seasonal state is used for the future period being forecast. In this way, the model combines the current estimate of the level, the projected trend, and the seasonal effect corresponding to the forecast horizon.

2.4.6 Prediction Intervals

A relevant advantage of the ETS framework is that it is formulated as a statistical state space model. Therefore, it can provide not only point forecasts, but also measures of forecast uncertainty. This is important because electricity market variables are subject to irregular shocks and sudden changes, making uncertainty quantification essential for a meaningful evaluation of the forecasts.

For linear models with additive errors, future observations are usually assumed to be normally distributed around the point forecast. The prediction interval for an h step ahead forecast can be written as

$$\hat{y}_{t+h|t} \pm c\sigma_h,$$

where $\hat{y}_{t+h|t}$ is the point forecast, c is the critical value from the standard normal distribution, and σ_h is the standard deviation of the forecast distribution at horizon h. For example, a standard 95% prediction interval uses approximately $c = 1.96$.

The term σ_h depends on the model structure and on the estimated smoothing parameters. In general, forecast uncertainty increases with the forecasting horizon, because future shocks affect not only the observation itself, but also the future evolution of the underlying states.

For example, in the simple exponential smoothing model ETS(A, N, N), the forecast variance at horizon h is

$$\sigma_h^2 = \sigma^2 \left[1 + (h-1)\alpha^2\right].$$

In more complex ETS models, such as those including trend and seasonality, the expression for the forecast variance also accounts for the interaction between the updating equations of the level, trend, and seasonal states. As a consequence, prediction intervals usually become wider as the forecast horizon increases.

2.4.7 Simulation and Bootstrapping

Prediction intervals are analytically computable for linear models with additive errors. However, when multiplicative errors or multiplicative seasonal components are included, the distribution of future observations may become asymmetric, and closed form formulas can be difficult to derive.

In these cases, prediction intervals can be obtained through simulation or bootstrapping. The model generates a large number of possible future paths by repeatedly drawing innovation terms and propagating them through the estimated state equations. The empirical distribution of these simulated future values is then used to compute the desired quantiles.

This approach is particularly useful when the assumptions required for analytical intervals are too restrictive. It allows the uncertainty of the forecasts to reflect the nonlinear structure of the model and the empirical behaviour of the residuals. In the context of this thesis, prediction intervals provide an additional tool for evaluating the reliability of the forecasts, complementing the analysis based on point forecast accuracy.

2.5 ARIMA Models

Together with exponential smoothing methods, ARIMA models, acronym for AutoRegressive Integrated Moving Average, represent one of the main classical approaches to time series forecasting. The framework was formalized in the work of Box and Jenkins [?] and is based on the idea that the future evolution of a time series can be described through its own past values and past forecast errors.

Differently from decomposition methods, which explicitly separate a series into trend, seasonal, and residual components, ARIMA models focus on the autocorrelation structure of the data. The model combines three elements. The autoregressive component describes the dependence of the current value on previous observations. The integrated component accounts for nonstationarity through differencing. The moving average component describes the dependence of the current value on previous innovation terms.

This framework is particularly useful when the series presents temporal dependence that can be captured through a relatively parsimonious linear model. In the context of electricity market data, ARIMA models provide a natural benchmark for forecasting prices and demand, since both variables often exhibit strong persistence and recurrent temporal patterns.

2.5.1 Stationarity and Differencing

A central assumption behind ARMA and ARIMA modelling is stationarity. A weakly stationary time series has statistical properties that remain stable over time. In particular, its mean and variance are constant, and its autocovariance depends only on the lag between two observations, not on the specific time at which they are observed. Formally, a time series y_t is weakly stationary if

$$\mathbb{E}[y_t] = \mu \quad \forall t, \tag{2.14}$$

$$\mathrm{Var}(y_t) = \sigma^2 \quad \forall t, \tag{2.15}$$

$$\mathrm{Cov}(y_t, y_{t-k}) = \gamma_k \quad \forall t. \tag{2.16}$$

In empirical applications, many time series are not stationary in their original form. Electricity prices and electricity demand may exhibit trends, changing levels, and persistent seasonal effects. In such cases, differencing is used to remove nonstationary behaviour in the mean.

Using the backshift operator B, defined by $B^k y_t = y_{t-k}$, the first difference of the series is

$$\Delta y_t = (1 - B)y_t = y_t - y_{t-1}.$$

More generally, the d order differenced series is defined as

$$\Delta^d y_t = (1 - B)^d y_t.$$

The purpose of differencing is to transform the original series into a more stable process that can be modelled using autoregressive and moving average components. In most practical applications, one or two differences are sufficient to remove the main nonstationary structure in the mean.

2.5.2 ARMA Models

Before introducing differencing, it is useful to consider the ARMA model, which is designed for stationary time series. An ARMA(p, q) model combines an autoregressive component of order p and a moving average component of order q:

$$y_t = c + \sum_{i=1}^{p} \phi_i y_{t-i} + \varepsilon_t + \sum_{j=1}^{q} \theta_j \varepsilon_{t-j}. \tag{2.17}$$

In backshift notation, the same model can be written as

$$\phi(B) y_t = c + \theta(B) \varepsilon_t, \tag{2.18}$$

where

$$\phi(B) = 1 - \phi_1 B - \phi_2 B^2 - \cdots - \phi_p B^p$$

is the autoregressive polynomial and

$$\theta(B) = 1 + \theta_1 B + \theta_2 B^2 + \cdots + \theta_q B^q$$

is the moving average polynomial.

The autoregressive part models the dependence of the current observation on previous values of the series. The moving average part models the effect of previous shocks or forecast errors. The innovation term ε_t is typically assumed to be a white noise process with mean zero and variance σ^2.

2.5.3 ARIMA Models

An ARIMA(p, d, q) model extends the ARMA framework to nonstationary time series by applying the ARMA structure to a differenced version of the original data. If

$$w_t = \Delta^d y_t,$$

then w_t is modelled as an ARMA(p, q) process. Equivalently, the ARIMA model can be written as

$$\phi(B)(1 - B)^d y_t = c + \theta(B) \varepsilon_t. \tag{2.19}$$

The three integers (p, d, q) define the structure of the model. The parameter p is the number of autoregressive lags, d is the number of differencing operations, and q is the number of moving average terms. Once the appropriate differencing has been applied, the remaining dynamics are captured through the autoregressive and moving average components.

This formulation allows ARIMA models to describe a wide range of temporal behaviours while maintaining a relatively simple and interpretable structure.

2.5.4 Seasonal ARIMA Models

When the time series contains a recurring seasonal pattern, the ARIMA framework can be extended to include seasonal autoregressive, differencing, and moving average terms. The resulting model is denoted by

$$\text{ARIMA}(p,d,q)(P,D,Q)_m,$$

where m is the seasonal period, and (P,D,Q) describe the seasonal autoregressive order, the seasonal differencing order, and the seasonal moving average order. The model can be written as

$$\phi(B)\Phi(B^m)(1-B)^d(1-B^m)^D y_t = c + \theta(B)\Theta(B^m)\varepsilon_t. \tag{2.20}$$

Here, $\Phi(B^m)$ and $\Theta(B^m)$ denote the seasonal autoregressive and seasonal moving average polynomials. The term $(1-B^m)^D$ represents seasonal differencing and is used to remove seasonal nonstationarity.

A limitation of standard seasonal ARIMA models is that they are naturally designed to handle a single seasonal period. This can be restrictive for electricity market data, where several seasonal patterns may coexist, such as daily, weekly, and annual cycles.

One way to address this limitation is to combine ARIMA with STL decomposition. In this approach, STL is first used to remove the seasonal components from the original series. Then, a nonseasonal ARIMA model is fitted to the seasonally adjusted series:

$$\phi(B)(1-B)^d y_t^{\text{adj}} = c + \theta(B)\varepsilon_t. \tag{2.21}$$

This strategy allows the ARIMA model to focus on the remaining nonseasonal temporal dependence, while the seasonal structure is handled separately through the decomposition step.

2.5.5 Parameter Estimation and Order Selection

The construction of an ARIMA model requires the selection of the orders (p,d,q) and, when seasonality is included, the seasonal orders (P,D,Q). In this thesis, the selection is performed using the Hyndman and Khandakar algorithm, implemented in [11].

The procedure first determines the appropriate level of differencing through statistical tests for unit roots and seasonal unit roots. Once the differencing orders have been selected, the algorithm searches over alternative autoregressive and moving average specifications and compares them using the corrected Akaike Information Criterion, denoted by AIC_c:

$$\text{AIC}_c = -2\mathcal{L} + 2k + \frac{2k(k+1)}{n-k-1}, \tag{2.22}$$

where $\mathcal{L}$ is the maximized log likelihood, k is the number of estimated parameters, and n is the sample size.

The use of AIC_c allows the model selection procedure to balance goodness of fit and model complexity. A model with a better fit is preferred, but excessive complexity is penalized in order to reduce the risk of overfitting.

After the orders have been selected, the model parameters are estimated by maximum likelihood. Under the assumption of Gaussian innovations, the log likelihood can be written as

$$\mathcal{L} = -\frac{n}{2}\log(2\pi\sigma^2) - \frac{1}{2\sigma^2}\sum_{t=1}^{n}\varepsilon_t^2. \tag{2.23}$$

The estimated parameters include the autoregressive coefficients, the moving average coefficients, the innovation variance, and, when present, the constant or drift term.

2.5.6 Forecasting

Once the ARIMA model has been estimated, forecasts are obtained by computing the conditional expectation of future observations given the information available up to time t:

$$\hat{y}_{t+h|t} = \mathbb{E}\left[y_{t+h} \mid y_t, y_{t-1}, \dots\right]. \tag{2.24}$$

The forecasts are generated recursively. Future values of the series that are not yet observed are replaced by their forecasts, while future innovation terms have conditional expectation equal to zero. This recursive structure makes it possible to produce forecasts for any horizon h.

Prediction intervals can also be derived from the forecast error variance. As the forecast horizon increases, the uncertainty generally becomes larger, because future shocks accumulate through the autoregressive and moving average structure of the model.

2.5.7 ARIMA with Exogenous Regressors

The ARIMA framework can also be extended to include external explanatory variables. This leads to a dynamic regression model with ARIMA errors. Let

$$\mathbf{x}_t = (x_{1t}, \dots, x_{rt})^\top$$

be a vector of external predictors, such as temperature, calendar effects, daylight hours, fuel prices, or other market variables. The model is written as

$$y_t = \mathbf{x}_t^\top \boldsymbol{\beta} + \eta_t, \tag{2.25}$$

where $\boldsymbol{\beta}$ is the vector of regression coefficients and η_t is an error term with ARIMA structure:

$$\phi(B)(1-B)^d \eta_t = \theta(B)\varepsilon_t, \quad \varepsilon_t \sim \mathrm{NID}(0, \sigma^2). \tag{2.26}$$

This formulation differs from standard linear regression because the errors are not assumed to be independent. Instead, their temporal dependence is explicitly modelled through the ARIMA component. As a result, the model can account both for the effect of external variables and for the remaining autocorrelation in the series.

In this thesis, dynamic regression with ARIMA errors provides a flexible framework for incorporating explanatory variables into the forecasting process, while preserving the time series structure of the residual component.

2.5.8 ARIMA with Fourier Terms

Multiple seasonal patterns can also be represented within an ARIMA based framework by introducing Fourier terms as external regressors. This approach is useful when the series contains several seasonal cycles and a direct seasonal ARIMA specification would be too restrictive.

For a seasonal period m, the Fourier representation uses pairs of sine and cosine functions:

$$\sin\left(\frac{2\pi kt}{m}\right), \quad \cos\left(\frac{2\pi kt}{m}\right), \quad k = 1, 2, \dots, K. \tag{2.27}$$

If the time series contains J seasonal periods, denoted by $m_1, \dots, m_J$, the model can be written as

$$y_t = \sum_{j=1}^{J} \sum_{k=1}^{K_j} \left[\alpha_{jk} \sin\left(\frac{2\pi kt}{m_j}\right) + \beta_{jk} \cos\left(\frac{2\pi kt}{m_j}\right) \right] + \eta_t. \tag{2.28}$$

The residual term η_t is then modelled as an ARIMA process:

$$\phi(B)(1-B)^d \eta_t = \theta(B)\varepsilon_t, \quad \varepsilon_t \sim \mathrm{NID}(0, \sigma^2). \tag{2.29}$$

The number of Fourier pairs K_j controls the flexibility of the seasonal representation. A larger value allows the model to capture more complex seasonal shapes, but also increases the number of parameters. Therefore, the choice of K_j must balance flexibility and parsimony.

This formulation is particularly useful for electricity data, where seasonal patterns may be strong but not necessarily simple. Fourier terms provide a smooth representation of recurrent cycles, while the ARIMA error structure captures the remaining autocorrelation.

2.6 GARCH Models

In addition to models designed to forecast the conditional mean of a time series, it is often important to model the evolution of its conditional variance. This is particularly relevant in electricity markets, where prices may exhibit periods of relative stability followed by intervals of intense fluctuations. Similar combinations of ARIMA-type mean models and GARCH volatility specifications have been used in short-term electricity price forecasting applications [20]. The GARCH framework, acronym for *Generalized Autoregressive Conditional Heteroskedasticity*, provides a statistical approach for modelling this time varying volatility.

The GARCH model was introduced by Bollerslev [1] as a generalization of the ARCH model proposed by Engle [7]. The main motivation behind these models is the empirical phenomenon known as volatility clustering. In many economic and financial time series, large shocks tend to be followed by large shocks, while small shocks tend to be followed by small shocks. This means that the variance of the series is not constant over time, but evolves according to its own dynamic structure.

Traditional models such as ARIMA usually assume that the innovation variance is constant. GARCH models relax this assumption by allowing the conditional variance to depend on past shocks and past values of the variance itself. Therefore, the model does not only describe the expected value of the series, but also the uncertainty surrounding that value.

2.6.1 Conditional Volatility

Let ε_t denote the innovation term obtained after modelling the conditional mean of the series. In a standard homoskedastic framework, this innovation is assumed to have constant variance σ^2. In contrast, GARCH models assume that the variance of ε_t changes over time and is conditional on the information available up to time $t-1$.

The innovation can be written as

$$\varepsilon_t = z_t \sigma_t, \tag{2.30}$$

where z_t is an independent and identically distributed random variable with zero mean and unit variance, and σ_t is the conditional standard deviation at time t. A common assumption is

$$z_t \sim \mathrm{NID}(0, 1),$$

although other distributions may be used when the empirical residuals display heavier tails.

The conditional variance is defined as

$$\sigma_t^2 = \mathrm{Var}\left(\varepsilon_t \mid \mathcal{F}_{t-1}\right),$$

where $\mathcal{F}_{t-1}$ denotes the information set available before observing the innovation at time t. This formulation makes the variance a predictable quantity based on past information.

2.6.2 ARCH and GARCH Specifications

The original ARCH model assumes that the conditional variance depends on past squared innovations. In an ARCH(q) model, the variance equation is

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2. \tag{2.31}$$

In this specification, large past shocks increase the current conditional variance, while small past shocks lead to lower predicted volatility. The coefficients α_i measure the effect of previous squared innovations on the current level of volatility.

The GARCH(p, q) model extends the ARCH structure by including lagged conditional variances. Its variance equation is given by

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2. \tag{2.32}$$

The parameter ω represents the baseline component of volatility and is usually required to be strictly positive. The coefficients α_i are the ARCH parameters and measure the short run impact of recent shocks. The coefficients β_j are the GARCH parameters and measure the persistence of past conditional variances.

This structure allows the model to capture persistent volatility dynamics with a relatively small number of parameters. For this reason, the GARCH$(1, 1)$ specification is often used as a benchmark in applications involving financial and energy market data.

2.6.3 Stationarity and Unconditional Variance

For a GARCH process to be stable, the conditional variance must not diverge over time. Under the standard nonnegativity restrictions on the parameters, a sufficient condition for covariance stationarity is

$$\sum_{i=1}^{q} \alpha_i + \sum_{j=1}^{p} \beta_j < 1. \tag{2.33}$$

When this condition is satisfied, the conditional variance may fluctuate over time, but it reverts toward a finite long run level. The unconditional variance is then given by

$$\sigma^2 = \frac{\omega}{1 - \sum_{i=1}^{q} \alpha_i - \sum_{j=1}^{p} \beta_j}. \tag{2.34}$$

The sum of the ARCH and GARCH coefficients determines the persistence of volatility. Values close to one indicate that volatility shocks decay slowly and may influence the series for a long period. If the sum is equal to one, the model corresponds to an integrated GARCH specification, in which shocks to volatility have a highly persistent effect.

2.6.4 Parameter Estimation

The parameters of a GARCH model are commonly estimated by maximum likelihood. Since the variance changes over time, the likelihood function must explicitly account for the conditional variance at each observation.

Assuming conditionally normally distributed innovations, the log likelihood is

$$\mathcal{L} = -\frac{1}{2} \sum_{t=1}^{n} \left[\log(2\pi) + \log(\sigma_t^2) + \frac{\varepsilon_t^2}{\sigma_t^2} \right]. \tag{2.35}$$

The estimation procedure selects the parameters that maximize this likelihood. Intuitively, the model is penalized when large squared residuals occur during periods in which the estimated conditional variance is low. Therefore, the fitted GARCH model attempts to assign higher volatility to periods characterized by larger shocks and lower volatility to more stable periods.

In practice, the GARCH variance equation is often combined with a separate mean equation. The mean equation may be a constant, an ARMA model, or another suitable specification. The GARCH component is then applied to the residuals of this mean equation in order to describe their conditional variance.

2.6.5 Forecasting Conditional Volatility

After the model has been estimated, it can be used to forecast future conditional variance. Consider the GARCH$(1,1)$ model:

$$\sigma_t^2 = \omega + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2. \tag{2.36}$$

The one step ahead forecast of the conditional variance, based on the information available at time t, is

$$\hat{\sigma}_{t+1|t}^2 = \omega + \alpha_1 \varepsilon_t^2 + \beta_1 \sigma_t^2. \tag{2.37}$$

For horizons greater than one, future squared innovations are not observed. Their expected value is replaced by the corresponding forecasted conditional variance. This leads to the recursive expression

$$\hat{\sigma}_{t+h|t}^2 = \omega + (\alpha_1 + \beta_1)\hat{\sigma}_{t+h-1|t}^2, \quad h > 1. \tag{2.38}$$

If the stationarity condition is satisfied, the volatility forecast converges toward the unconditional variance as the forecast horizon increases. This property reflects the mean reverting behaviour of stationary GARCH volatility.

In the context of electricity price modelling, GARCH models are useful because they provide a direct representation of volatility dynamics. While models such as ARIMA or ETS focus primarily on the expected value of the series, GARCH models describe how uncertainty changes over time. This makes them particularly relevant when the objective is not only to forecast prices, but also to evaluate periods of higher risk and instability in the market.

2.7 Prophet

Prophet is a forecasting model introduced by Taylor and Letham and developed for time series that display strong seasonal patterns, trend changes, missing observations, and calendar effects. Unlike autoregressive models, such as ARIMA, Prophet does not primarily describe the series through its dependence on past values and past errors. Instead, it adopts a decomposable regression structure, in which the observed series is represented as the combination of several interpretable components.

This approach is particularly useful in applications where the time series is influenced by recurrent calendar patterns and irregular events. In electricity market data, for example, demand and prices may be affected by daily and weekly cycles, annual seasonality, public holidays, and changes in the underlying market trend. Prophet provides a flexible framework for incorporating these effects explicitly.

2.7.1 General Formulation

In its standard additive formulation, Prophet represents the observed value y_t as

$$y_t = g(t) + s(t) + h(t) + \varepsilon_t, \tag{2.39}$$

where $g(t)$ is the trend component, $s(t)$ is the seasonal component, $h(t)$ is the holiday component, and ε_t is an error term. The error term is usually assumed to be normally distributed with zero mean and variance σ^2:

$$\varepsilon_t \sim \mathcal{N}(0, \sigma^2).$$

The trend component captures the long term, non periodic evolution of the series. The seasonal component represents systematic patterns that repeat over fixed periods, such as daily, weekly, or yearly cycles. The holiday component accounts for the effect of known calendar events that may cause deviations from the usual behaviour of the series.

Prophet also allows a multiplicative specification, which can be useful when the magnitude of seasonal or holiday effects changes proportionally with the level of the trend. In that case, the model can be written as

$$y_t = g(t)\,(1 + s(t) + h(t)) + \varepsilon_t. \tag{2.40}$$

The additive formulation is appropriate when seasonal and holiday effects have approximately constant magnitude, while the multiplicative formulation is more suitable when these effects increase or decrease with the level of the series.

2.7.2 Trend Component

The trend component $g(t)$ describes the underlying direction of the series over time. Prophet provides two main alternatives for modelling the trend: a piecewise linear trend and a saturating logistic trend.

The piecewise linear trend is generally appropriate when the series can grow or decline without a predefined upper bound. In this case, the model allows the slope of the trend to change at a set of potential changepoints. The trend is written as

$$g(t) = \left(k + \mathbf{a}(t)^\top \boldsymbol{\delta}\right) t + \left(m + \mathbf{a}(t)^\top \boldsymbol{\gamma}\right), \tag{2.41}$$

where k is the initial growth rate and m is an offset parameter. The vector $\mathbf{a}(t)$ indicates which changepoints have occurred before time t, while $\boldsymbol{\delta}$ contains the corresponding changes in the growth rate. The vector $\boldsymbol{\gamma}$ is introduced to ensure that the trend remains continuous at the changepoints.

To avoid excessive flexibility in the trend, Prophet uses a prior distribution on the rate adjustments. In particular, the changepoint effects are regularized through a Laplace prior:

$$\delta_j \sim \text{Laplace}(0, \tau), \tag{2.42}$$

whose density is given by

$$p(\delta_j \mid \tau) = \frac{1}{2\tau} \exp\left(-\frac{|\delta_j|}{\tau}\right). \tag{2.43}$$

The parameter τ controls the flexibility of the trend. Larger values allow the model to adapt more strongly to changes in the historical trajectory, while smaller values produce a smoother and more stable trend.

When the series is expected to approach a limiting value, Prophet can instead use a saturating growth model. This specification is based on a logistic curve with carrying capacity $C(t)$:

$$g(t) = \frac{C(t)}{1 + \exp\left(-\left(k + \mathbf{a}(t)^\top \boldsymbol{\delta}\right)\left(t - \left(m + \mathbf{a}(t)^\top \boldsymbol{\gamma}\right)\right)\right)}. \tag{2.44}$$

This formulation is useful when the process is subject to a natural or structural limit. In electricity applications, such limits may arise from physical capacity constraints, infrastructure restrictions, or saturation effects in demand.

2.7.3 Seasonal Component

The seasonal component $s(t)$ is modelled using Fourier series. This allows Prophet to represent periodic patterns through smooth combinations of sine and cosine functions. For a seasonal period P, the seasonal component can be written as

$$s(t) = \sum_{k=1}^{K} \left[a_k \cos\left(\frac{2\pi kt}{P}\right) + b_k \sin\left(\frac{2\pi kt}{P}\right)\right]. \tag{2.45}$$

The parameter K determines the number of Fourier terms included in the model. A small value of K produces a smoother seasonal pattern, while a larger value allows the model to capture more complex seasonal shapes. However, increasing K also increases the number of parameters and may lead to overfitting if the seasonal component becomes too flexible.

One advantage of this formulation is that several seasonal periods can be included simultaneously. For example, the total seasonal component may be written as

$$s(t) = s_{\text{daily}}(t) + s_{\text{weekly}}(t) + s_{\text{yearly}}(t). \tag{2.46}$$

This is particularly relevant for electricity market data, where multiple seasonal patterns may coexist. Prices and consumption can vary within the day, across days of the week, and across seasons of the year. The Fourier representation allows these effects to be modelled in a flexible and interpretable way.

2.7.4 Holiday Component

The holiday component $h(t)$ accounts for the effect of known calendar events. Public holidays may significantly modify electricity demand and, consequently, electricity prices. These effects are often not captured adequately by standard seasonal components, because holidays do not always occur at the same position within weekly or yearly cycles.

Prophet represents holidays through a matrix of indicator variables. Let $\mathbf{Z}(t)$ be a vector indicating whether time t belongs to a specific holiday or holiday window. The holiday component is then written as

$$h(t) = \mathbf{Z}(t)^\top \boldsymbol{\kappa}, \tag{2.47}$$

where $\boldsymbol{\kappa}$ is the vector of holiday effects.

This structure allows each holiday to have its own estimated impact on the forecast. In the context of this thesis, the holiday component is especially useful for electricity data because national holidays may alter consumption profiles, working schedules, and market activity. For the Italian market, the implementation relies on the holiday calendar generated through the `prophet` package, so that national public holidays are explicitly included in the model.

2.7.5 Parameter Estimation and Forecasting

Prophet estimates its parameters using a Bayesian framework. In practice, the model is commonly fitted through Maximum A Posteriori estimation. This means that the estimation procedure combines the likelihood of the observed data with prior assumptions on the model parameters. The priors play an important regularization role, especially for changepoints and seasonal effects, helping to avoid excessive sensitivity to short term fluctuations.

Once the model components have been estimated, the h step ahead point forecast is obtained by evaluating the trend, seasonal, and holiday components at the future time point $t + h$:

$$\hat{y}_{t+h|t} = g(t+h) + s(t+h) + h(t+h). \tag{2.48}$$

In the additive case, the forecast is therefore obtained by summing the extrapolated trend, the expected seasonal pattern, and the holiday effect if a relevant calendar event occurs at the forecast horizon.

Prophet also provides uncertainty intervals by accounting for observation noise and uncertainty in future trend changes. These intervals are useful because they allow the forecasts to be interpreted not only in terms of expected values, but also in terms of the uncertainty associated with future market behaviour.

Chapter 3

Dataset Description

The development and validation of the forecasting models presented in this thesis are based on three main datasets. Together, they provide a complementary representation of the electricity market from both a physical and an economic perspective. The first dataset describes the physical side of the Italian power system, through historical electricity load data. The second dataset captures the Italian wholesale electricity price, through the *Prezzo Unico Nazionale*. The third dataset extends the price analysis to other European markets, France, Germany, and Denmark.

The following sections describe the sources, temporal coverage, frequency, and main statistical characteristics of the datasets used in the empirical analysis.

3.1 Prezzo Unico Nazionale Dataset

The first dataset contains historical wholesale electricity prices for the Italian market. In particular, the analysis focuses on the *Prezzo Unico Nazionale*, commonly referred to as PUN, which represents the main national reference price for electricity purchased on the Italian day ahead market.

While the Italian electricity market is physically divided into several geographical bidding zones, the PUN provides an aggregate national price signal. It is therefore widely used as a reference for studying the economic behaviour of the Italian power market.

3.1.1 Source and Temporal Scope

The price data were acquired from the ENTSO E, European Network of Transmission System Operators for Electricity, Transparency Platform, https://transparency.entsoe.eu/. The dataset refers to the Italian day ahead market and covers the period from January 1, 2004, to the most recent available observations in 2026.

The long temporal coverage of this dataset is one of its main strengths. It allows the analysis to include different market regimes, structural changes, and periods of exceptional volatility. As a result, the dataset provides a broad historical perspective on the evolution of wholesale electricity prices in Italy.

3.1.2 Resolution and Main Characteristics

The PUN time series is recorded at an hourly frequency, corresponding to 24 observations per day. This resolution is consistent with the historical structure of the day ahead electricity market, where prices have traditionally been determined for each delivery hour.

The hourly frequency makes the series suitable for analysing both short term price movements and longer term market dynamics. Within each day, prices may vary significantly according to demand levels, generation availability, renewable production, and cross border exchanges.

The extensive historical depth of the dataset makes it possible to observe several important episodes in the Italian and European energy markets. For example, the series includes the period of the 2008 global financial crisis, which affected industrial activity and electricity demand. It also includes the 2022 energy crisis, characterized by exceptional price increases and high volatility, largely connected to tensions in natural gas markets and supply uncertainty across Europe.

3.2 European Electricity Price Dataset

The second dataset extends the analysis of electricity prices beyond the Italian market. It contains hourly electricity market data for France, Germany, and Denmark, and is used to compare the behaviour of wholesale electricity prices across different European power markets.

The inclusion of these countries allows the empirical analysis to move from a purely national perspective to a broader European comparison. This is relevant because European electricity markets are increasingly interconnected through cross-border transmission capacity and market coupling mechanisms. At the same time, each country retains specific structural characteristics, such as its generation mix, demand profile, bidding zone configuration, and exposure to renewable production. These differences may lead to distinct price dynamics, even within the common European market framework.

3.2.1 Source and Temporal Scope

The data were obtained from the Energy Charts platform, available at `energy-charts.info`. Energy Charts is a public data platform developed by the Fraunhofer Institute for Solar Energy Systems ISE, which provides data on electricity production and spot market prices for several European countries. The platform offers data on power generation, spot market prices, emissions, and related electricity system indicators.

The dataset includes hourly observations spanning from 2019 to 2024. The dependent variable is the day-ahead electricity price, expressed in EUR/MWh. The explanatory variables describe electricity demand, generation by source, cross-border electricity exchanges, and the cost of carbon emission allowances.

3.2.2 Country Representation

The French price series represents the day-ahead market price for the French bidding zone. Since France is generally treated as a single bidding zone in the day-ahead market, the series can be interpreted as the national wholesale price signal for electricity delivery in France.

The German price series refers to the German-Luxembourg bidding zone. This bidding zone represents the standard reference area for German day-ahead electricity prices. The resulting price reflects the interaction between supply and demand within the German-Luxembourg market area, together with the effect of cross-border exchanges and market coupling with neighbouring zones.

The Danish case requires a more specific definition because Denmark is divided into two main bidding zones, DK1 and DK2. DK1 corresponds to Western Denmark, while DK2 corresponds to Eastern Denmark. Since the two zones may display different prices, the Danish series used in this thesis is represented by DK1. This choice provides a clear and consistent reference for Denmark and avoids ambiguity in the construction of the price series.

3.2.3 Covariates Used in the Regression Analysis

The analysis uses only variables that describe structural and market-related drivers of electricity prices. Variables representing other electricity prices, such as neighbouring bidding-zone prices, are excluded in order to avoid using price variables as explanatory factors for the target price itself, since the price series are highly correlated.

The variable *Load* represents the total electricity demand in each country or bidding zone at a given hour. It captures the level of consumption that must be met by the electricity system. Higher load is generally expected to increase electricity prices, since additional demand may require the activation of more expensive generation units.

The variable *Renewable* represents electricity generation from renewable sources. In the Energy Charts framework, renewable generation includes sources such as wind, solar, hydro, biomass, and other renewable technologies, depending on the country-specific generation mix. This variable is relevant because renewable technologies typically have low marginal generation costs. Therefore, an increase in renewable production may reduce wholesale electricity prices through the merit-order effect.

The variable *Non-Renewable* represents electricity generation from conventional non-renewable sources. This category generally includes fossil-fuel-based generation and other non-renewable technologies, depending on the country. It is included because non-renewable generation is often associated with higher marginal costs and greater exposure to fuel and carbon prices.

The variable *Cross-border electricity trading* represents the net contribution of international electricity exchanges. It captures the role of imports and exports in balancing domestic supply and demand. This variable is important because European electricity markets are interconnected, and cross-border flows may affect domestic prices by increasing available supply or by linking the domestic market to neighbouring price areas.

The variable *European CO_2 emission allowances* represents the price of emission allowances in the European carbon market. It is included because fossil-fuel power plants must account for the cost of CO_2 emissions. An increase in the carbon price can raise the marginal cost of non-renewable electricity generation and, consequently, may increase wholesale electricity prices.

For France, the variable *Nuclear* is also included. It represents electricity generation from nuclear power plants. This variable is particularly relevant for the French market because nuclear generation accounts for a large share of domestic electricity production. Higher nuclear availability can decrease electricity prices.

For Italy, France, and Germany, the variable *Pumped-storage hydro consumption* is included when available. This variable represents electricity consumed by pumped-storage hydroelectric plants in order to pump water to an upper reservoir. It is not final electricity demand, but a storage-related consumption component. It can be relevant for price dynamics because pumped-storage plants tend to consume electricity when prices are relatively low and release electricity when prices are higher.

3.2.4 Resolution and Main Characteristics

The European electricity market series are recorded at an hourly frequency, providing 24 observations per day for each market. This makes them directly comparable with the Italian PUN series, which is also available at hourly resolution.

From a statistical perspective, these price series display strong intra-day variation, recurrent weekly and annual patterns, and periods of high volatility. However, the intensity and timing of these patterns may differ across countries because of differences in demand behaviour, renewable penetration, generation technologies, interconnection capacity, and market structure.

France is strongly influenced by the role of nuclear generation and by domestic demand patterns. Germany is characterized by a high penetration of renewable generation, especially wind and solar power, which can strongly affect short-term price variability and the occurrence of low or negative prices. Denmark, and in particular DK1, is closely connected to the Nordic and Central European electricity systems and is strongly affected by renewable generation and cross-border exchanges.

For these reasons, the European electricity price dataset is useful not only for forecasting, but also for comparing how different market structures and generation mixes affect the temporal behaviour of electricity prices. The dataset allows the thesis to evaluate whether price dynamics are consistent across countries or whether they change according to the characteristics of each national market.

3.3 Italian Electricity Load Dataset

The third dataset contains historical observations of the Italian electricity load. This series represents the total amount of electricity demanded by the country, Sardegna and Sicilia over time and is therefore directly connected to the physical operation of the national grid.

3.3.1 Source and Temporal Scope

The load data were obtained from the official data portal of Terna https://dati.terna.it/download-center, the Italian Transmission System Operator responsible for managing the national high voltage transmission grid. The dataset covers the period from January 1, 2021, to the most recent available observations, in 2026.

3.3.2 Resolution and Main Characteristics

The load series is available at a fifteen minute resolution. This means that each day contains 96 observations, allowing the dataset to capture the intra day dynamics of electricity consumption with a high level of detail.

This temporal granularity is particularly relevant for short term load forecasting. Electricity demand can change substantially within the same hour, especially during morning and evening ramps. A fifteen minute resolution makes it possible to observe these variations directly, whereas hourly averages may smooth out part of the short term variability.

The dataset is also aligned with the increasing importance of finer time intervals in European electricity markets. Shorter settlement and trading periods have become particularly relevant because of the growing penetration of renewable generation, whose production can change rapidly and is more difficult to predict.

From a statistical point of view, the load series displays several overlapping seasonal patterns. A daily cycle reflects the regular alternation between hours of low and high consumption. A weekly cycle is associated with differences between working days and weekends. A broader annual component captures the effect of seasons, weather conditions, and changes in social and economic activity throughout the year.

These characteristics make the load dataset particularly suitable for evaluating forecasting models that are able to handle multiple seasonalities and changing demand patterns.

Chapter 4

PUN Results

4.1 Preprocessing

Before extending the analysis to a broader European setting, this chapter focuses on the Italian *Prezzo Unico Nazionale*, usually referred to as PUN. The PUN represents the main reference price for the Italian day ahead electricity market and is therefore a natural starting point for the empirical study of electricity price forecasting.

Accurate price forecasts are important for several operational and financial decisions, including the definition of supply contracts, portfolio management, hedging strategies, and the evaluation of market risk, as explained in [14]. For this reason, the analysis of the PUN requires a dataset that is sufficiently long to capture both regular market dynamics and exceptional periods of volatility.

The PUN dataset used in this thesis has hourly resolution and covers more than two decades, from 2004 to 2026. This long temporal span makes it possible to observe different market regimes, recurrent seasonal patterns, and major structural shocks affecting the Italian electricity market.

A particularly important preprocessing issue in hourly electricity price data is related to Daylight Saving Time. Since the electricity market operates continuously, the seasonal change of clocks can create irregularities in the time index. These anomalies must be corrected because most statistical forecasting models require a regular temporal structure.

Two different cases were considered:

- **Spring time change**: when clocks move forward by one hour, one hourly timestamp is missing from the series. This creates a gap in the regular sequence of observations. The missing value was replaced through linear interpolation, using the observed prices immediately before and after the gap.
- **Autumn time change**: when clocks move backward by one hour, the same timestamp may appear twice. In this case, two market prices are associated with the same hourly index. To preserve a unique and regular time index without discarding information, duplicated values were aggregated by taking their arithmetic average.

4.2 Preliminary Analysis

Before fitting the forecasting models, an exploratory analysis was conducted to understand the main characteristics of the electricity price series. Visual inspection provides useful information on the structure that forecasting models must be able to capture.

4.2.1 Whole Time Series

The first exploratory step consists of plotting the complete PUN time series from 2004 to 2026. The plot shows that the series is clearly nonstationary. During the first part of the sample, prices fluctuate around a relatively stable level, although with visible seasonal movements and occasional spikes. In contrast, the later part of the sample contains much stronger volatility.

The most evident feature is the exceptional increase in prices observed in 2022. This period corresponds to the European energy crisis, during which wholesale electricity prices increased sharply and displayed extreme volatility. The crisis represents a relevant structural shock in the series and has a substantial effect on model evaluation. Forecasting models trained on historical data may struggle to reproduce such an extreme episode, since its magnitude is very different from the regular behaviour observed in previous years.

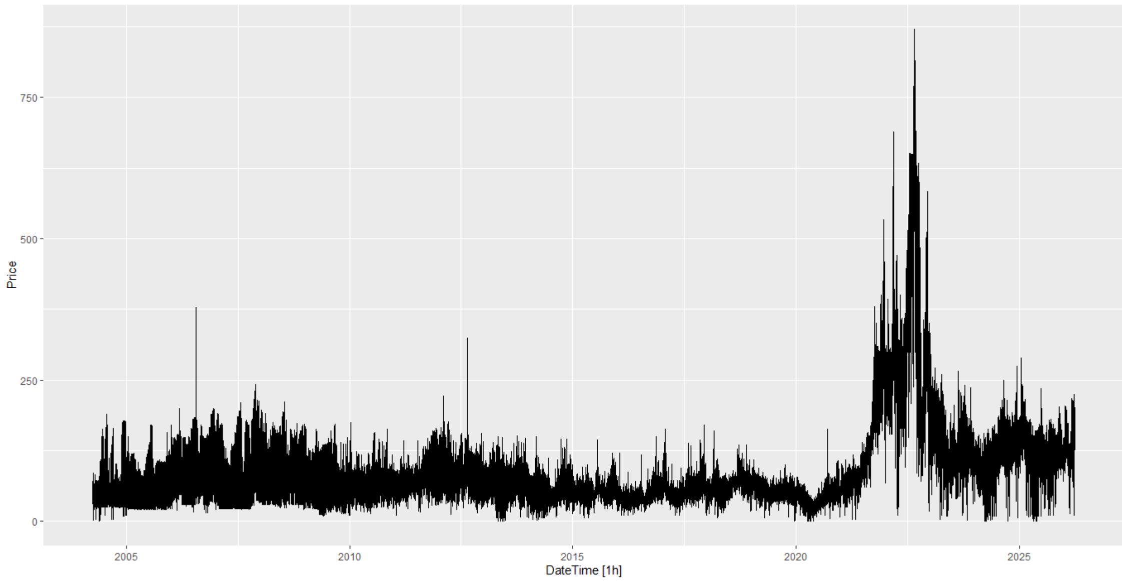

Figure 4.1: PUN: Time series

4.2.2 Seasonal Plots

The complete time series provides a global view of the data, but it is not sufficient to identify the recurring seasonal patterns in detail. For this reason, seasonal plots were used to analyse the behaviour of prices over different calendar frequencies.

These plots make it possible to compare the same seasonal position across different years, weeks, or days. Some trajectories deviate substantially from the regular pattern, especially those associated with the exceptional market conditions of 2022. These observations confirm the presence of outlying periods that must be taken into account when interpreting the forecasting results.

- **Annual seasonality**: the yearly seasonal plot overlays the observations across the months of the year. Although the absolute price level changes over time, the plot suggests the presence of recurrent annual behaviour. Prices tend to be affected by seasonal demand conditions, with different patterns during mild months and during periods of higher heating or cooling needs.

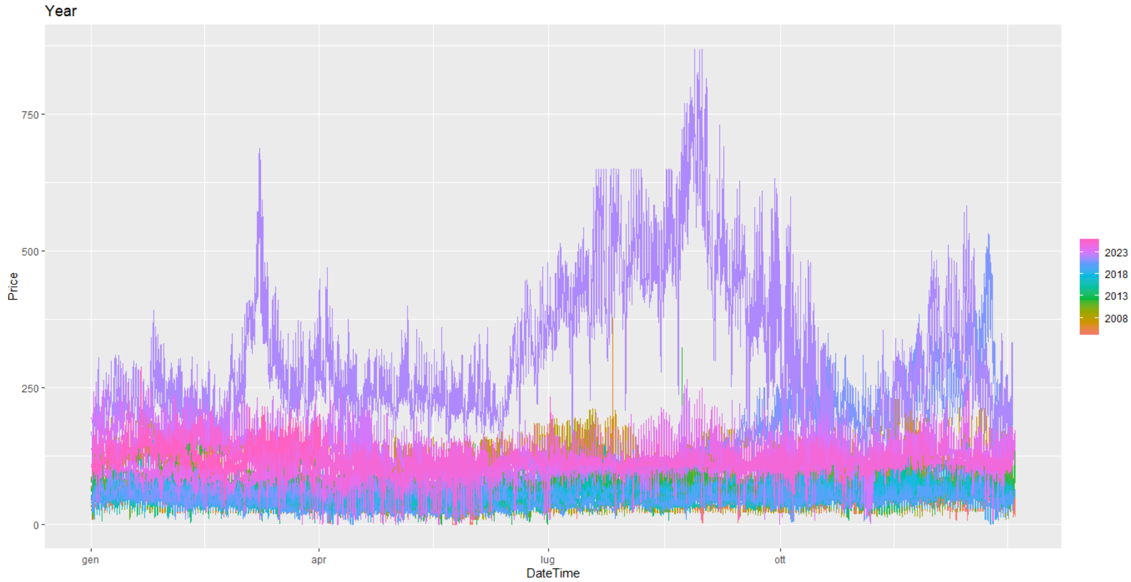

Figure 4.2: PUN: Yearly seasonality

- **Weekly seasonality**: the weekly seasonal plot highlights the difference between working days and weekends. Prices usually decrease during weekends, when industrial and commercial electricity consumption is lower. This confirms the presence of a clear weekly structure in the PUN series.

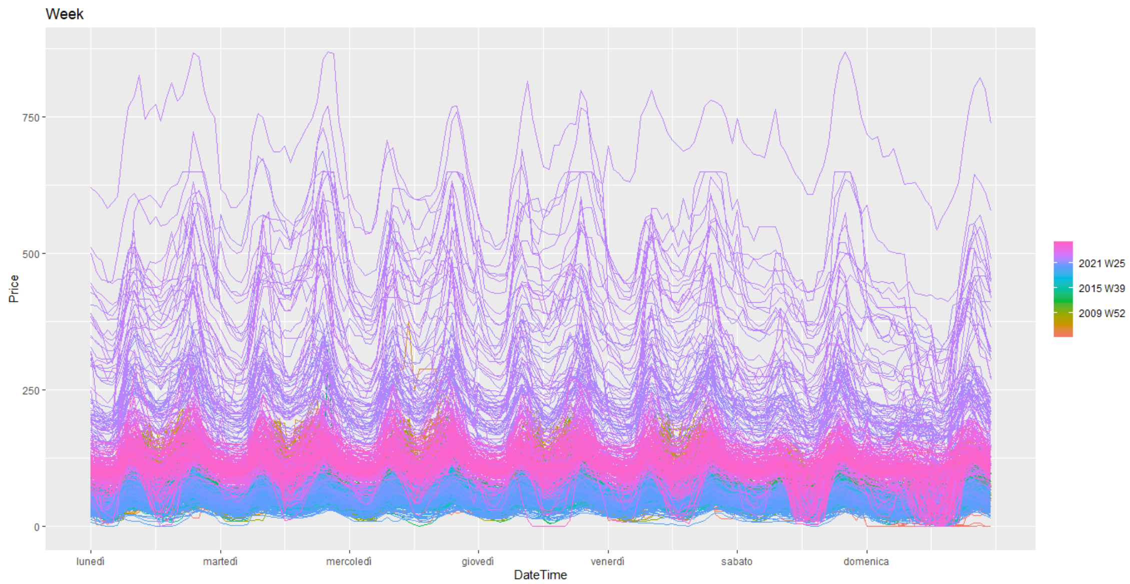

Figure 4.3: PUN: Weekly seasonality

- **Daily seasonality**: the daily seasonal plot shows the intra day structure of electricity prices. Prices generally follow the daily demand profile, with higher values during the hours in which consumption is stronger. In recent years, the increasing role of solar production may also affect the shape of the daily curve, especially during central hours of the day.

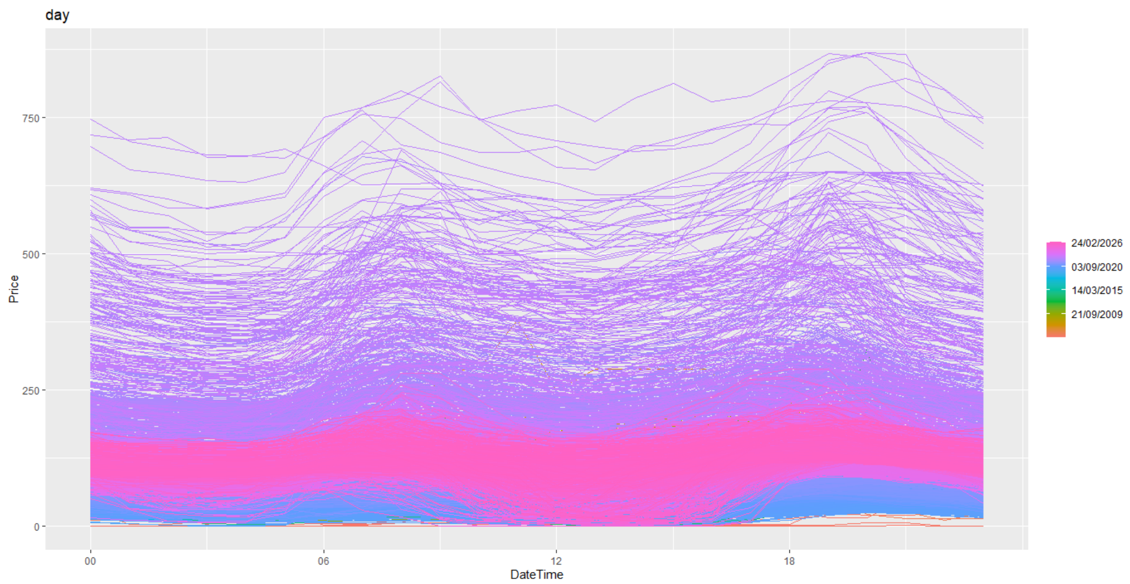

Figure 4.4: PUN: Daily seasonality

Overall, the exploratory analysis shows that the PUN series is characterized by multiple seasonalities, nonstationarity, and extreme observations. These features justify the use of forecasting methods able to account for complex temporal patterns.

At the same time, the daily seasonal plot also suggests a possible limitation of modelling the PUN as a single hourly time series. Prices observed at different hours of the day may follow different dynamics. Night hours, morning ramps, central hours, and evening peaks are affected by different combinations of demand, renewable generation, and market clearing conditions. This point will become central in the following chapter, where the hourly structure of electricity prices is modelled explicitly.

4.3 Data Preparation and Cross Validation Settings

Before model estimation, a Box Cox transformation was applied to stabilize the variance of the series. This step is useful because the variability of electricity prices is not constant over time and tends to increase during periods of high price levels. The transformation parameter was selected using the Guerrero algorithm, which provides a data driven criterion for choosing the value that makes the variance more homogeneous across different portions of the sample.

After the transformation, a Time Series Cross Validation strategy was implemented to evaluate the forecasting performance of the models. The cross validation procedure was configured as follows:

- **Initial training window**: the first estimation window contained 80 weeks of consecutive hourly observations. This initial sample provides enough information to estimate the main daily and weekly seasonal patterns.
- **Expanding step**: after each validation step, the training sample was expanded by adding 80 additional weeks of observations. The model was then estimated again on the enlarged sample and evaluated on the following forecast horizon.

This expanding window design allows the evaluation to reproduce a realistic forecasting setting, in which more historical information becomes available as time progresses.

4.3.1 Preliminary Model Comparison

The first cross validation results show clear differences in predictive performance across models. Prophet produced substantially larger RMSE values than the other approaches. This suggests that, in this application, its flexible curve fitting structure was not sufficient to reproduce the sharp intra day movements and highly autocorrelated dynamics of the PUN series.

The best overall result in the initial comparison was obtained by the standard ARIMA model, followed by the ARIMA model combined with STL decomposition and the ETS model combined with STL decomposition. The results are reported in Table 4.1.

Table 4.1: PUN: Model comparison

Model	**RMSE**
ARIMA	25.9
ARIMA + STL	30.9
ETS + STL	31.0
ETS	41.6
Prophet	126.0
Prophet + STL	126.0

Because of the poor performance of Prophet in the preliminary evaluation, the following analysis focuses mainly on the classical statistical models and on their STL based variants.

4.3.2 Model Selection and Effect of the 2022 Crisis

A more detailed inspection of the cross validation results shows that one validation fold has a disproportionate impact on the average RMSE. In particular, fold 12 corresponds to the period affected by the exceptional price dynamics of 2022 and produces an error much larger than those observed in the remaining folds.

Table 4.2 reports the fold specific RMSE values for the ARIMA model combined with STL decomposition. The value associated with fold 12 is clearly separated from the rest of the distribution, confirming the effect of the energy crisis on the evaluation procedure.

Table 4.2: PUN: Cross validation RMSE for ARIMA + STL

ID Fold	Model	RMSE
1	ARIMA + STL	14.10
2	ARIMA + STL	24.10
3	ARIMA + STL	26.20
4	ARIMA + STL	14.40
5	ARIMA + STL	17.90
6	ARIMA + STL	15.90
7	ARIMA + STL	11.90
8	ARIMA + STL	9.68
9	ARIMA + STL	15.00
10	ARIMA + STL	9.30
11	ARIMA + STL	18.20
12	ARIMA + STL	336.00
13	ARIMA + STL	20.70
14	ARIMA + STL	19.50

For this reason, an additional comparison was performed by excluding fold 12. This does not remove the importance of the crisis period, but it allows the models to be compared under more regular market conditions. The resulting RMSE values are reported in Table 4.3.

Table 4.3: PUN: Model comparison excluding fold 12

Model	Mean RMSE
ARIMA + STL	16.7
ARIMA	42.5
ETS + STL	206.0
ETS	439.0

After excluding the crisis fold, the ARIMA model combined with STL decomposition provides the lowest average RMSE. The use of ARIMA models for electricity price forecasting is consistent with early applications to next-day electricity prices [5]. This result suggests that, under ordinary market conditions, removing the seasonal components before fitting the ARIMA model improves forecasting accuracy.

The same conclusion is supported by the comparison across different forecast horizons. Figure 4.5 shows how the RMSE changes as the prediction horizon increases from one to sixteen weeks. The ARIMA + STL specification remains the most reliable model across the considered horizons, confirming its suitability for the PUN dataset when the series is modelled as a single hourly sequence.

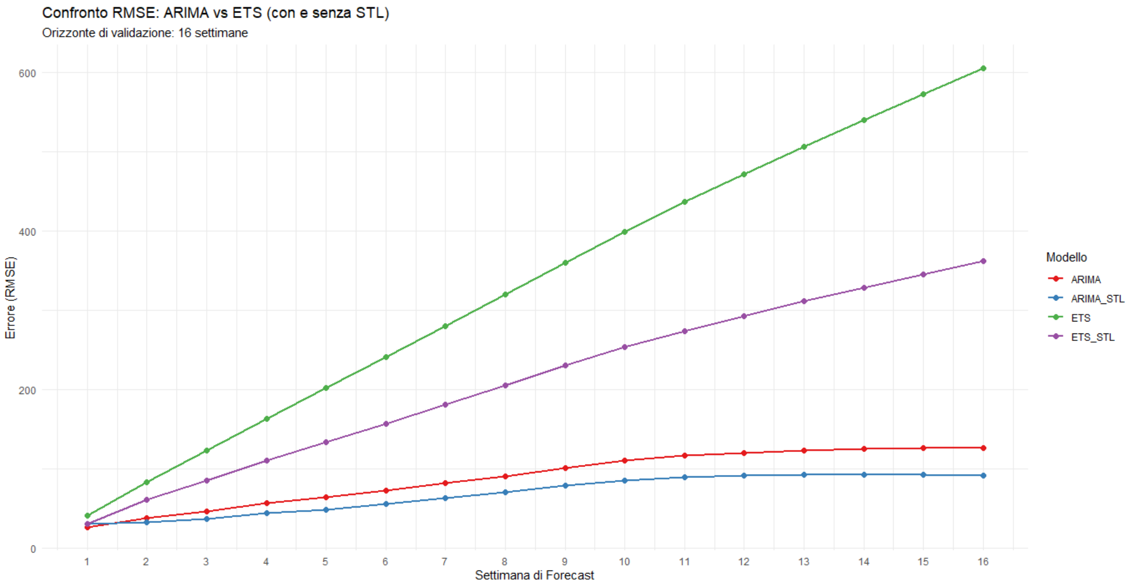

Figure 4.5: PUN: RMSE comparison across forecast horizons

The selected ARIMA specifications during cross validation are reported in Table 4.4. The results show that the automated selection procedure frequently chooses models with first order differencing, confirming that differencing is required to stabilize the mean of the series. The most frequently selected specification is ARIMA$(2, 1, 2)$.

Table 4.4: PUN: Selected ARIMA specifications

Model Specification	**Times Selected**
ARIMA(2,1,2)	4
ARIMA(1,1,5)	3
ARIMA(2,1,2) with drift	3
ARIMA(4,1,2)	2
ARIMA(2,1,3)	1
ARIMA(3,1,3)	1

4.4 Prediction Interval

To evaluate the uncertainty associated with the forecasts, prediction intervals were analysed using a classical training and test set structure. The training set contains the historical observations from 2004 up to the beginning of 2017, while the test set contains the observations from 2017 to 2026.

The ARIMA + STL model was fitted on the training set and then used to forecast the first six months of 2017. Figure 4.6 compares the observed PUN values with the corresponding prediction intervals.

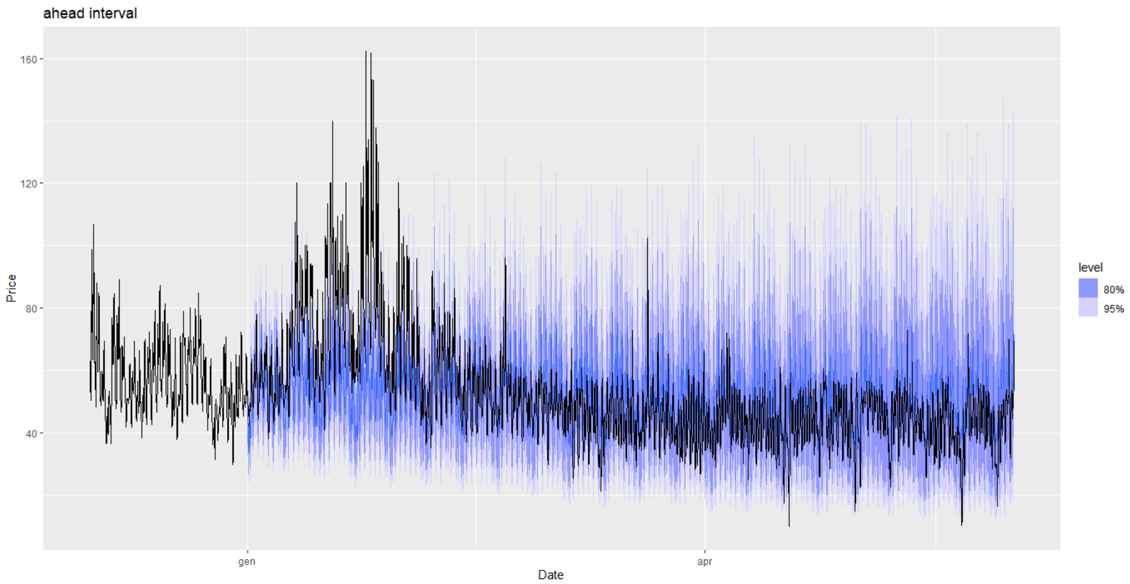

Figure 4.6: PUN: Prediction interval

The prediction interval captures the general seasonal structure of the series, indicating that the model is able to reproduce part of the regular temporal behaviour of the PUN. However, some sudden price movements remain outside the interval, especially at the beginning of the forecast period. This suggests that, although the ARIMA + STL model performs well in terms of average accuracy, it does not fully capture all short term dependencies and market shocks.

4.5 Discussion and Link with the European Analysis

The results obtained in this chapter provide a first benchmark for electricity price forecasting on the Italian market. When the PUN is modelled as a complete hourly time series, ARIMA+STL model provide the most reliable performance.

This confirms the usefulness of decomposing the series into trend, seasonal, and remainder components before applying autoregressive modelling.

At the same time, the results also reveal the limitations of the whole series approach. A single model estimated on the full hourly sequence has to represent night prices, morning ramps, central hours, and evening peaks within the same specification. This can make the forecasting task very complex.

This suggests that further improvements may be obtained either by changing the structure of the forecasting problem. Instead of modelling the PUN as one complete hourly time series, the next analysis divides each national price series into 24 separate time series, as suggested by [18], one for each hour of the day. This hourly modelling strategy is then applied not only to Italy, but also to France, Germany, and Denmark. This cross-country perspective is relevant because European electricity markets are increasingly interconnected, while still presenting different national generation structures and market conditions [13]. The purpose is to verify whether separating the intra day structure before model estimation improves forecasting accuracy and whether the relative performance of ARIMA, ETS, and GARCH changes across European markets, while trying to capture prices behaviour during 2022 crisis.

Chapter 5

European Prices Results, and time specific models

5.1 Introduction

After analysing the Italian PUN series in detail, this chapter extends the empirical study to a broader European setting. The objective is to compare the forecasting behaviour of electricity spot prices across four markets: Italy, France, Germany, and Denmark. This extension is useful because European electricity markets are increasingly interconnected through market coupling and cross border exchanges, but they still differ substantially in terms of generation mix, demand structure, renewable penetration, and bidding zone configuration [13].

The analysis focuses on hourly day observations. For Italy, the reference series is the *Prezzo Unico Nazionale*. For France, Germany, and Denmark, the price series were obtained from the Energy Charts platform. In the Danish case, the DK1 bidding zone is used as the reference series.

A central methodological choice in this chapter is the treatment of the hourly structure of electricity prices. Rather than modelling each country as a single hourly time series, the data are divided into 24 separate series, one for each hour of the day.

This approach is motivated by the strong intra day structure of electricity prices. Prices at different hours of the day may follow substantially different dynamics.

The choice of separating the series by hour is inspired by the approach discussed in the scientific papers [21] and [18], where electricity spot prices in several European markets are analysed by considering different seasonalities of price data. The 24 hourly series provide a natural way to study price behaviour across the different hours of the day.

The empirical results support this modelling choice, even if tested on 2022 crisi, In particular, the Italian case shows a substantial improvement. Table 5.1 compares the Mean Absolute Error obtained by models estimated on the whole time series with the average Mean Absolute Error obtained across the 24 hourly models. The evaluation is performed using Time Series Cross Validation.

Table 5.1: Italy: MAE comparison between whole series models and hourly models

Model	**MAE**
ETS hourly models	15.463
GARCH hourly models	16.975
ARIMA hourly models	17.025
ARIMA + STL whole series	92.656
ARIMA whole series	120.203
ETS + STL whole series	292.596
ETS whole series	471.360

The difference between the two modelling strategies is substantial. Models estimated on the whole hourly series produce much larger errors. This suggests that the hourly decomposition reduces the complexity of the original series and allows each model to focus on a more homogeneous temporal pattern.

An additional result emerges from the comparison among hourly models. The results obtained in this thesis indicate that ETS models are highly competitive in this setting. In the Italian case, the ETS specification estimated separately for each hour achieves the lowest average MAE among the tested models.

The following sections analyse the results for Italy, France, Germany, and Denmark. For each country, the forecasting models are evaluated under the hourly decomposition strategy, and their performance is compared in terms of prediction accuracy. This allows the chapter to assess not only which model performs best overall, but also whether different European markets require different forecasting approaches.

5.2 Construction of the Hourly Series

For each country, the original hourly price dataset is divided according to the variable identifying the hour of the day. Before model estimation, the target variable is transformed through first order differencing This transformation compares the price observed at the same hour on two consecutive days. Differencing is introduced to reduce nonstationarity and to make the dynamics of each hourly series easier to model.

5.3 Model Specifications and Cross Validation

The hourly forecasting analysis considers three model classes: ARIMA, ETS, and GARCH. Each model is estimated separately for every country and every hour of the day. Therefore, for each country, 24 models are estimated for each forecasting method.

The evaluation period starts on January 1, 2020. All observations before this date are used as the initial training sample. The remaining observations are used for the cross validation procedure.

5.3.1 ARIMA Model

The ARIMA model is estimated on the differenced hourly price series. The specification used in the empirical analysis is ARIMA(1,1,1). The model includes one autoregressive term and one moving average term. The ARIMA model is evaluated through an expanding window Cross Validation.

At each iteration, the model is fitted on the available training data and then used to forecast the following 90 daily observations of the same hour. After each fold, the training window is expanded by 90 observations.

5.3.2 ETS Model

The ETS model is also estimated on the differenced hourly price series. The specification used in the analysis is an additive error, additive trend, and additive seasonal model with weekly period, ETS(A,A,A)7. The presence of a weekly seasonal component is important because, after splitting the original series by hour, each hourly series has daily frequency. Therefore, the main calendar seasonality that remains in each series is the weekly cycle.

The ETS model is evaluated using the same expanding window structure used for ARIMA. The initial training period contains all observations before January 1, 2020, the forecast horizon is equal to 90 days, and the training window is expanded by 90 observations at each step.

5.3.3 GARCH Model

The GARCH model is included to account for time varying volatility in electricity prices. While ARIMA and ETS mainly focus on the conditional mean of the series, GARCH models are designed to represent changes in conditional variance over time.

The model is estimated on the differenced hourly price series. The conditional mean equation is specified as an ARiMA$(1, 1, 1)$ process.

The conditional variance follows a standard sGARCH$(1, 1)$ model.

The innovations are assumed to follow a Student t distribution. This choice is useful because price variations may display occasional large deviations.

In the implementation, the model is estimated using `R` library rugarch. The rolling evaluation uses an expanding window. The model produces one step ahead forecasts and is refitted every 90 observations. This refitting frequency reduces the computational burden while still allowing the model parameters to adapt over time.

5.4 Differences in Average Electricity Prices

The average electricity prices observed during the study period are reported in Table 5.2.

Table 5.2: Average electricity prices by country.

Country	Average Price (EUR/MWh)
Italy	126.02
France	101.88
Germany	95.65
Denmark	87.98

Italy exhibits the highest average electricity price among the four countries. This is consistent with the amount of energy the country imports with respect to the others.

France has a lower average price than Italy. A possible explanation is the important role of nuclear generation, nuclear power provides a large amount of low marginal cost electricity and can therefore contribute to reducing prices.

Germany has an even lower average price. This result is consistent with the strong negative correlation of price, with renewable energy production, which indicates that renewable production has a substantial price reducing effect in the German market.

Denmark records the lowest average price among the countries considered. This may be related to the high degree of market integration, the role of renewable generation, and the strong connection with neighboring electricity systems.

5.5 Hourly Forecasting Performance

After splitting each national price series into 24 hourly time series, the forecasting models were estimated and evaluated. This representation is particularly informative because it allows the analysis to identify whether forecasting accuracy is homogeneous throughout the day or whether some hours are systematically more difficult to predict. Electricity prices are strongly affected by the daily demand profile, renewable production, and market clearing conditions.

Table 5.3 reports the MAE obtained for each hour, country, and model. The comparison includes ARIMA, ETS, and GARCH specifications.

Table 5.3: European prices: Hourly MAE by country and model

Hour	ITA GARCH	ITA ARIMA	ITA ETS	FRA GARCH	FRA ARIMA	FRA ETS	GER GARCH	GER ARIMA	GER ETS	DK GARCH	DK ARIMA	DK ETS
0	11.4	11.7	11.6	18.0	18.3	18.4	22.7	23.0	23.3	24.0	24.1	24.3
1	11.5	11.7	11.5	18.4	18.7	18.9	22.8	23.0	23.8	24.0	24.0	24.7
2	11.8	12.1	12.0	18.9	19.1	19.2	22.6	22.7	23.1	24.2	24.1	24.2
3	12.2	12.6	12.6	19.9	20.0	20.1	23.3	23.2	24.0	25.7	25.3	25.6
4	12.1	12.5	12.6	23.7	23.2	22.6	29.0	27.6	27.2	30.9	29.5	29.9
5	11.4	11.6	11.7	31.9	30.1	26.7	35.8	34.2	30.7	37.7	36.1	34.3
6	14.2	14.3	12.9	36.5	33.7	28.0	39.3	37.4	32.2	40.9	39.6	36.0
7	20.9	20.9	16.8	33.0	31.0	24.8	38.1	36.0	32.0	39.4	38.0	35.0
8	25.6	25.6	19.8	30.4	28.8	24.5	36.7	35.2	32.3	37.0	35.6	33.7
9	22.9	23.7	19.2	28.7	27.2	23.6	36.0	34.7	32.7	35.1	34.1	32.9
10	20.0	20.5	17.7	26.3	25.3	22.6	35.7	34.4	33.0	34.6	33.6	33.0
11	19.9	20.2	17.9	27.2	25.8	23.4	36.5	35.1	33.4	34.5	33.4	32.6
12	18.4	18.5	17.3	29.5	27.9	24.6	37.6	36.0	34.1	35.2	34.0	33.0
13	21.3	21.2	18.6	30.0	28.0	24.4	36.7	34.7	32.3	35.1	33.4	31.9
14	23.7	23.1	19.6	27.8	26.4	23.5	34.7	32.7	30.4	34.4	32.9	31.4
15	23.4	22.7	19.4	25.5	24.6	22.0	32.4	30.8	29.0	34.4	32.8	31.7
16	21.2	20.4	17.8	23.5	23.0	20.9	31.5	30.9	30.0	36.3	35.1	34.5
17	18.8	18.2	16.3	23.6	23.3	22.0	34.6	34.6	33.7	40.0	39.0	38.5
18	18.6	18.4	17.4	20.9	20.7	19.6	31.4	31.3	30.1	36.4	35.3	34.6
19	19.1	19.1	18.3	18.5	18.5	18.0	25.2	25.4	24.8	29.5	28.9	28.4
20	15.9	16.1	16.3	16.5	16.5	16.1	21.7	22.2	22.0	25.7	25.7	25.5
21	12.6	12.8	12.8	16.1	16.3	16.5	20.9	21.4	21.2	24.0	24.1	24.5
22	10.3	10.3	10.4	17.3	17.9	17.9	22.4	23.0	22.8	25.2	25.5	25.4
23	10.2	10.4	10.6	18.5	18.7	18.9	23.0	23.3	23.4	24.6	24.8	24.7

A first result is that the hourly modelling strategy produces much more stable results than the models estimated on the whole hourly series, also on 2022 crisi. This confirms the evidence already discussed for the Italian case.

To summarize the information contained in Table 5.3, Table 5.4 reports the average MAE across the 24 hourly models for each country and model.

Table 5.4: European prices: Average MAE across hourly models

Country	GARCH	ARIMA	ETS
Italy	16.975	17.025	15.463
France	24.192	23.458	21.550
Germany	30.442	29.700	28.396
Denmark	32.033	31.204	30.429

The results show that ETS obtains the lowest average MAE in all four countries. In Italy, the ETS average MAE is equal to 15.46, improving over both GARCH and ARIMA. The same pattern is observed in France, Germany, and Denmark, where ETS remains the best performing model in terms of average error.

The empirical evidence obtained in this thesis suggests that ETS models can be highly effective when they are applied to hourly separated electricity price series.

5.6 Best Model by Hour

A more detailed comparison can be obtained by identifying, for each country and each hour, the model with the lowest MAE. This provides information not only on average performance, but also on the consistency of each model across the 24 hours of the day.

Table 5.5 reports how many times each model is selected as the best performing specification across the 24 hourly series.

Table 5.5: European prices: Number of hourly series in which each model obtains the lowest MAE

Country	ARIMA	ETS	GARCH
Denmark	4	16	4
France	0	17	7
Germany	1	16	7
Italy	1	15	8

The dominance of ETS is visible in all countries. This means that ETS is not only the best model on average, but also the most frequently preferred model across the daily price profile.

GARCH remains competitive in several hours, ARIMA is the best model only rarely.

5.7 Hourly Error Profiles

The behaviour of the forecasting errors over the 24 hours of the day is shown in Figure 5.1. The graph plots the MAE as a function of the hour of the day, focusing on ETS model.

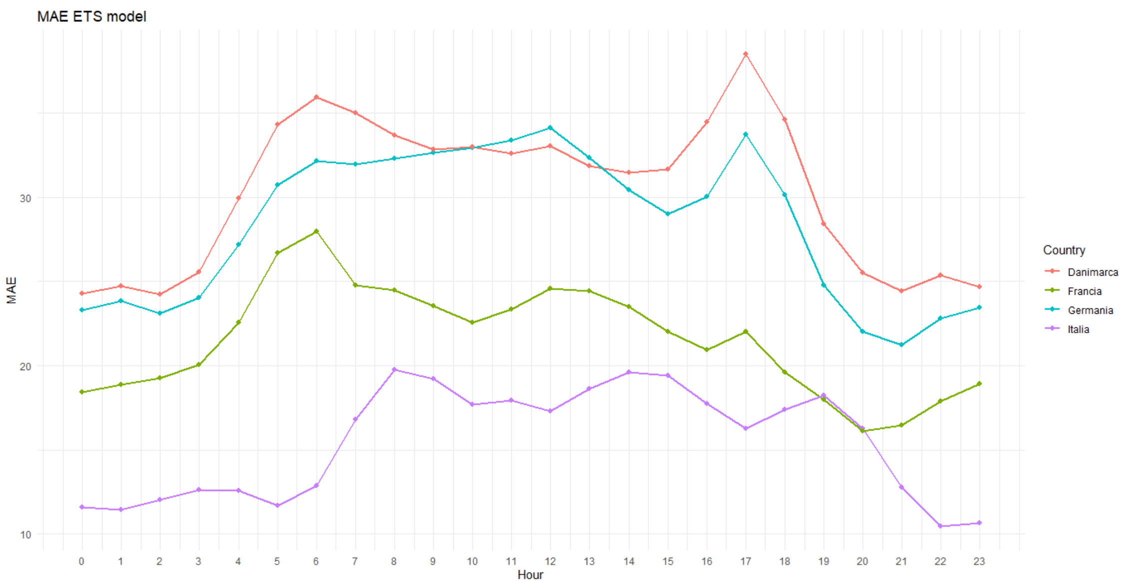

Figure 5.1: European prices: Hourly MAE profile for ETS

The figure highlights two important patterns. First, forecasting accuracy differs substantially across countries. The lowest errors are observed for Italy, followed by France, Germany, and Denmark.

Second, the forecasting error is not constant throughout the day. In all countries, MAE values are generally lower during night hours and higher during daytime hours. Night prices tend to be more stable because demand is lower and the system is usually less affected by rapid changes in consumption. During the day, instead, electricity prices are influenced by stronger demand variation, industrial and commercial activity, renewable production, and market ramping effects. These factors increase the variability of prices and make forecasting more difficult.

The increase in MAE during daytime hours is particularly visible around the morning and central hours of the day. These periods often correspond to rapid changes in load and generation conditions. In markets with high renewable penetration, such as Germany and Denmark, the interaction between demand and wind or solar generation may further increase price variability. This can explain why these countries display higher errors than Italy and France.

5.8 Discussion of the European Comparison

The European comparison provides three main findings. First, separating the price data into 24 hourly time series substantially improves forecasting performance with respect to modelling the whole hourly series directly, even on the year following 2020, and during 2022 energy crisis.

Second, ETS models achieve the best overall performance across all countries. The evidence obtained in this thesis suggests that exponential smoothing should be included among the candidate models for European electricity price forecasting.

Third, the forecasting difficulty varies across countries, and across hours.

The presence of higher errors during specific hours and in specific countries suggests that further improvements could be obtained by incorporating other covariates.

Chapter 6

Determinants of Electricity Prices in European Markets

This chapter investigates the main determinants of electricity prices in the four electricity markets considered in this thesis, namely Italy, France, Germany, and Denmark. The objective is not to develop a forecasting model, but rather to identify which structural and market related variables are most strongly associated with electricity price formation in each country.

The analysis is based on linear regression models estimated separately for each country. The estimated coefficients provide information about the direction and relative magnitude of the relationship between electricity prices and the explanatory variables, while the corresponding p-values are used to assess statistical significance.

6.1 Regression Model Specification

For each country, a multiple linear regression model was estimated using the hourly electricity price as the dependent variable. The explanatory variables include electricity demand, renewable and non-renewable generation, cross-border electricity exchanges, carbon emission allowance prices, and country-specific generation technologies.

The general specification can be written as

$$P_t = \beta_0 + \sum_{i=1}^{k} \beta_i X_{i,t} + \sum_{h} \gamma_h H_h + \sum_{d} \delta_d D_d + \varepsilon_t,$$

where:

- P_t denotes the electricity price at hour t;
- $X_{i,t}$ represents the exogenous variables;
- H_h denotes hourly dummy variables;
- D_d denotes "day of week" dummy variables;
- ε_t is the random error term.

All continuous explanatory variables were standardized before estimation. Consequently, the estimated coefficients can be directly compared within each country. A larger coefficient in absolute value indicates a stronger association with electricity prices.

Hourly and weekly dummy variables were included as control variables in order to capture recurrent temporal effects associated with demand patterns, peak hours, and weekend behavior. The reference categories are Monday for the day-of-week effects and hour 0 for the hourly effects.

6.2 Italian Electricity Market

The Italian market shows a strong positive correlation between electricity prices and the European CO_2 emission allowance price. The standardized coefficient of the CO_2 variable is equal to 0.667 and is statistically significant, since its p-value is lower than 0.05. This is the largest coefficient among the structural regressors for Italy, suggesting that carbon costs are the most relevant factor in the Italian model.

Non-renewable generation also has a positive and statistically significant coefficient, equal to 0.288. This result is consistent with the structure of the Italian electricity system, where thermal generation plays an important role in price formation.

Renewable generation has a negative and statistically significant coefficient, equal to -0.233. This suggests that higher renewable production is associated with lower prices.

Cross border electricity net import also have a negative coefficient, equal to -0.120. This indicates that electricity exchanges with neighboring countries are associated with lower Italian prices.

Electricity demand has a positive coefficient equal to 0.100. However, its magnitude is smaller than those of CO_2, non-renewable generation, and renewable generation.

Pumped-storage hydro consumption has a small negative coefficient equal to -0.044.

Regarding temporal effects, several hourly and daily coefficients are statistically significant. Weekend has a positive effect on electricity price, which is surprising, because load is usually higher during weekends, and positive correlated with energy price.

Table 6.1: Italy regression coefficients. Continuous regressors are standardized.

Variable	**Est.**	**SE**	**t-stat.**	**p-val.**
Exogenous regressors				
CO_2 allowances	0.667	0.003	220.676	0.000
Non-renewable gen.	0.288	0.030	9.485	0.000
Renewable gen.	-0.233	0.022	-10.335	0.000
Cross-border trade	-0.119	0.011	-10.604	0.000
Load	0.100	0.037	2.671	0.008
Pumped-storage cons.	-0.044	0.004	-10.478	0.000
Daily effects				
Monday ref.	0.000	–	–	–
Tuesday	-0.031	0.011	-2.786	0.005
Wednesday	-0.047	0.011	-4.246	0.000
Thursday	-0.015	0.011	-1.412	0.158
Friday	0.046	0.011	4.166	0.000
Saturday	0.193	0.011	16.730	0.000
Sunday	0.165	0.012	13.548	0.000
Hourly effects				
Hour 0 ref.	0.000	–	–	–
Hour 1	-0.057	0.020	-2.823	0.005
Hour 2	-0.100	0.020	-4.929	0.000

Variable	Est.	SE	t-stat.	p-val.
Hour 3	-0.150	0.020	-7.399	0.000
Hour 4	-0.215	0.020	-10.575	0.000
Hour 5	-0.218	0.021	-10.436	0.000
Hour 6	-0.063	0.022	-2.864	0.004
Hour 7	0.142	0.023	6.196	0.000
Hour 8	0.277	0.024	11.816	0.000
Hour 9	0.263	0.024	11.085	0.000
Hour 10	0.198	0.024	8.349	0.000
Hour 11	0.157	0.024	6.681	0.000
Hour 12	0.077	0.023	3.285	0.001
Hour 13	-0.008	0.023	-0.342	0.732
Hour 14	-0.066	0.023	-2.886	0.004
Hour 15	-0.131	0.022	-5.819	0.000
Hour 16	-0.164	0.023	-7.241	0.000
Hour 17	-0.134	0.023	-5.783	0.000
Hour 18	-0.054	0.023	-2.349	0.019
Hour 19	0.052	0.022	2.347	0.019
Hour 20	0.031	0.021	1.465	0.143
Hour 21	-0.017	0.021	-0.801	0.423
Hour 22	-0.001	0.021	-0.040	0.968
Hour 23	0.003	0.020	0.122	0.903

6.3 French Electricity Market

The French market is characterized by several important drivers.

The CO_2 allowance price has a positive coefficient equal to 0.448 lower then than italy. Electricity demand similarly has a positive coefficient, equal to 0.448. These two variables are therefore the most important positive drivers in the French model.

Nuclear generation has a negative coefficient equal to -0.412. This is one of the most relevant results for France, since it reflects the central role of nuclear power in the French electricity system. Higher nuclear generation is associated with lower electricity prices.

Renewable generation also has a negative effect. This indicates that renewable production contributes to reducing prices.

Cross-border electricity exchanges and non-renewable generation also have a positive but lower effect.

The temporal controls show that Sunday and Saturday both have negative coefficients relative to Monday, indicating lower prices during the weekend.

Table 6.2: France regression coefficients. Continuous regressors are standardized.

Variable	Est.	SE	t-stat.	p-val.
Exogenous regressors				
CO_2 allowances	0.448	0.003	149.568	0.000
Load	0.448	0.033	13.655	0.000
Nuclear gen.	-0.412	0.022	-18.886	0.000
Renewable gen.	-0.316	0.014	-22.869	0.000
Cross-border trade	0.100	0.018	5.692	0.000

Variable	**Est.**	**SE**	**t-stat.**	**p-val.**
Non-renewable gen.	0.077	0.010	7.781	0.000
Pumped-storage cons.	0.050	0.005	10.235	0.000
Daily effects				
Monday ref.	0.000	–	–	–
Tuesday	0.023	0.010	2.265	0.024
Wednesday	0.035	0.010	3.412	0.001
Thursday	0.037	0.010	3.606	0.000
Friday	0.038	0.010	3.712	0.000
Saturday	-0.052	0.011	-5.020	0.000
Sunday	-0.141	0.011	-13.088	0.000
Hourly effects				
Hour 0 ref.	0.000	–	–	–
Hour 1	0.021	0.019	1.102	0.270
Hour 2	0.047	0.019	2.443	0.015
Hour 3	0.064	0.019	3.317	0.001
Hour 4	0.076	0.019	3.985	0.000
Hour 5	0.106	0.019	5.536	0.000
Hour 6	0.141	0.019	7.213	0.000
Hour 7	0.123	0.020	6.215	0.000
Hour 8	0.100	0.020	5.029	0.000
Hour 9	0.100	0.020	4.970	0.000
Hour 10	0.084	0.020	4.153	0.000
Hour 11	0.075	0.020	3.714	0.000
Hour 12	0.105	0.020	5.282	0.000
Hour 13	0.145	0.020	7.371	0.000
Hour 14	0.191	0.019	9.759	0.000
Hour 15	0.246	0.019	12.595	0.000
Hour 16	0.275	0.020	13.966	0.000
Hour 17	0.269	0.020	13.467	0.000
Hour 18	0.235	0.020	11.686	0.000
Hour 19	0.195	0.020	9.828	0.000
Hour 20	0.122	0.020	6.216	0.000
Hour 21	0.016	0.019	0.834	0.404
Hour 22	-0.029	0.019	-1.478	0.140
Hour 23	-0.019	0.019	-0.981	0.327

6.4 German Electricity Market

Germany displays the strongest renewable effect among all countries considered, surprisingly even higher than denmark. Renewable generation has a coefficient equal to -0.744, which represents the largest coefficient in absolute value among the structural variables of the German model. This indicates a very strong negative correlation between renewable generation and electricity prices.

The CO_2 allowance price is also highly relevant, with a positive effect.

Cross-border electricity exchanges have a negative coefficient equal to -0.370. This suggests that electricity exchanges are associated with lower German prices.

Electricity demand has a positive coefficient equal to 0.272.

Non-renewable generation has a negative and statistically significant coefficient equal to -0.127. This result is surprising but could be explained by lower cost energy production methods, from non renewable sources.

Sunday and Saturday have negative coefficients, as expected.

Table 6.3: Germany regression coefficients. Continuous regressors are standardized.

Variable	**Est.**	**SE**	**t-stat.**	**p-val.**
Exogenous regressors				
Renewable gen.	-0.744	0.015	-50.427	0.000
CO_2 allowances	0.667	0.004	179.016	0.000
Cross-border trade	-0.370	0.008	-46.787	0.000
Load	0.272	0.013	21.476	0.000
Non-renewable gen.	-0.127	0.013	-9.973	0.000
Pumped-storage cons.	0.053	0.004	12.666	0.000
Daily effects				
Monday ref.	0.000	–	–	–
Tuesday	0.040	0.011	3.612	0.000
Wednesday	0.042	0.011	3.774	0.000
Thursday	0.031	0.011	2.759	0.006
Friday	-0.003	0.011	-0.276	0.782
Saturday	-0.157	0.012	-12.983	0.000
Sunday	-0.243	0.013	-18.761	0.000
Hourly effects				
Hour 0 ref.	0.000	–	–	–
Hour 1	0.006	0.021	0.307	0.759
Hour 2	0.009	0.021	0.424	0.672
Hour 3	0.012	0.021	0.582	0.560
Hour 4	0.084	0.021	3.984	0.000
Hour 5	0.185	0.022	8.494	0.000
Hour 6	0.277	0.023	12.191	0.000
Hour 7	0.307	0.023	13.256	0.000
Hour 8	0.337	0.023	14.420	0.000
Hour 9	0.379	0.024	16.024	0.000
Hour 10	0.397	0.024	16.730	0.000
Hour 11	0.377	0.023	16.081	0.000
Hour 12	0.338	0.023	14.654	0.000
Hour 13	0.310	0.023	13.646	0.000
Hour 14	0.283	0.022	12.544	0.000
Hour 15	0.284	0.023	12.470	0.000
Hour 16	0.340	0.023	14.581	0.000
Hour 17	0.373	0.024	15.828	0.000
Hour 18	0.295	0.023	12.783	0.000
Hour 19	0.167	0.022	7.521	0.000
Hour 20	0.084	0.022	3.881	0.000
Hour 21	0.015	0.021	0.731	0.465
Hour 22	-0.011	0.021	-0.510	0.610
Hour 23	-0.003	0.021	-0.146	0.884

6.5 Danish Electricity Market

In the Danish model, the CO_2 allowance price is the most important structural variable. Its coefficient is equal to 0.571.

Non-renewable generation also has a positive coefficient, as expected, equal to 0.368.

Electricity demand has a negative equal to -0.293. This differs from the results obtained for Italy, France, and Germany. The result may reflect the specific structure of the Danish electricity system, where sometimes, because of overloading created by renewable forces, electricity price drops largely.

Renewable generation has a negative but lower effect.

The Danish model exhibits the strongest negative effect on energy price among the countries, caused by weekends.

Table 6.4: Denmark regression coefficients. Continuous regressors are standardized.

Variable	**Est.**	**SE**	**t-stat.**	**p-val.**
Exogenous regressors				
CO_2 allowances	0.571	0.003	175.291	0.000
Non-renewable gen.	0.368	0.005	79.225	0.000
Load	-0.293	0.008	-38.543	0.000
Renewable gen.	-0.139	0.010	-13.792	0.000
Cross-border trade	0.062	0.009	6.864	0.000
Daily effects				
Monday ref.	0.000	–	–	–
Tuesday	0.050	0.012	4.174	0.000
Wednesday	0.045	0.012	3.718	0.000
Thursday	0.034	0.012	2.814	0.005
Friday	-0.049	0.012	-4.086	0.000
Saturday	-0.318	0.013	-25.181	0.000
Sunday	-0.376	0.013	-29.867	0.000
Hourly effects				
Hour 0 ref.	0.000	–	–	–
Hour 1	-0.034	0.022	-1.511	0.131
Hour 2	-0.045	0.022	-2.028	0.043
Hour 3	0.003	0.022	0.155	0.877
Hour 4	0.184	0.022	8.242	0.000
Hour 5	0.421	0.023	18.312	0.000
Hour 6	0.588	0.024	24.736	0.000
Hour 7	0.600	0.024	24.768	0.000
Hour 8	0.572	0.025	23.325	0.000
Hour 9	0.543	0.025	22.019	0.000
Hour 10	0.502	0.025	20.419	0.000
Hour 11	0.455	0.025	18.552	0.000
Hour 12	0.413	0.024	16.983	0.000
Hour 13	0.405	0.024	16.847	0.000
Hour 14	0.440	0.024	18.444	0.000
Hour 15	0.591	0.024	24.274	0.000
Hour 16	0.736	0.025	29.838	0.000
Hour 17	0.745	0.024	30.875	0.000
Hour 18	0.623	0.024	26.566	0.000

Variable	Est.	SE	t-stat.	p-val.
Hour 19	0.467	0.023	20.280	0.000
Hour 20	0.349	0.023	15.376	0.000
Hour 21	0.212	0.022	9.486	0.000
Hour 22	0.108	0.022	4.841	0.000
Hour 23	0.051	0.022	2.298	0.022

6.6 Cross-Country Comparison

The regression results show that the determinants of electricity prices differ substantially across the four European markets.

The CO_2 allowance price is significant in all countries and has a positive large coefficient in every model. This confirms that carbon costs represent a common driver of electricity prices across European electricity markets.

Renewable generation has a negative coefficient in all four countries. This result is consistent with the merit-order effect: when renewable generation increases, electricity prices tend to decrease. The strongest renewable effect is observed in Germany, where the coefficient is equal to -0.744, followed by France, Italy, and, surprisingly last, Denmark.

Non-renewable generation has a positive effect in Italy, France, and Denmark. In Germany, instead, the coefficient negative, but relatively low.

Electricity demand has a positive coefficient in Italy, France, and Germany, while it is negative in Denmark. The Danish result shows the particular state of Denmark energy overload, caused by renewable sources.

Cross-border electricity import reduce prices in Italy and Germany, while they are positively associated with prices in France and Denmark. This indicates that the impact of interconnection depends on the role played by each country within the European electricity market. The more one country overproduces, as France with nuclear and Denmark with renewables, the more import is associated with an increase of the price. while Italy and Germany import large amounts of energy, because it's advantageous to them, indeed it decreases energy price.

Chapter 7

ARIMAX on European Prices

7.1 Introduction

The previous analysis of European electricity prices focused on univariate forecasting models applied separately to each hour of the day. In that setting, the price series of Italy, France, Germany, and Denmark were divided into 24 hourly time series, and each model was estimated independently for every hour. This strategy made it possible to compare the forecasting performance of ARIMA, ETS, and GARCH models under a framework in which the intra day price structure was treated explicitly.

The aim of this chapter is to extend that analysis by introducing exogenous regressors into the ARIMA framework. The resulting model, usually referred to as ARIMAX, combines the autoregressive and moving average structure of ARIMA with additional explanatory variables that may help describe the evolution of electricity prices. The purpose is to verify whether including market related information improves forecasting accuracy with respect to models based only on the past behaviour of the target series.

Electricity prices, are not driven only by their own past values. They are influenced by several market fundamentals, including electricity demand, renewable and non renewable generation, cross border electricity exchanges, and the cost of emission allowances.

The empirical analysis is performed on the same four European markets considered in the previous price chapter: Italy, France, Germany, and Denmark. For Italy, the target variable is the PUN, while for the other countries the day ahead prices refer to the corresponding national or bidding zone series used in the previous analysis. As before, the original hourly datasets are divided into 24 separate time series, one for each hour of the day.

For each hourly series, the target variable is defined as the daily price difference, This transformation is applied in order to make the series more stationary.

7.2 Models specifications

The ARIMAX model is estimated by combining an ARIMA$(1, 0, 1)$ specification with a set of exogenous regressors. The included variables differ slightly across countries, depending on the available data. The common regressors include load, renewable generation, non renewable generation, cross border electricity exchanges, European CO_2 emission allowance prices, and weekend dummy variables. For France, nuclear generation is also included as an additional explanatory variable, given its specific importance in the French electricity system.

The ARIMAX model is then compared with three benchmark models. The first benchmark is a standard ARIMA$(1, 0, 1)$ model without exogenous regressors. This allows the direct contribution of the external variables to be evaluated. The second benchmark is an ETS model with additive error, additive trend, and additive weekly seasonality. The third benchmark is a GARCH model, specified with an ARMA$(0, 1)$ conditional mean, an sGARCH$(1, 1)$ conditional variance, and Student t distributed innovations.

7.3 Cross validation setting, and one step ahead forecasting

The year 2022 is excluded from the estimation sample because it was characterized by highly irregular price dynamics related to the European energy crisis. Including this period in the initial training set could distort the estimation of the parameters and reduce the interpretability of the comparison under more regular market conditions. For this reason, the initial training period starts in 1 January 2023, after the crisis, and extends until 1 January 2024, while 2024 is used as the test period.

The forecasting evaluation is based on a time series cross validation design. The initial training set is used to estimate the first model, which is then tested through a one step ahead forecast. After each forecast, one additional observation is added to the training sample, and the model is estimated again on the expanded dataset. This procedure is repeated recursively throughout the test period. In this way, each forecast is produced using only the information that would have been available until the day before, at the same hour.

The one step ahead structure is particularly appropriate in this context because the objective is to evaluate short term forecasting accuracy. Electricity day ahead prices are strongly affected by recent market conditions, and the ability to update the model as new observations become available is therefore an important part of the forecasting procedure.

7.4 Average MAE by Country

Table 7.1 reports the average MAE obtained for each country and model. The results show that ARIMAX provides the best performance in the analysed markets. This suggests that the inclusion of exogenous regressors could improve forecasting accuracy. Moreover, using only data after 2022 for training, results in far better precision metrics.

Table 7.1: Average MAE by country and model

Country	**ARIMAX**	**ARIMA raw**	**GARCH**	**ETS**
Italy	–	17.025	16.975	15.463
Italy_crisisExcluded	10.59	10.75	11.17	10.76
France	–	23.458	24.192	21.550
France_crisisExcluded	18.57	19.46	20.36	19.06
Germany	–	29.700	30.442	28.396
Germany_crisisExcluded	24.18	25.27	27.05	25.89
Denmark	–	31.204	32.033	30.429
Denmark_crisisExcluded	26.44	26.94	28.86	27.80

7.5 Hourly Error Profiles

The average MAE values summarize the overall performance of each model, but they do not show how the forecasting error changes during the day. For this reason, the MAE is also analysed separately for each hour. Figures 7.1, 7.2, 7.3, and 7.4 report the hourly MAE profiles for Italy, France, Germany, and Denmark respectively. In each figure, one line is reported for each model. For each country we notice that daytime hours prices are harder to predict, than night hours ones. MAE on Italian series has three peaks, indicating there is times in the day when forecasting is harder. France has one visible morning peak. Germany and Denmark have a lower morning peak, and an higher afternoon peak.

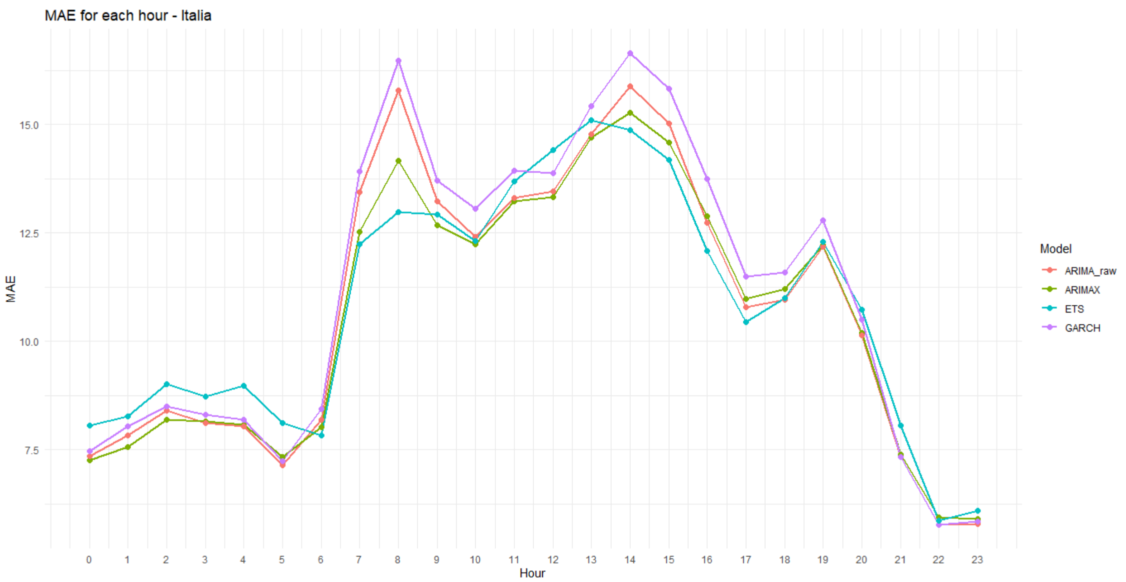

Figure 7.1: MAE for each hour Italy

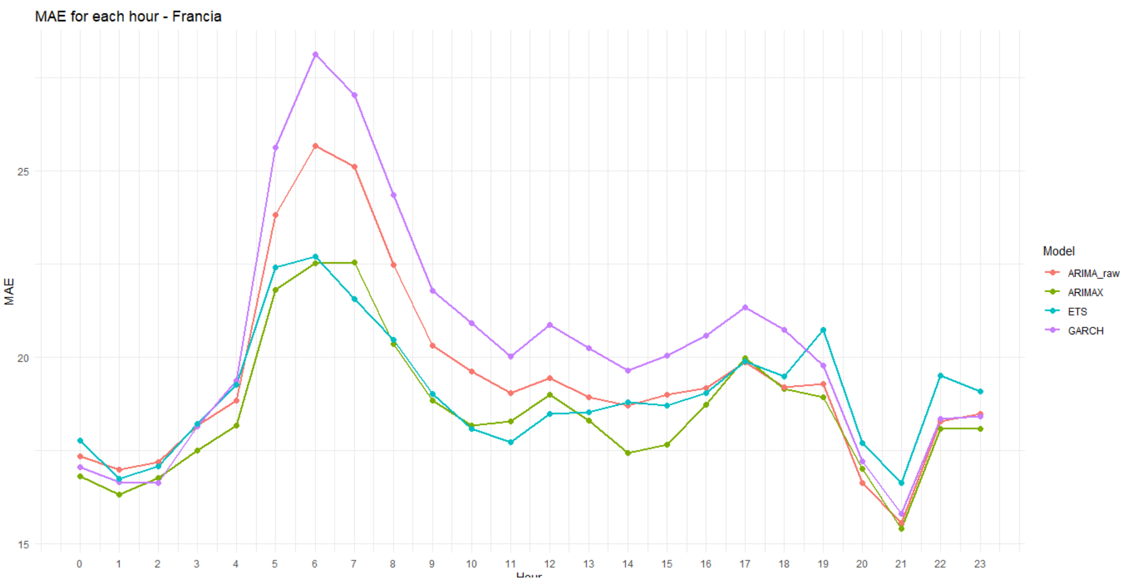

Figure 7.2: MAE for each hour France

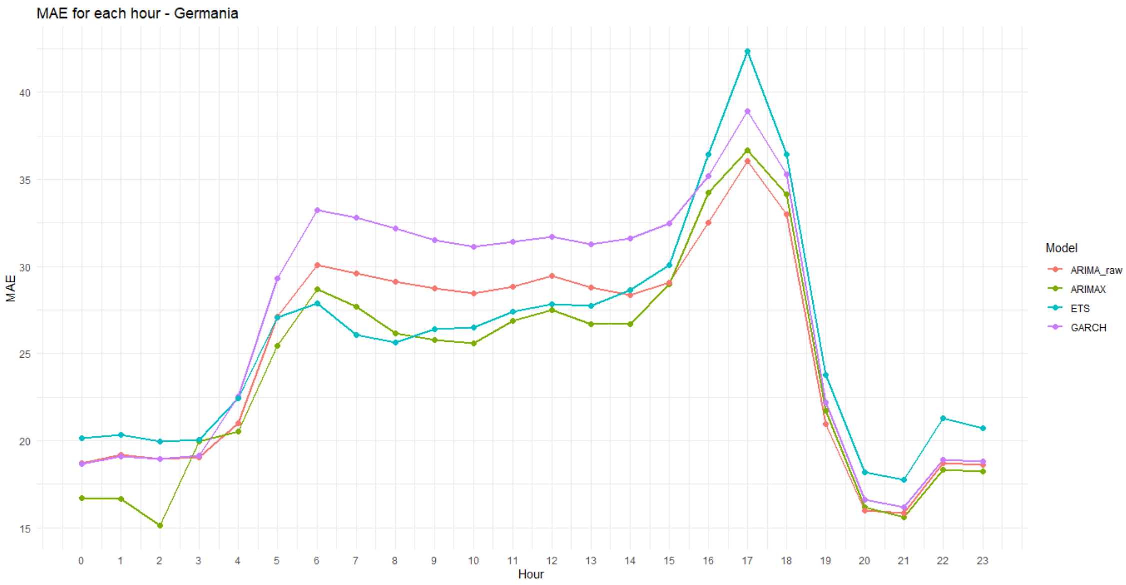

Figure 7.3: MAE for each hour Germany

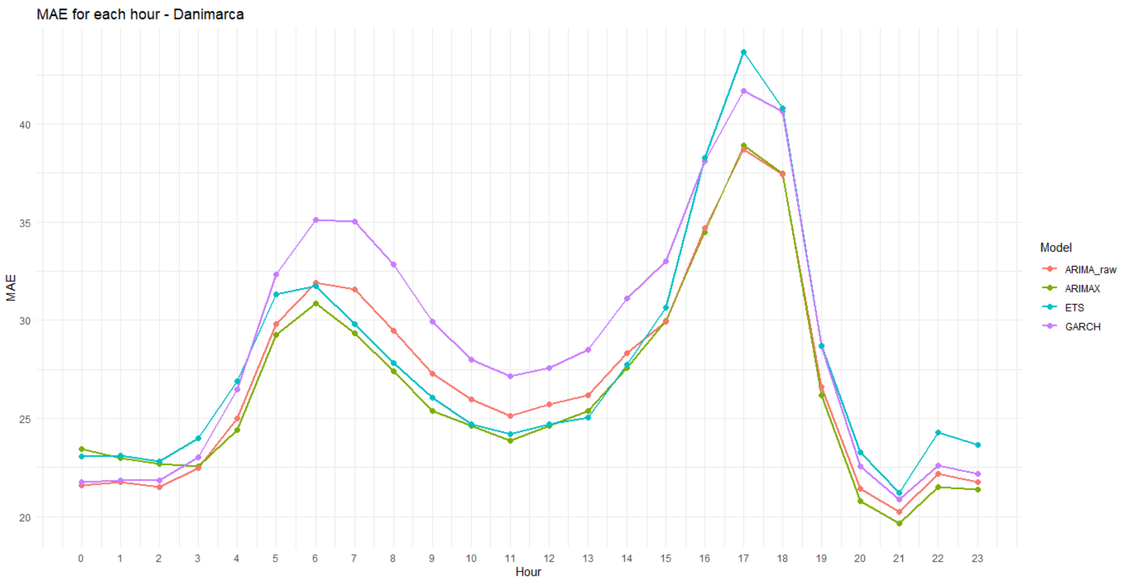

Figure 7.4: MAE for each hour Denmark

Table 7.2: MAE by hour, country and model

Hour	IT AX	IT AR	IT G	IT ETS	FR AX	FR AR	FR G	FR ETS	DE AX	DE AR	DE G	DE ETS	DK AX	DK AR	DK G	DK ETS
0	7.26	7.34	7.48	8.06	16.81	17.34	17.06	17.77	16.70	18.72	18.66	20.16	23.42	21.57	21.76	23.05
1	7.56	7.83	8.04	8.28	16.32	16.99	16.66	16.75	16.65	19.18	19.09	20.34	22.97	21.75	21.83	23.09
2	8.20	8.39	8.50	9.02	16.76	17.18	16.63	17.08	15.11	18.94	18.93	19.96	22.66	21.51	21.83	22.78
3	8.16	8.12	8.30	8.73	17.50	18.16	18.14	18.20	19.94	19.07	19.13	20.03	22.54	22.45	23.00	23.99
4	8.07	8.04	8.19	8.98	18.18	18.84	19.37	19.26	20.53	20.99	22.55	22.45	24.41	25.01	26.48	26.91
5	7.34	7.14	7.24	8.13	21.81	23.82	25.62	22.41	25.46	27.12	29.30	27.09	29.23	29.79	32.31	31.32
6	8.03	8.19	8.45	7.84	22.52	25.66	28.12	22.70	28.71	30.07	33.22	27.90	30.84	31.91	35.09	31.75
7	12.52	13.45	13.92	12.24	22.54	25.11	27.03	21.55	27.69	29.62	32.80	26.08	29.33	31.58	35.04	29.80
8	14.17	15.79	16.47	12.98	20.35	22.46	24.34	20.47	26.17	29.11	32.18	25.65	27.39	29.48	32.83	27.82
9	12.68	13.24	13.70	12.93	18.83	20.31	21.78	19.01	25.80	28.77	31.53	26.43	25.38	27.25	29.94	26.04
10	12.24	12.41	13.06	12.31	18.16	19.62	20.91	18.07	25.60	28.45	31.12	26.49	24.61	25.97	28.01	24.71
11	13.23	13.31	13.94	13.68	18.27	19.04	20.02	17.73	26.89	28.84	31.40	27.42	23.84	25.14	27.15	24.18
12	13.32	13.46	13.87	14.42	19.00	19.43	20.87	18.48	27.49	29.48	31.72	27.83	24.62	25.70	27.56	24.69
13	14.70	14.77	15.41	15.09	18.30	18.92	20.24	18.52	26.71	28.79	31.28	27.77	25.37	26.18	28.48	25.03
14	15.27	15.88	16.65	14.86	17.42	18.71	19.64	18.79	26.72	28.38	31.61	28.63	27.56	28.31	31.10	27.73
15	14.59	15.02	15.83	14.17	17.65	18.99	20.05	18.69	29.01	29.08	32.49	30.09	29.98	29.93	33.02	30.63
16	12.88	12.74	13.75	12.08	18.72	19.16	20.57	19.05	34.25	32.54	35.19	36.45	34.47	34.67	38.11	38.25
17	10.97	10.79	11.48	10.45	19.97	19.87	21.33	19.88	36.66	36.08	38.91	42.38	38.90	38.70	41.70	43.68
18	11.20	10.97	11.58	11.01	19.15	19.19	20.72	19.48	34.17	32.98	35.29	36.42	37.47	37.43	40.63	40.79
19	12.22	12.17	12.79	12.29	18.93	19.28	19.77	20.72	21.73	20.96	22.19	23.80	26.18	26.60	28.66	28.69
20	10.20	10.13	10.51	10.72	17.01	16.62	17.22	17.70	16.18	16.00	16.62	18.21	20.78	21.42	22.57	23.26
21	7.40	7.34	7.33	8.05	15.39	15.57	15.81	16.64	15.61	15.86	16.18	17.78	19.63	20.22	20.87	21.19
22	5.94	5.78	5.76	5.86	18.08	18.28	18.34	19.51	18.32	18.73	18.91	21.32	21.51	22.16	22.61	24.27
23	5.90	5.79	5.85	6.10	18.09	18.48	18.41	19.09	18.22	18.64	18.79	20.72	21.38	21.76	22.17	23.66

Table 7.2 reports the complete hourly MAE values. The table confirms the graphical evidence.

7.5.1 Italian ARIMAX Forecasts

To better illustrate the behaviour of the ARIMAX model, two additional figures are reported for Italy. Figure 7.5 compares the actual and predicted values for the hourly series corresponding to midnight over the full forecasting horizon.

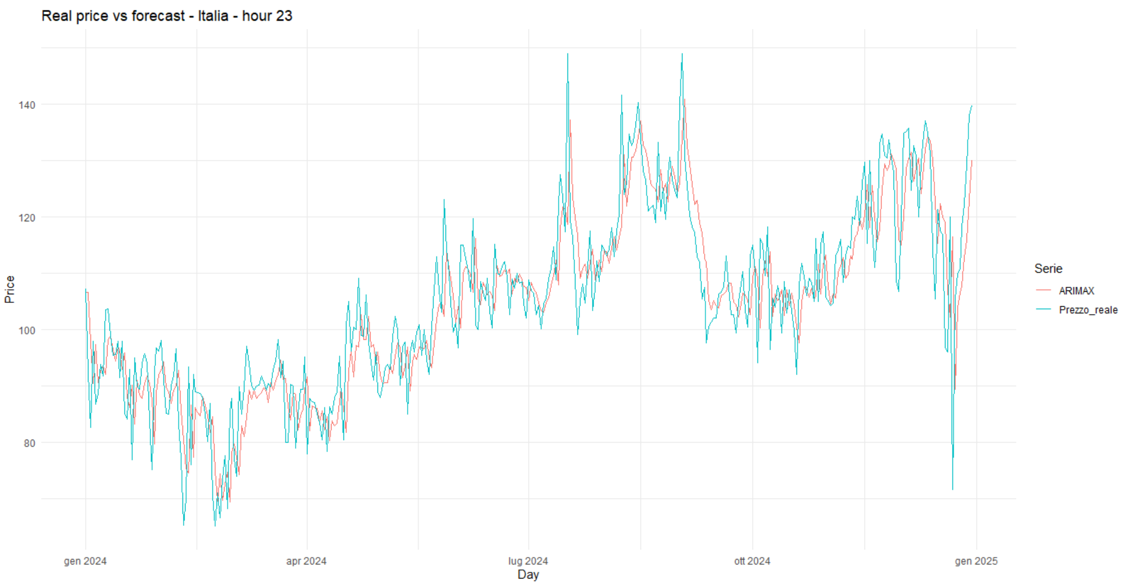

Figure 7.5: Real vs ARIMAX forecast, Italy, hour 23

Figure 7.6 reports the actual and predicted values for all 24 hourly series together, but only for the period from 14 September to 29 September. This shorter horizon is used to improve readability, since plotting all hourly series over the entire test period would make the graph too dense.

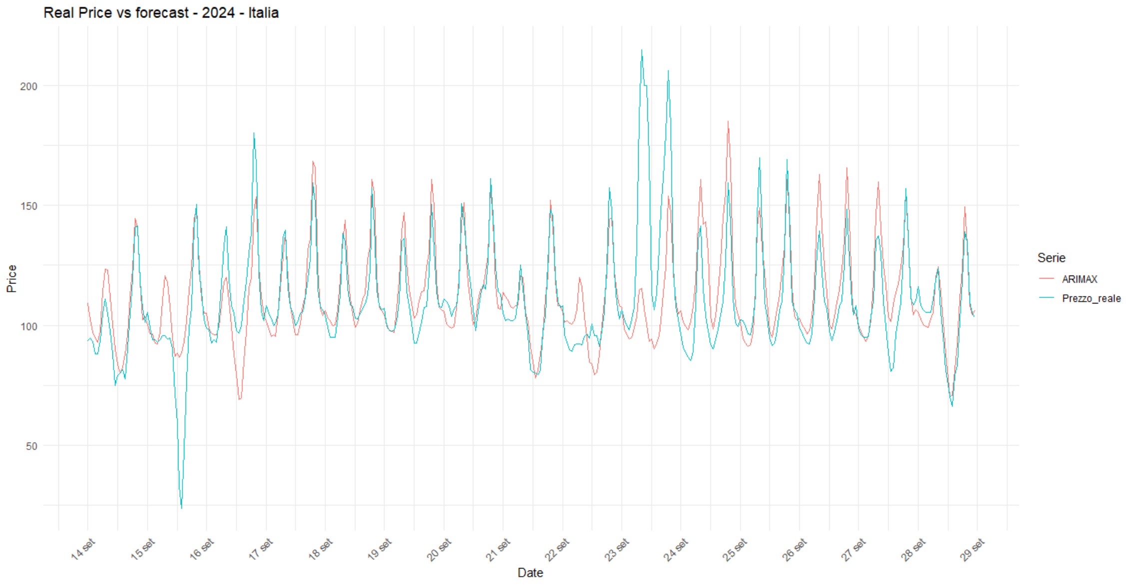

Figure 7.6: Real vs ARIMAX forecast, Italy, 14 sept - 19 sept

These figures show that ARIMAX is able to reproduce the main movements of the Italian price series, but still fails to catch some sudden jump of the serie, indicating that some other covariates should be taken into consideration

Chapter 8

Electricity Load Results: Italy, Sardinia and Sicily

8.1 Introduction

After analysing the financial dynamics of the wholesale electricity market, the focus of the thesis shifts to the physical side of the power system: electricity load. While price forecasting is mainly related to market decisions, hedging strategies, and financial exposure, load forecasting plays a central role in the electricity grid. Accurate demand forecasts are essential for generation scheduling, system balancing, reserve planning, and the overall reliability of the power system.

This chapter studies electricity load at two different geographical levels. First, the aggregated national Italian load is analysed, providing a system wide view of electricity demand. Then, the analysis is extended to Sardinia and Sicily, the two major Italian islands. This regional perspective is relevant because island systems may display specific local dynamics and are subject to additional operational constraints due to their geographical separation from the mainland.

All the load series considered in this chapter are available at a fifteen minute resolution. This means that each day is represented by 96 observations. Such high frequency data provide a detailed description of electricity consumption.

The objective of this chapter is to evaluate forecasting models on the complete fifteen minute time series. This provides a first benchmark for national and regional load forecasting. The following chapter will then investigate whether a different modelling strategy, based on splitting the original series into 96 separate time series, one for each fifteen minute interval of the day, can further improve forecasting performance.

8.2 Preprocessing

The preprocessing procedure was applied consistently to the Italian, Sardinian and Sicilian load series.

A relevant preprocessing issue was related to Daylight Saving Time. Since the electricity system operates continuously, clock changes introduce artificial irregularities in the time index. In fifteen minute data, one hour corresponds to four observations, so the effect of these anomalies is more visible than in hourly datasets.

Two cases were handled separately, similarly to price case:

- **Spring time change**: when clocks move forward by one hour, four observations are missing from the time series. To preserve the regular structure required by the forecasting models, the missing values were replaced through linear interpolation, using the observations immediately before and after the gap.
- **Autumn time change**: when clocks move backward by one hour, four timestamps are duplicated. In this case, two different load values correspond to the same timestamp. To maintain a unique temporal index without discarding information, duplicated observations were aggregated by taking their average.

8.3 Preliminary Analysis

Before fitting the forecasting models, an exploratory analysis was performed to understand the main temporal characteristics of the load series. This step is particularly important for high frequency electricity demand data, where the observed values are strongly influenced by human activity, working schedules, weather conditions, and seasonal effects.

Compared with wholesale electricity prices, electricity load displays a more regular and structured behaviour. Load, is more directly connected to physical consumption patterns and tends to follow more stable daily, weekly, and annual cycles. This suggests that a large part of the temporal structure can be captured through appropriate seasonal modelling.

8.3.1 Load Series

The complete load time series from 2021 to 2026 shows a dense sequence of regular oscillations, reflecting the high frequency nature of the data and the strong intra day variation of electricity demand.

Unlike the PUN series, the load series does not show extreme price like spikes that dominate the scale of the graph. Its behaviour is more stable and continuous, although the amplitude of the oscillations changes over time.

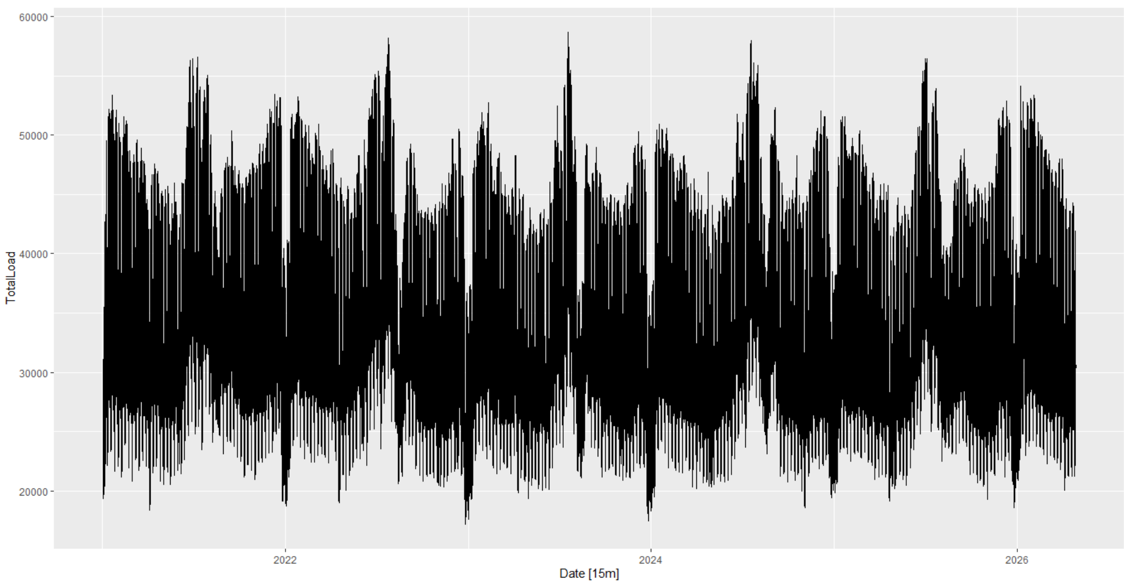

Figure 8.1: Italian load: Time series

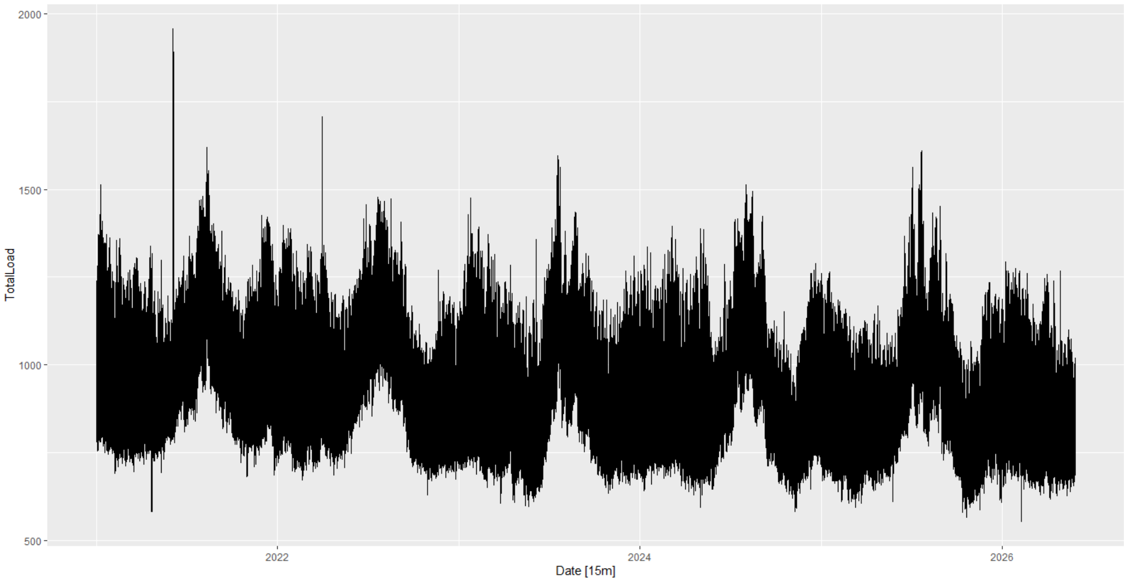

Figure 8.2: Sardinia load: Time series

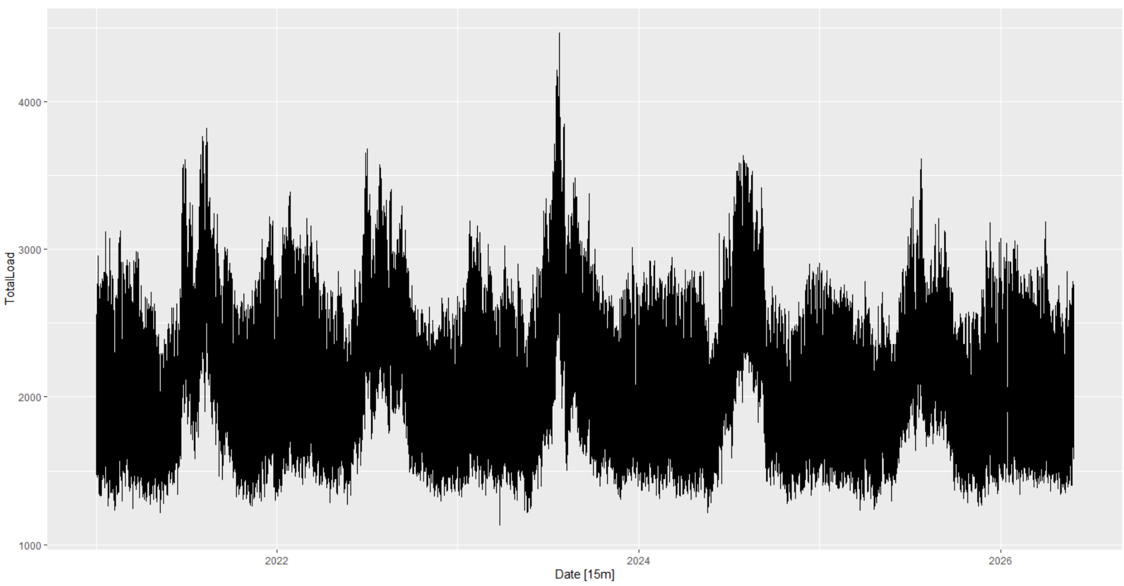

Figure 8.3: Sicily load: Time series

Electricity demand typically exhibits multiple seasonal patterns, including intra-day, weekly and annual cycles, which are central features in short-term load forecasting [16]. The seasonal plots confirm the presence of multiple and well defined seasonal patterns. The yearly plot shows higher demand during summer and winter, and lower demand during spring and autumn. The summer peak is mainly associated with cooling needs, while the winter peak reflects heating and lighting.

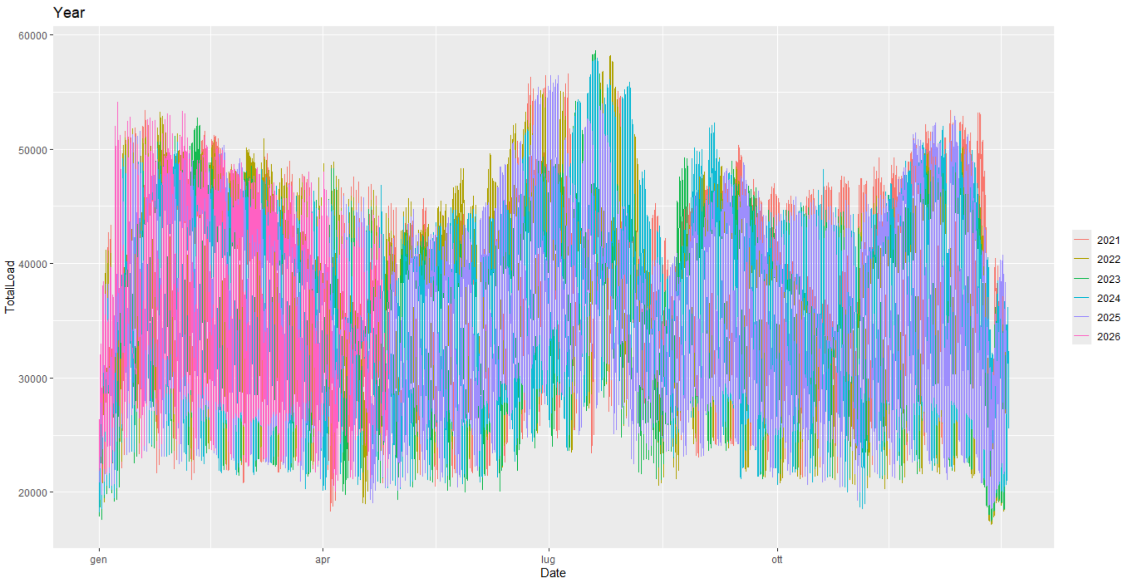

Figure 8.4: Italian load: Yearly seasonality

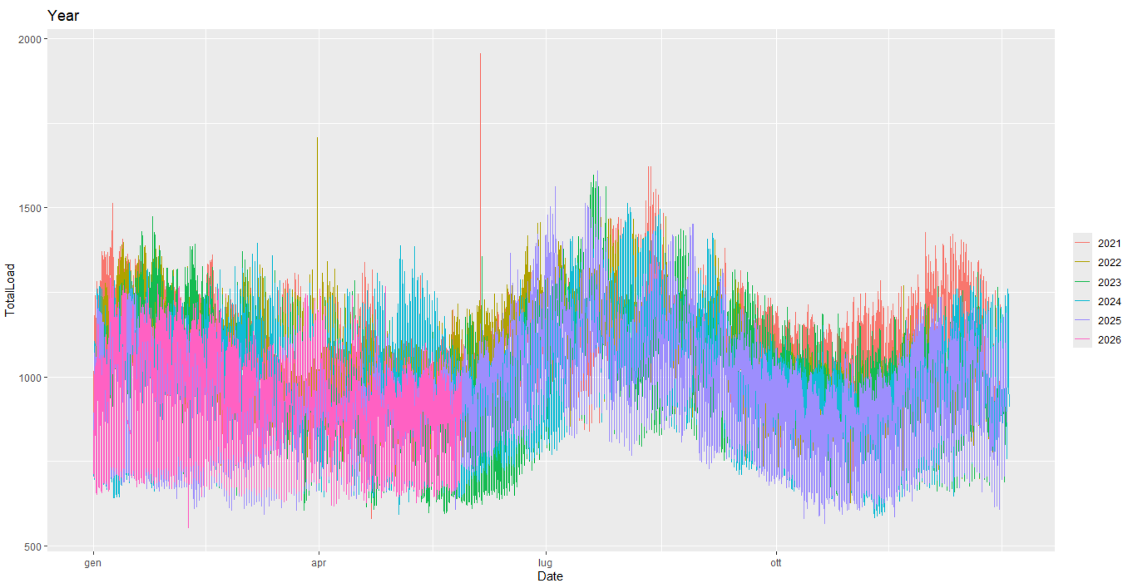

Figure 8.5: Sardinia load: Yearly seasonality

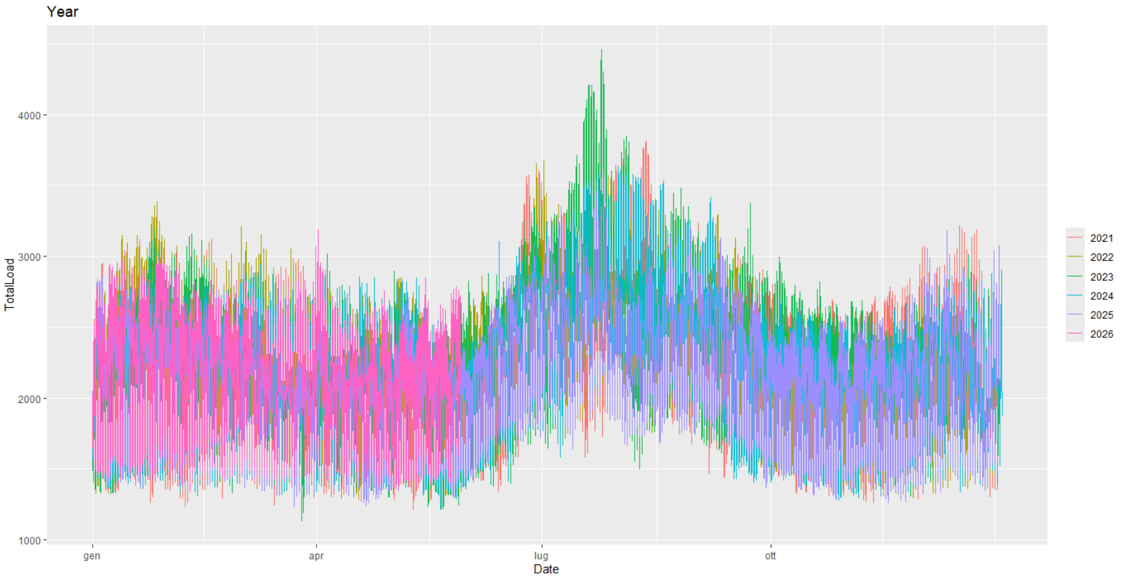

Figure 8.6: Sicily load: Yearly seasonality

The weekly seasonal plot highlights the effect of the working week. From Monday to Friday, demand remains structurally higher because of industrial production, commercial activity, and office work. During the weekend, load decreases, reflecting the reduction in economic activity.

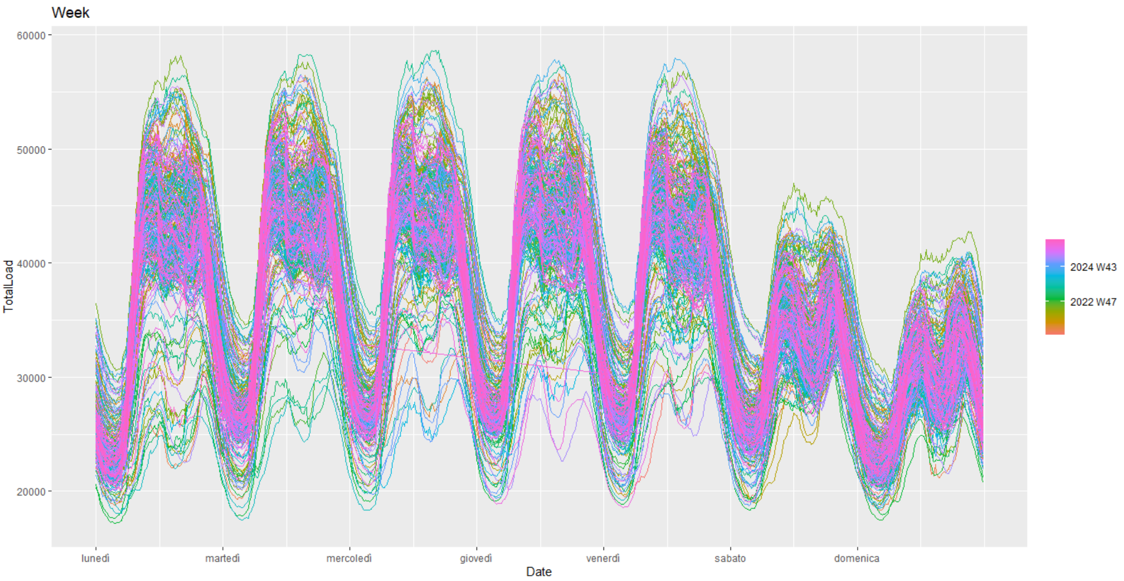

Figure 8.7: Italian load: Weekly seasonality

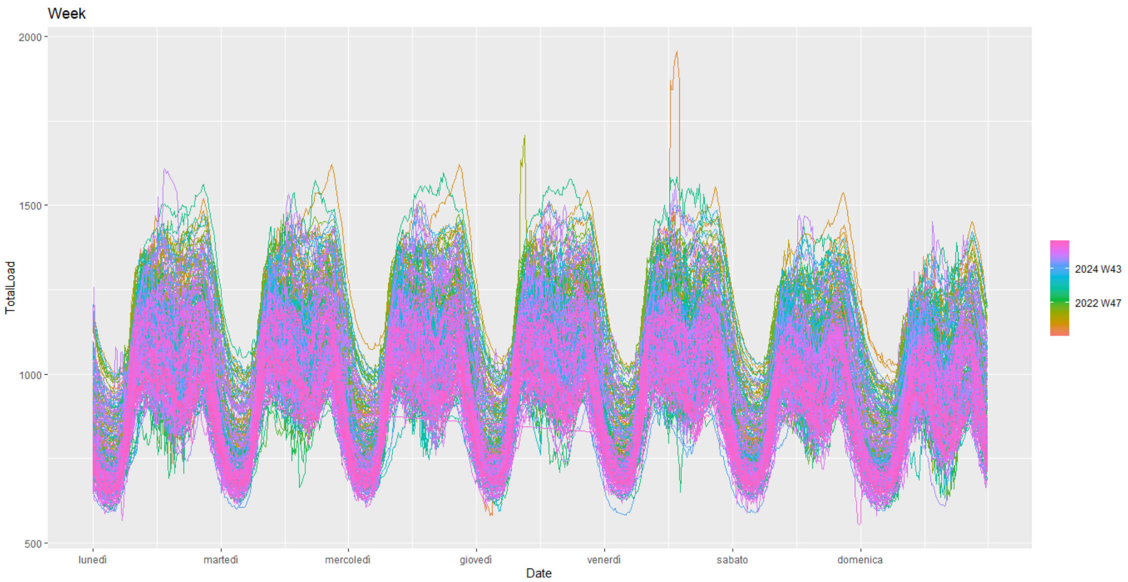

Figure 8.8: Sardinia load: Weekly seasonality

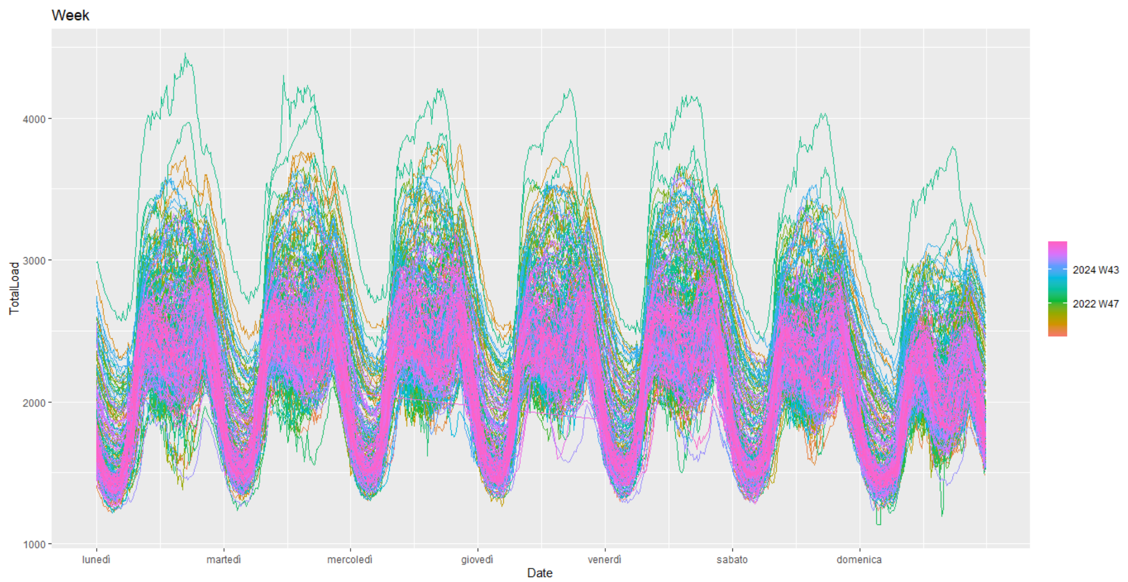

Figure 8.9: Sicily load: Weekly seasonality

The daily seasonal plot shows the intra day load profile. Demand increases sharply in the morning, as households, services, and industrial activities become active. During the central part of the day, the load remains at a relatively high level. In the evening, demand may increase again due to domestic consumption, before declining during the night.

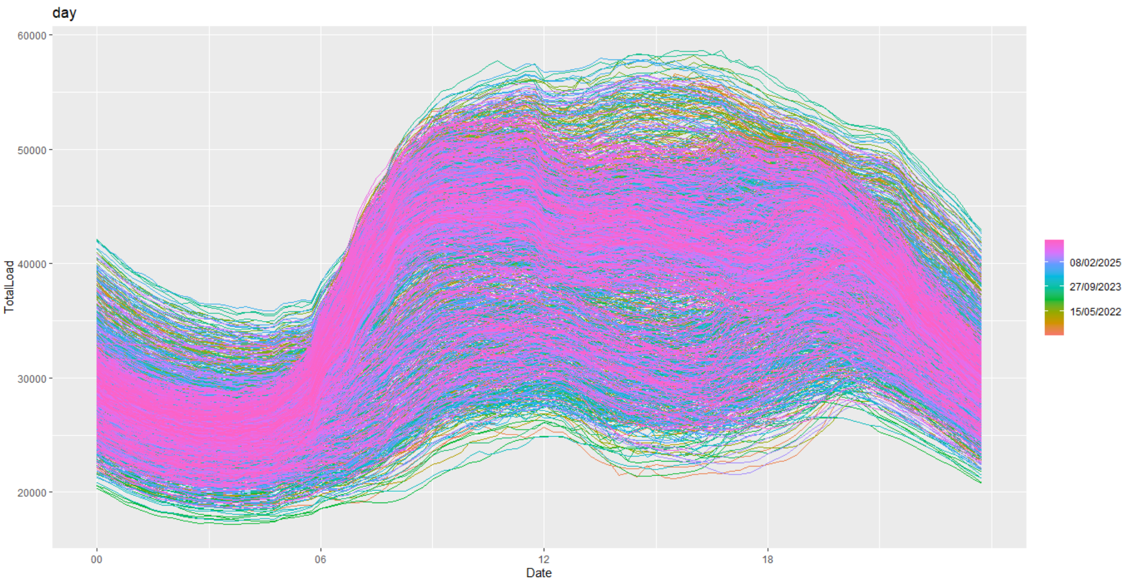

Figure 8.10: Italian load: Daily seasonality

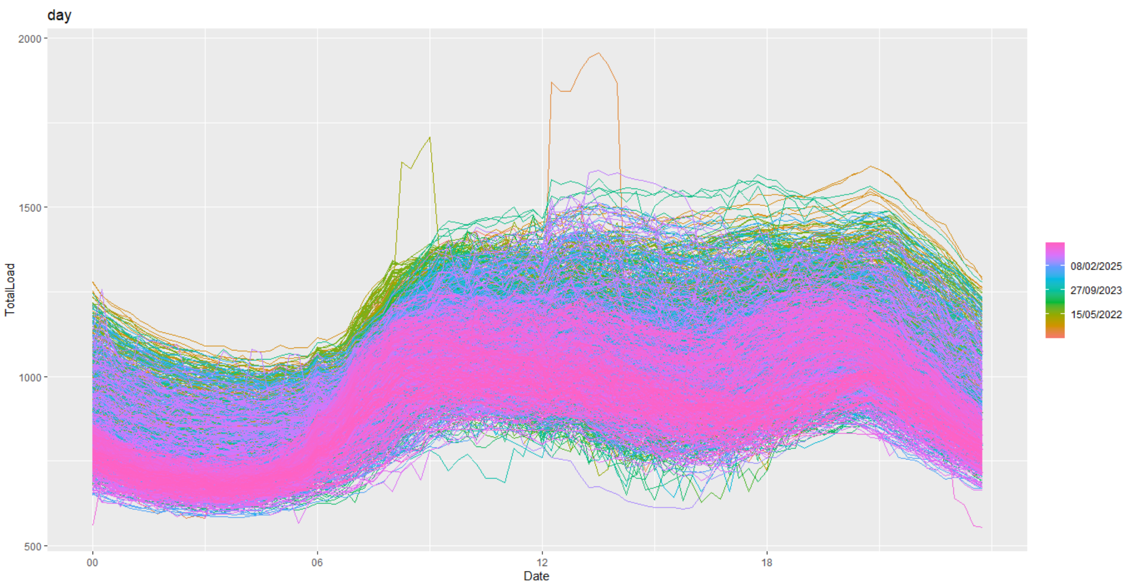

Figure 8.11: Sardinia load: Daily seasonality

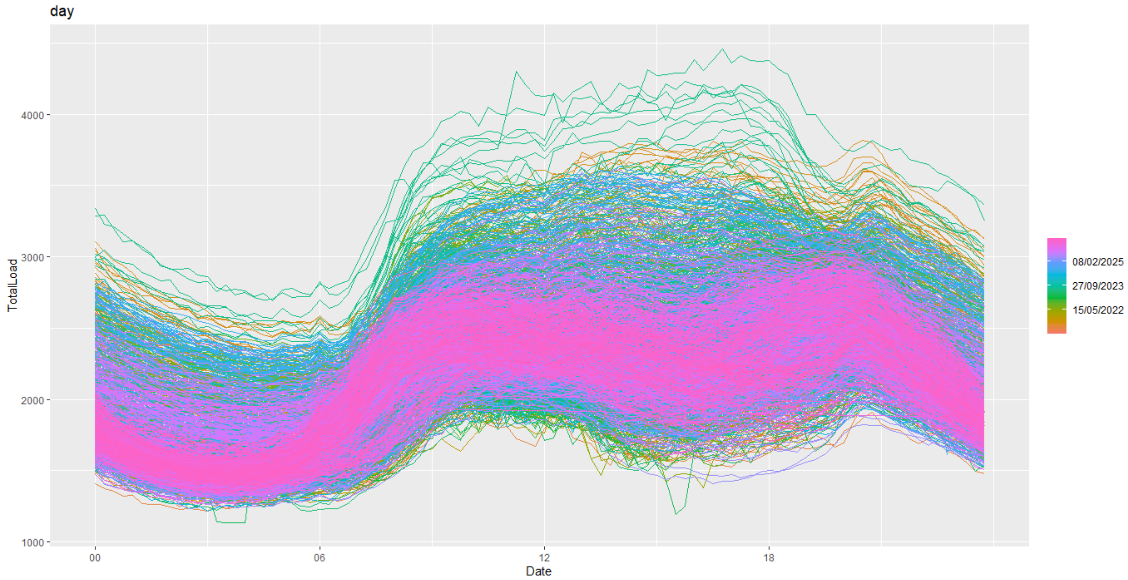

Figure 8.12: Sicily load: Daily seasonality

8.4 Data Preparation and Cross Validation Settings

Before estimating the forecasting models, the load series were transformed in order to improve their statistical properties. Although electricity demand is more regular than electricity prices, its variability is not constant over time. In particular, the amplitude of daily fluctuations tends to become larger during peak seasons and smaller during milder periods.

To reduce this heteroscedastic behaviour, a Box Cox transformation was applied. The transformation parameter was selected using the Guerrero algorithm, which provides a data driven criterion for choosing the value of λ that stabilizes the variance across different portions of the time series.

After the transformation, the forecasting models were evaluated using Time Series Cross Validation.

Two cross validation configurations were used, reflecting the different objectives of the national and regional analyses.

For the national Italian load, the cross validation procedure was configured as follows:

- **Initial training window**: the first estimation window contained 24 weeks of consecutive observations. At fifteen minute resolution, this corresponds to more than 16,000 observations, providing the models with a sufficiently rich sample to learn daily and weekly demand patterns.

- **Expanding step**: after each validation step, the training window was expanded by adding 24 additional weeks of observations. The models were then re-estimated on the enlarged sample and evaluated on the next forecast horizon.

For Sardinia and Sicily, the cross validation strategy was adapted to a short term operational forecasting objective. In this setting, the goal is to evaluate the models over a horizon of approximately one week, which is more directly connected to local dispatching and balancing decisions.

A standard seven day expanding step would have shifted the forecast origin by exactly one week at each iteration. This would imply that the forecast always starts on the same day of the week, potentially introducing a systematic bias in the evaluation. To avoid this issue, an eight day step was used. This choice allows the forecast origin to move across different weekdays over successive folds.

The regional cross validation settings were defined as follows:

- **Initial training window**: the first training sample contained 53 weeks of observations. This provides slightly more than one full year data, allowing the models to learn the annual, weekly, and daily seasonal structure before producing the first forecast.
- **Expanding step**: after each validation fold, the training window was expanded by 8 days, corresponding to 768 observations. This makes the starting point of the forecast move progressively across the days of the week.

This configuration produced more than 200 validation models for the island datasets, providing a detailed and robust estimate of forecasting performance.

8.5 Models and Evaluation Metrics

The predictive accuracy of the forecasting models was evaluated using the Root Mean Square Error. The national analysis considered three main forecasting frameworks: ARIMA, ETS, and Prophet. Each model was evaluated both as a standalone specification and in combination with STL decomposition.

The use of STL is motivated by the structure observed in the preliminary analysis. The series contains several overlapping seasonal patterns, including daily, weekly, and annual components. Standard forecasting models are not always able to represent all these seasonalities simultaneously in an effective way. By applying STL decomposition before model estimation, the main seasonal components are removed from the original series, allowing the forecasting model to focus on the seasonally adjusted dynamics. The final forecast is then reconstructed by adding back the extrapolated seasonal components.

For the regional analysis of Sardinia and Sicily, the model set was selected based on the evidence obtained from the national load results. Since standalone ARIMA, ETS, and Prophet models were clearly outperformed by models combined with STL decomposition, the island analysis focuses on the two most promising specifications: ARIMA + STL and ETS + STL. This choice improves computational efficiency while keeping the analysis focused on the models that showed the best behaviour for high frequency electricity load data.

The next chapter will examine an alternative strategy, in which the original series is split into 96 separate time series, one for each fifteen minute position within the day, similarly to the strategy adopted for prices.

8.6 National Load Forecasting Results

The first comparison across models for the national Italian load is reported in Table 8.1. The results show a clear difference between models estimated directly on the original series and models combined with STL decomposition.

Table 8.1: Italian load: RMSE comparison

Model	RMSE
ETS + STL	2347.16
ARIMA + STL	2351.37
Prophet + STL	2937.00
Prophet	3346.00
ARIMA	16254.49
ETS	19018.41

The best results are obtained by ETS + STL and ARIMA + STL, whose RMSE values are very close. This confirms that STL decomposition is a crucial preprocessing step for this time serie. Without the decomposition, both ARIMA and ETS perform much worse, indicating that the original high frequency series is too complex to be modelled effectively by the standard models.

Prophet performs better when combined with STL, but its error remains higher than those of ETS + STL and ARIMA + STL. This suggests that, in this application, Prophet is less able to capture the sharp and recurrent load variations observed in the Italian demand series. One possible explanation is that Prophet tends to produce smoother forecasts, which may be less effective when the series contains pronounced intra day peaks and troughs.

The comparison across different forecast horizons provides additional information on model performance. Figure 8.13 shows how the RMSE changes as the prediction horizon increases from one to sixteen weeks.

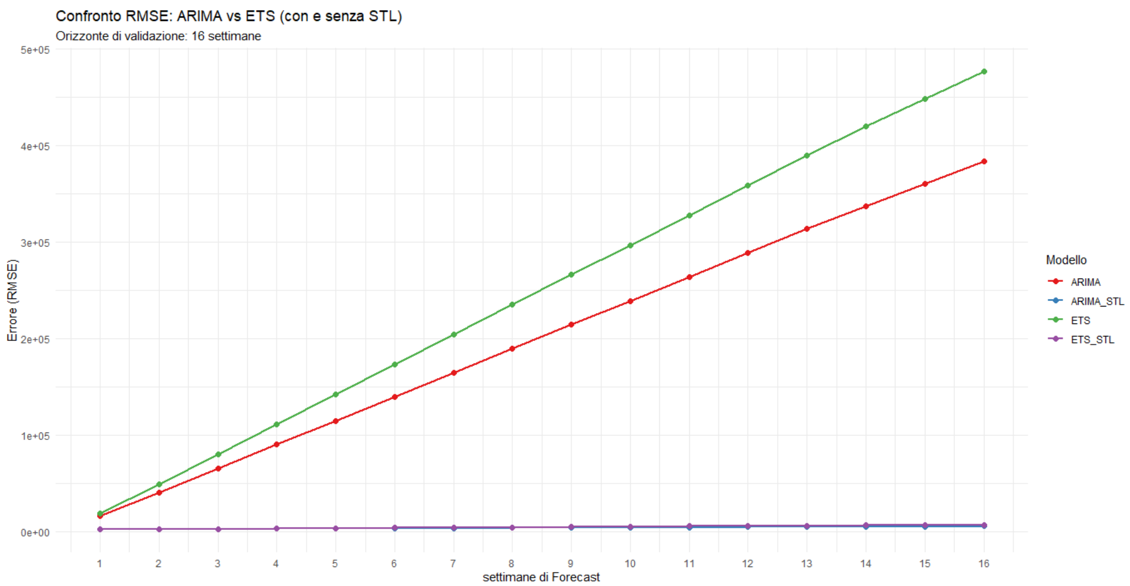

Figure 8.13: Italian load: RMSE comparison across forecast horizons

The figure confirms that models combined with STL decomposition remain clearly superior over longer horizons. A more focused comparison between the two best STL based models is shown in Figure 8.14. The graph compares ARIMA + STL and ETS + STL over the same range of forecast horizons.

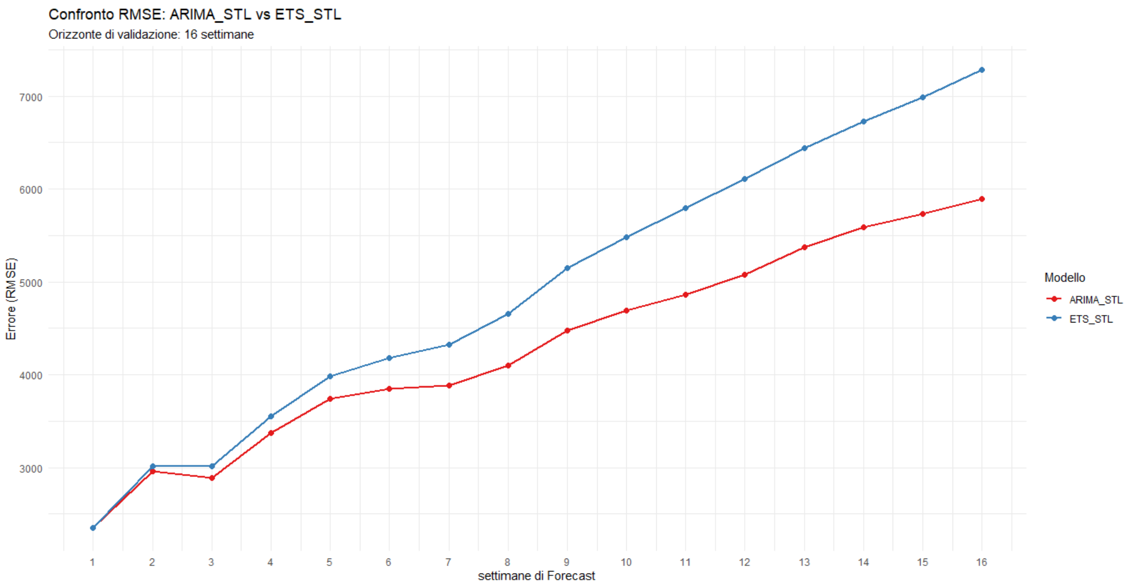

Figure 8.14: Italian load: RMSE comparison of STL based models

Although the initial RMSE values reported in Table 8.1 are very close, the horizon based comparison suggests that ARIMA + STL performs slightly better over longer horizons. For this reason, ARIMA combined with STL decomposition can be considered the most reliable specification for medium term national load forecasting in this analysis.

Short term forecasts are essential for generation scheduling, balancing operations, and reserve management, as explained in [17, 9]. The importance of evaluating load forecasting methods on realistic energy datasets is also highlighted by the Global Energy Forecasting Competition [10]. The importance of evaluating load forecasting methods on realistic energy datasets is also highlighted by the Global Energy Forecasting Competition [10]. In particular, accurate demand predictions help the system operator maintain the balance between generation and consumption and support the secure operation of the grid.

For this reason, the best performing STL based models were also compared over shorter forecasting horizons.

Figure 8.15 reports the RMSE for horizons from one to seven days. The two models show very similar performance, confirming that both ARIMA + STL and ETS + STL are effective in short term load forecasting.

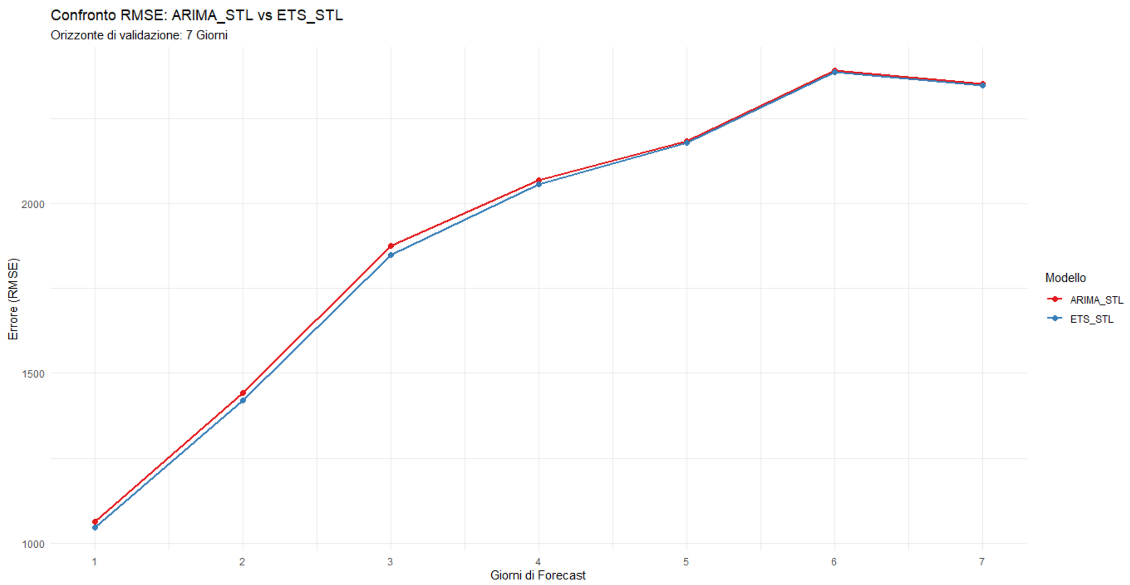

Figure 8.15: Italian load: RMSE comparison for horizons from one to seven days

Figure 8.16 further reduces the forecasting horizon and compares the models from one to twenty four hours ahead.

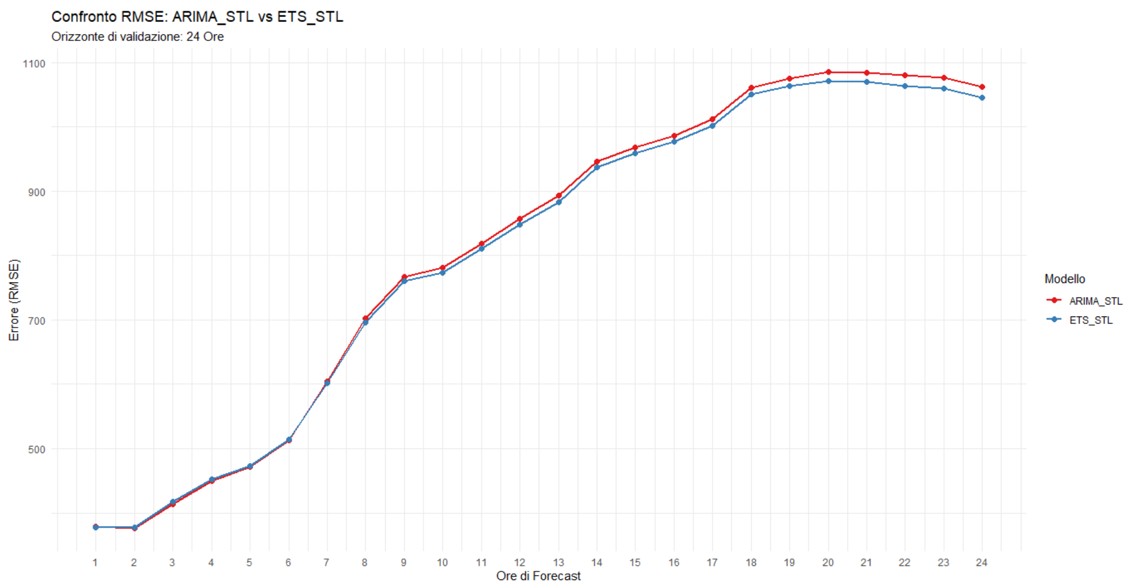

Figure 8.16: Italian load: RMSE comparison for horizons from one to twenty four hours

The short horizon results show that the two models remain very close. However, ETS + STL appears to be slightly more accurate for very short term forecasts.

The ARIMA specifications selected during cross validation are reported in Table 8.2. The results show that the selected models generally include first order differencing, confirming the need to stabilize the mean of the series before modelling the remaining autocorrelation structure. The most frequently selected models are ARIMA$(1, 1, 5)$ and ARIMA$(5, 1, 1)$.

Table 8.2: Italian load: Selected ARIMA specifications

Model Specification	**Times Selected**
ARIMA(1,1,5)	3
ARIMA(5,1,1)	3
ARIMA(2,1,4)	2
ARIMA(1,1,4)	1
ARIMA(3,1,3)	1
ARIMA(4,1,1)	1

8.7 National Prediction Intervals

To evaluate not only the point forecasts but also the uncertainty associated with them, prediction intervals were analysed using a classical training and test split. The training set contains the observations from 2021 to 2024, while the remaining observations are used as the test set.

The selected model was fitted on the training set and used to generate forecasts over the following period. Figure 8.17 shows the prediction interval for the forecast horizon considered.

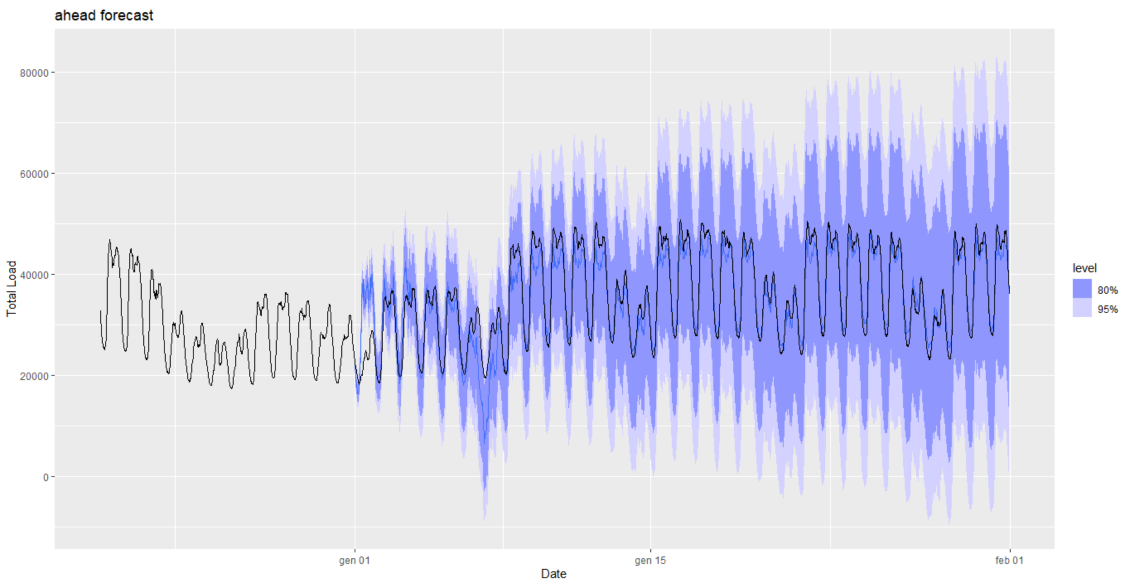

Figure 8.17: Italian load: Prediction interval

The prediction interval captures the main seasonal structure of the load series, indicating that the model is able to reproduce the regular temporal behaviour of electricity demand. At the same time, the interval becomes wider as the forecast horizon increases. This is expected, since uncertainty accumulates over time and the model has to predict increasingly distant observations.

8.8 Sardinia and Sicily Forecasting Results

The forecasting performance for Sardinia and Sicily was first evaluated over horizons ranging from the next hour to 192 hours, corresponding to 8 days. For each horizon, the RMSE was computed in order to analyse how prediction accuracy changes as the forecast moves further away from the estimation point.

As expected, the RMSE increases with the forecasting horizon for both ARIMA + STL and ETS + STL. This reflects the natural accumulation of uncertainty over time. However, the two models do not behave in exactly the same way. The comparison shows that ARIMA + STL provides more stable performance over the longer horizon and generally achieves lower errors than ETS + STL.

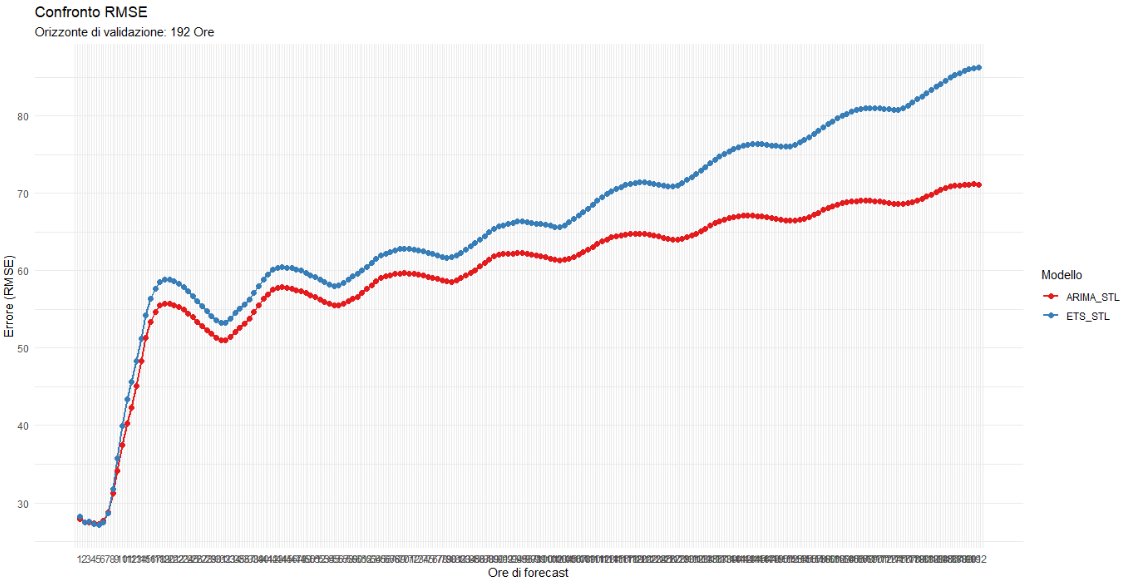

Figure 8.18: Sardinia load: RMSE comparison over an 8 day horizon

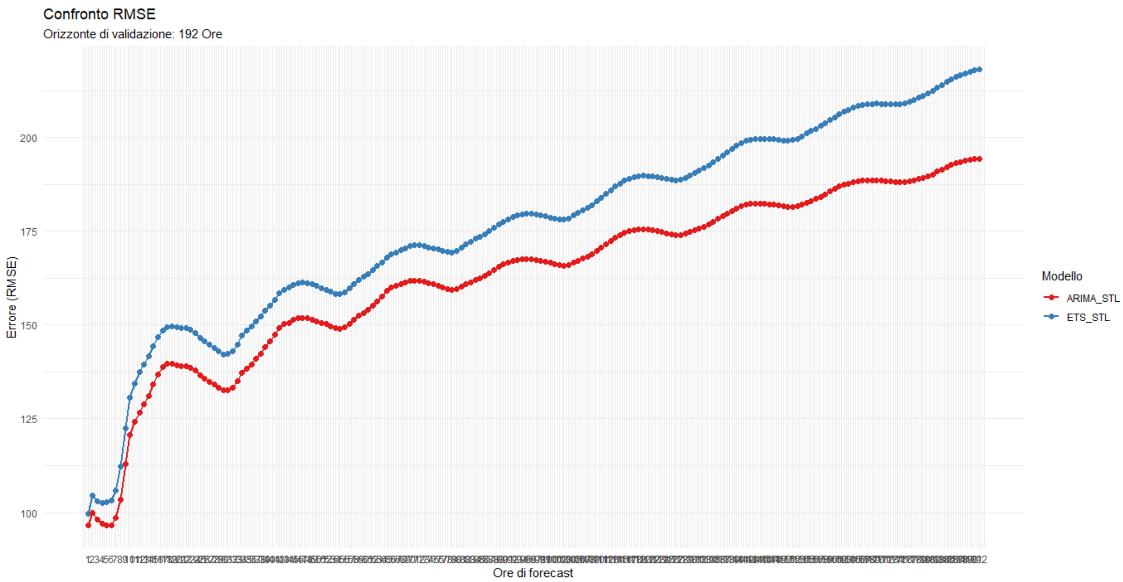

Figure 8.19: Sicily load: RMSE comparison over an 8 day horizon

The results suggest that ARIMA + STL is better suited for the medium term horizon considered in the regional analysis. After the seasonal components are removed, the ARIMA structure appears able to capture the remaining autocorrelation of the seasonally adjusted load series more effectively.

The same comparison was then restricted to the first 24 forecast hours. This shorter horizon is particularly relevant for operational applications, since it is closely connected to day ahead scheduling, reserve planning, and short term balancing.

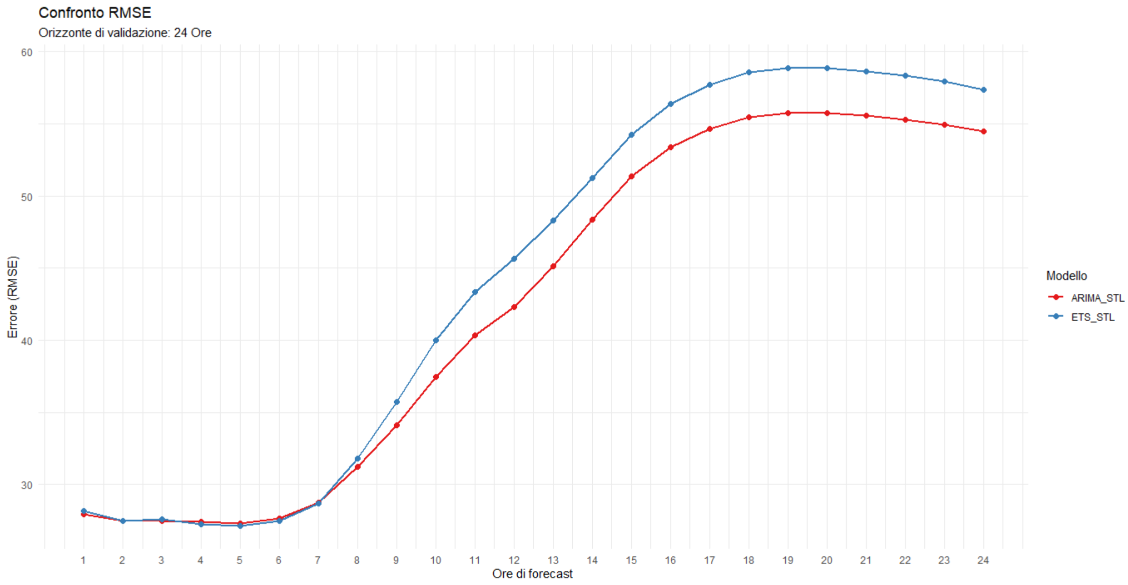

Figure 8.20: Sardinia load: RMSE comparison over a 24 hour horizon

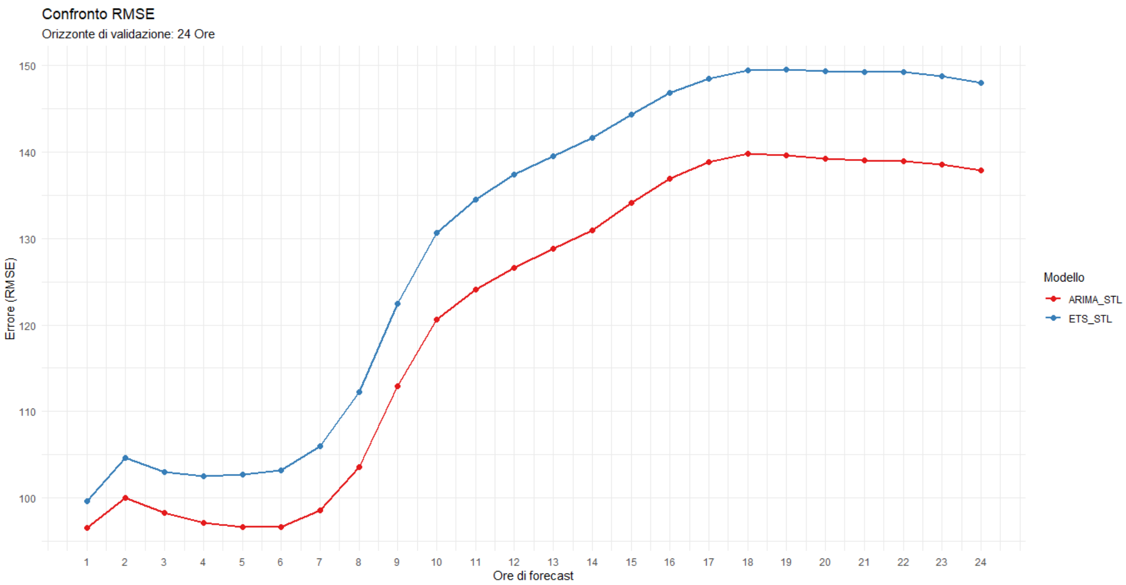

Figure 8.21: Sicily load: RMSE comparison over a 24 hour horizon

Even at the shorter horizon, ARIMA + STL tends to achieve lower RMSE values than ETS + STL. These findings indicate that ARIMA + STL is the preferred model for both Sardinia and Sicily in the considered forecasting setting. The ARIMA specifications selected along cross validation are reported in Table 8.3. In both islands, all selected ARIMA models include first order differencing. This confirms that differencing is required to stabilize the mean of the seasonally adjusted series before modelling the remaining autocorrelation structure.

For Sardinia, the most frequently selected specifications are ARIMA$(3, 1, 3)$ and ARIMA$(2, 1, 3)$. For Sicily, the most frequently selected model is ARIMA$(2, 1, 4)$, followed by ARIMA$(2, 1, 2)$ and ARIMA$(3, 1, 3)$. This indicates that the two island series require slightly different autoregressive and moving average structures.

Table 8.3: Sardinia and Sicily load: ARIMA specifications selected for ARIMA + STL

Area	**Model Specification**	**Times Selected**
Sardinia	ARIMA(3,1,3)	59
Sardinia	ARIMA(2,1,3)	41
Sardinia	ARIMA(2,1,2)	22
Sardinia	ARIMA(2,1,4)	19
Sardinia	ARIMA(2,1,4) with drift	18
Sardinia	ARIMA(2,1,3) with drift	16
Sardinia	ARIMA(3,1,3) with drift	9
Sardinia	ARIMA(1,1,5)	8
Sardinia	ARIMA(4,1,2)	6
Sardinia	ARIMA(1,1,3)	1
Sardinia	ARIMA(1,1,4)	1
Sardinia	ARIMA(5,1,1)	1
Sicily	ARIMA(2,1,4)	72
Sicily	ARIMA(2,1,2)	45
Sicily	ARIMA(3,1,3)	35
Sicily	ARIMA(2,1,3)	28
Sicily	ARIMA(2,1,4) with drift	18
Sicily	ARIMA(2,1,3) with drift	1
Sicily	ARIMA(3,1,3) with drift	1
Sicily	ARIMA(4,1,2) with drift	1

Table 8.4 reports the ETS specifications selected during cross validation. In both islands, the most frequently selected model is ETS(A, N, N), which corresponds to additive errors, no trend, and no seasonal component in the seasonally adjusted series.

This result is coherent with the use of STL decomposition. Since the main seasonal patterns are removed before the ETS model is estimated, the residual series often does not require an explicit seasonal component. In most folds, a simple local level structure is sufficient to model the seasonally adjusted load.

Table 8.4: Sardinia and Sicily load: ETS specifications selected for ETS + STL

Area	**Model Specification**	**Times Selected**
Sardinia	ETS(A,N,N)	183
Sardinia	ETS(M,N,N)	18
Sicily	ETS(A,N,N)	190
Sicily	ETS(M,N,N)	8
Sicily	ETS(A,Ad,N)	3

8.9 Transition to the Hour Specific Analysis

The results presented in this chapter show that modelling the full fifteen minute load series is feasible, but also challenging. The complete time series contains 96 observations per day, and the model must reproduce the entire intra day structure within a single specification. Even when STL decomposition is used to remove the main seasonal components, the residual dynamics remain complex.

This motivates the analysis developed in the following chapter, inspired from the approach used for prices. Instead of modelling the entire fifteen minute series directly, the next step consists of splitting each load series into 96 separate time series, one for each fifteen minute interval of the day. The purpose of the next chapter is to evaluate whether this time specific modelling strategy improves forecasting performance with respect to the full series approach analysed here.

Chapter 9

Time Specific Load Forecasting Results

9.1 Introduction

The previous chapter analysed electricity load forecasting for Italy, Sardinia, and Sicily by modelling each area as a complete time series. This approach made it possible to evaluate the general forecasting performance of classical statistical models on the full load profile. However, electricity demand has a very strong intra day structure.

For this reason, as suggested by [3], [15], which analyzes Brazilian load, on 24 series and, and by [6], which does the same analysis for French load, this chapter introduces a more detailed forecasting strategy. Instead of modelling the complete series as a single object, the load data are divided into 96 separate daily time series, one for each fifteen minute interval of the day. Each of these series contains the observations corresponding to the same time slot across consecutive days.

This procedure is applied separately to the Italian, Sardinian, and Sicilian load datasets. The aim is to verify whether decomposing the daily load profile into homogeneous fifteen minute components improves forecasting accuracy with respect to the whole series approach, and see which model performs better in this settings. This chapter therefore follows the same logic used in the European price analysis, In the present case, however, the decomposition produces 96 interval specific series instead of 24.

9.2 Construction of the Interval Specific Series

Let $L_{d,s}^{(r)}$ denote the electricity load observed in region r on day d and fifteen minute slot s, where $s = 0, \ldots, 95$. The index s identifies the position of the observation within the day. Therefore, $s = 0$ corresponds to 00:00, $s = 1$ to 00:15, and $s = 95$ to 23:45.

For each region, the original dataset was divided according to the variable `SlotQuarto`. This produces 96 daily time series for each area. The resulting structure is conceptually different from the whole series approach used in the previous chapter. In the complete fifteen minute series, the model has to learn both the evolution over time and the systematic variation between different moments of the day. In the interval specific approach, each model is estimated only on observations referring to the same fifteen minute slot. This removes most of the deterministic intra day variation from the modelling problem.

For the ARIMA and ETS models, the target variable is defined as the daily first difference of load within the same slot:

$$\Delta L^{(r)}_{d,s} = L^{(r)}_{d,s} - L^{(r)}_{d-1,s}.$$

This transformation is obtained after filtering the observations for a fixed slot and ordering them by date. The first missing value generated by the lag operator is removed.

For the GARCH model, a weekly difference is used instead:

$$\Delta_7 L^{(r)}_{d,s} = L^{(r)}_{d,s} - L^{(r)}_{d-7,s}.$$

This choice was made trying to reduce weekly dependence, since garch doesn't include seasonal components. The errore is stil evaluated on the original series, which makes the models comparable.

9.3 Model Specifications and Cross Validation

The forecasting experiment compares three statistical models: ARIMA, ETS, and GARCH. The same modelling structure is applied to Italy, Sardinia, and Sicily, and each model is estimated independently for all 96 fifteen minute slots.

The validation period starts on 1 January 2023. For each interval specific series, the initial training set contains all observations before this date. Forecasts are then generated through a rolling or expanding validation procedure, depending on the model.

9.3.1 ARIMA Model

The ARIMA model is estimated on the daily first difference of load for each fifteen minute slot. The specification used in the experiment is:

$$ARIMA(0,1,1)(1,0,0)_7.$$

The non seasonal part contains a moving average component of order one, while the seasonal part includes an autoregressive component with weekly period equal to seven. Since each slot specific series has daily frequency, the seasonal period captures weekly dependence in the same fifteen minute interval.

The cross validation procedure is implemented using an expanding window. The initial training sample includes all observations before 1 January 2023. At each iteration, the model is re estimated after adding eight new observations, and an eight day forecast horizon is produced. The forecasted values are then matched with the corresponding observed values in order to compute the forecast errors.

9.3.2 ETS Model

The ETS model is also estimated on the daily first difference of load for each fifteen minute slot. The selected specification is:

$$ETS(A,A,A)_7.$$

This model includes additive error, additive trend, and additive weekly seasonality. The weekly seasonal component is again defined with period seven, reflecting the daily frequency of each interval specific series.

The cross validation procedure is the same as the one used for ARIMA. The initial training set contains all observations before 1 January 2023, the training window expands over time, and the model is refitted every eight observations, as for ARIMA. At each validation step, an eight day forecast horizon is generated.

9.3.3 GARCH Model

The GARCH model is used to capture time varying volatility in the differenced load series. Unlike ARIMA and ETS, the GARCH specification is estimated on the weekly difference of load. The mean equation follows an ARMA$(0, 1)$ structure, while the conditional variance is modelled with a standard GARCH$(1, 1)$ process. The innovations are assumed to follow a Student t distribution. Rolling validation is performed, The model produces one step ahead forecasts, uses an expanding window, and is refitted every 90 observations.

Table 9.1: Interval specific load forecasting: Model specifications

Model	**Target variable**	**Specification**	**Validation setting**
ARIMA	Daily load difference	ARIMA$(0, 0, 1)(1, 0, 0)_7$	8 day expanding window
ETS	Daily load difference	ETS$(A, A, A)_7$	8 day expanding window
GARCH	Weekly load difference	ARMA$(0, 1)$ + sGARCH$(1, 1)$	One step ahead rolling forecast

9.4 Comparison with the Whole Series Approach

The first objective of this chapter is to verify whether the interval specific approach improves the forecasting results obtained in the previous chapter. For this reason, the best model estimated on the whole Italian load series is compared with the best model obtained after splitting the data into 96 interval specific series.

The comparison is reported in Table 9.2. The best whole series model is ARIMA combined with STL decomposition, with a MAE equal to 3372.592. In contrast, the best interval specific model is ETS, with a MAE equal to 1278.000, on the same forecast horizon.

The reduction in MAE is substantial. This confirms that the separation of the fifteen minute slots simplifies the forecasting problem.

This result is important because it shows that the gain is not only methodological but also empirical. Dividing the Italian load series into 96 homogeneous daily series leads to a much more accurate forecast than modelling the complete fifteen minute sequence directly.

Table 9.2: Italian load: Best model comparison between whole series and interval specific approach

Approach	**Best Model**	**MAE**
Whole time series	ARIMA + STL	3372.592
96 interval specific time series	ETS	1278.000

9.5 Average Results Across Regions

The interval specific strategy was then applied to Sardinia, Sicily, and Italy. For each region and each model, the errors were computed for the slot specific forecasts and then summarized through average MAE and MAPE values. The resulting comparison is shown in Table 9.3.

The results show a clear and consistent pattern. Among the statistical models, ETS provides the best average performance in all three areas.

ARIMA is the second best statistical model in all three areas, while GARCH produces the largest errors.

TERNA's forecasts are also included as an external benchmark. TERNA achieves the lowest MAE and MAPE in all three areas. This result is expected, since operational forecasts are likely to rely on richer information sets.

Nevertheless, the comparison remains informative. The statistical models, especially ETS, are able to reproduce a relevant part of the load dynamics using only historical load observations.

Table 9.3: Load forecasting: Average MAE and MAPE for interval specific models and TERNA forecasts

Area	**Model**	**MAE**	**MAPE**
Sardinia	ARIMA	41.9	4.27
	GARCH	50.8	5.17
	ETS	36.5	3.73
	TERNA	19.4	1.97
Sicily	ARIMA	117.0	5.33
	GARCH	132.0	5.97
	ETS	98.2	4.51
	TERNA	43.1	1.97
Italy	ARIMA	1682.0	4.85
	GARCH	1825.0	5.24
	ETS	1278.0	3.80
	TERNA	697.0	1.98

9.6 Slot Level Error Analysis

The average results summarize the overall performance of each model, but they do not show how forecast accuracy changes within the day. For this reason, the errors were also analysed separately for each fifteen minute slot. The complete slot level results are reported in Tables 9.4, 9.5, and 9.6.

In addition to the numerical tables, the MAE profiles are plotted across the day for each region. The graphs report one line for each model, allowing the intra day variation of the forecasting error to be inspected visually. These figures make clear that the errors are generally lower during the night and higher during the daytime hours, especially during the morning increase and the central part of the day. The graphs also demonstrate how ETS isn't just the best model on average, but it seems the best model for each time series.

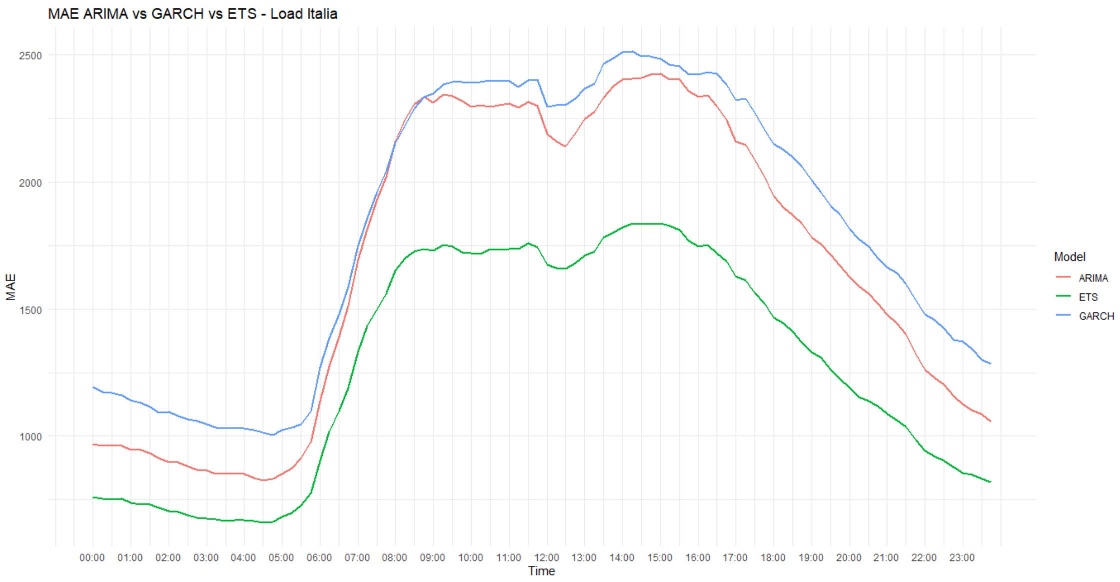

Figure 9.1: MAE Italy

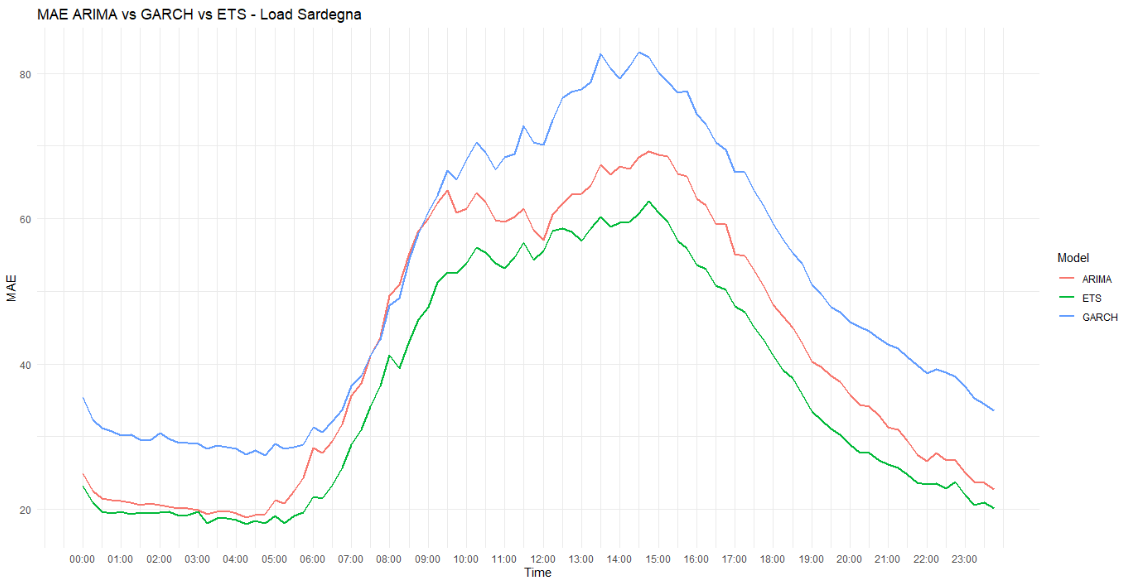

Figure 9.2: MAE Sardinia

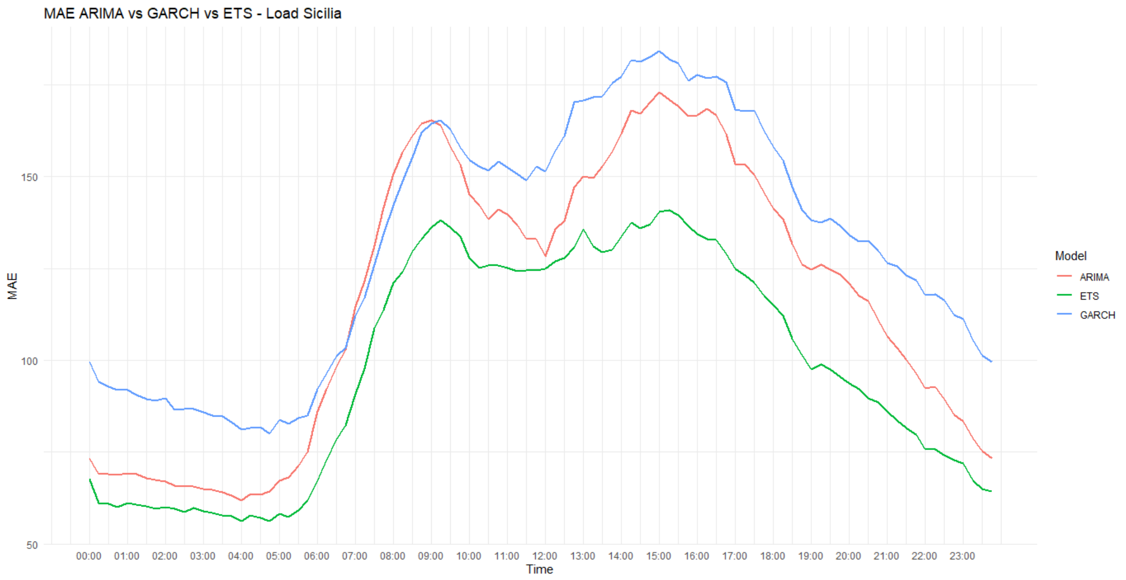

Figure 9.3: MAE Sicily

Figure 9.4 compares the observed Italian load values with the corresponding ETS forecasts during 2023. The plot provides a direct visual assessment of the best statistical model. ETS is able to reproduce the general trajectory of the load and the main recurrent patterns.

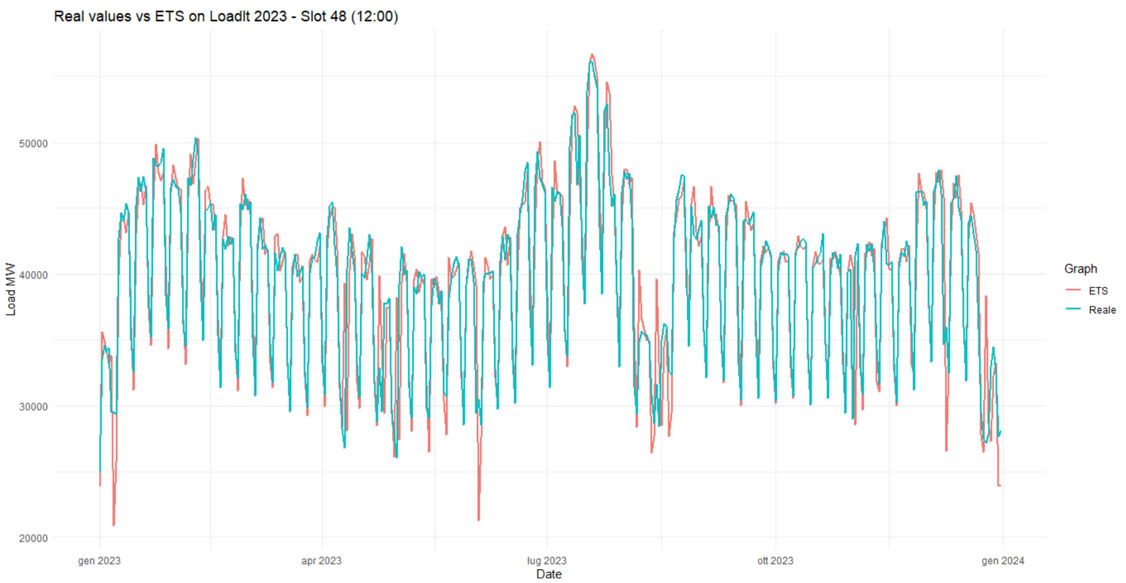

Figure 9.4: Real vs ETS It

Figure 9.5 compares actual and predicted values for the Sardinian ETS forecasts in 2023. The figure shows that the model follows the main evolution of the regional load.

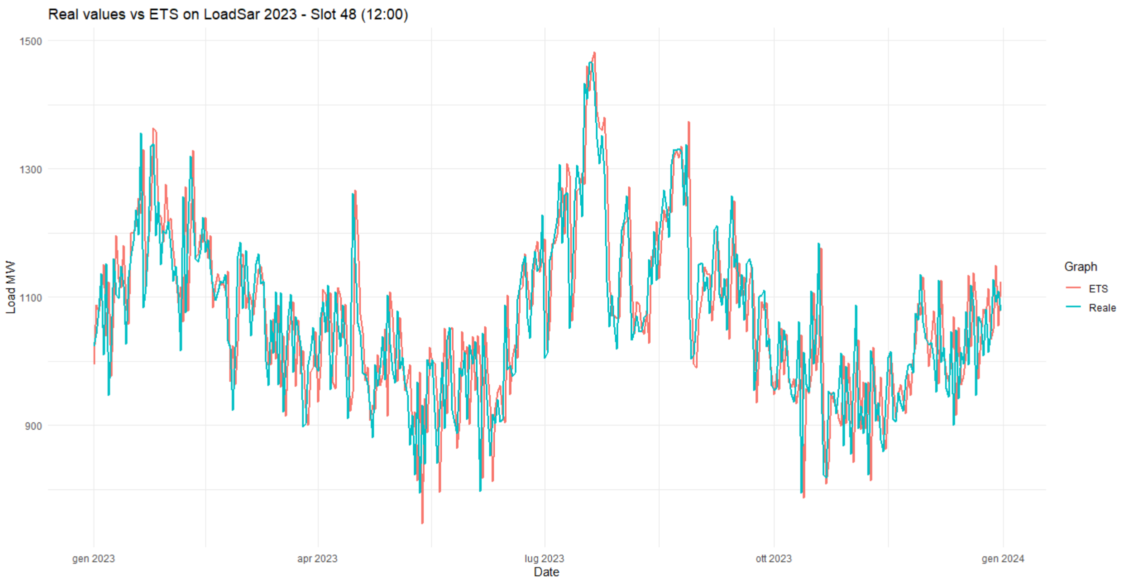

Figure 9.5: Real vs ETS Sar

Figure 9.6 compares the observed Sicilian load values with the ETS forecasts in 2023. The plot highlights that ETS captures the main movements of the series.

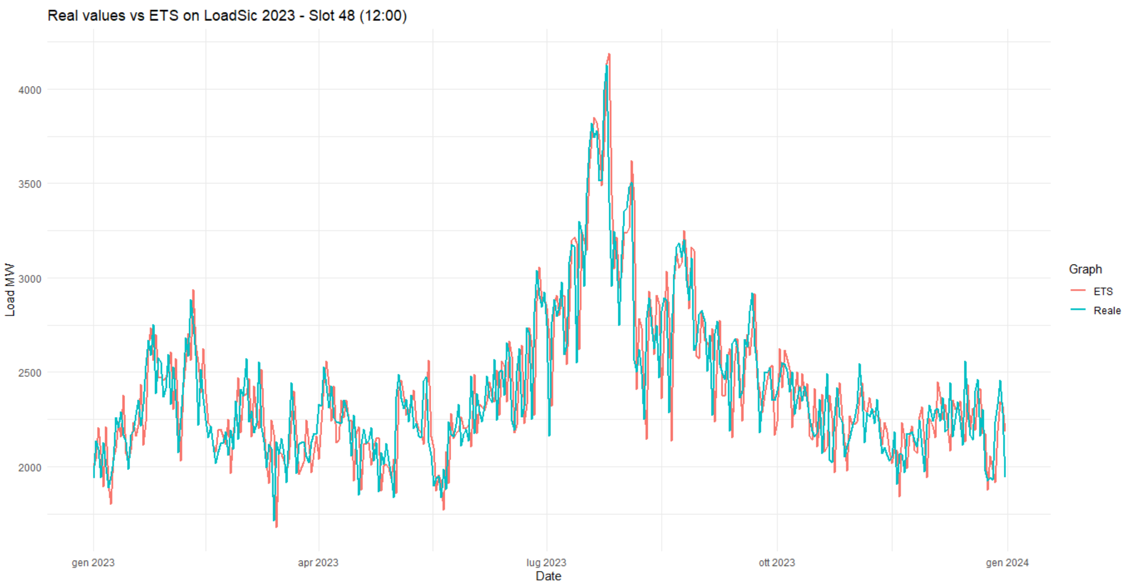

Figure 9.6: Real vs ETS Sic

The slot level results also show that the difference between the statistical models and TERNA is larger during the most variable parts of the day. This reinforces the interpretation that purely univariate models can capture regular temporal dynamics, but they are limited when the load is driven by exogenous factors.

9.7 Discussion

The interval specific analysis provides three main findings.

First, dividing the fifteen minute load series into 96 daily series improves forecasting performance. The comparison for Italy shows that the best interval specific model, ETS, substantially outperforms the best model estimated on the complete time series. This confirms that the intra day structure of electricity load is not only visually evident but also relevant for model accuracy.

Second, ETS is the most effective statistical model in this setting. Its additive error, trend, and weekly seasonal components provide a simple but robust structure for modelling the daily evolution of each time slot. This result differs from the whole series analysis, where ARIMA combined with STL decomposition was the best performing model.

Third, TERNA forecasts remain more accurate than all statistical models. This does not invalidate the interval specific approach. Rather, it shows that historical load alone is not sufficient to fully reproduce operational forecasting performance. The gap suggests that future improvements should include explanatory variables, or furtherly splitting the series.

The graphical analysis supports the numerical evidence. The MAE profiles show that the forecasting error is systematically higher during the daytime hours than during the night for all three regions. This behaviour is consistent with the fact that daytime load is more affected by production activity, temperature conditions, and changes in consumption patterns.

9.8 Appendix

Table 9.4: Italy load: MAE and MAPE by fifteen minute slot

Slot	Time	MAE A.	MAE G.	MAE E.	MAE T.	MAPE A.	MAPE G.	MAPE E.	MAPE T.
0	00:00	969.0	1194.0	759.0	663.0	3.37	4.09	2.68	2.26
1	00:15	962.0	1173.0	754.0	620.0	3.42	4.10	2.72	2.17
2	00:30	965.0	1172.0	750.0	649.0	3.49	4.18	2.75	2.32
3	00:45	963.0	1161.0	754.0	650.0	3.55	4.22	2.83	2.37
4	01:00	945.0	1142.0	738.0	620.0	3.54	4.21	2.81	2.28
5	01:15	946.0	1132.0	732.0	597.0	3.59	4.23	2.83	2.23
6	01:30	933.0	1116.0	732.0	570.0	3.58	4.23	2.86	2.16
7	01:45	913.0	1093.0	720.0	541.0	3.55	4.19	2.84	2.08
8	02:00	898.0	1094.0	705.0	509.0	3.51	4.22	2.80	1.96
9	02:15	896.0	1078.0	704.0	505.0	3.53	4.20	2.82	1.97
10	02:30	883.0	1066.0	688.0	491.0	3.49	4.18	2.77	1.93
11	02:45	869.0	1061.0	678.0	487.0	3.46	4.19	2.75	1.92
12	03:00	866.0	1047.0	678.0	453.0	3.45	4.14	2.75	1.79
13	03:15	853.0	1034.0	674.0	450.0	3.41	4.10	2.74	1.78
14	03:30	852.0	1030.0	668.0	451.0	3.42	4.10	2.72	1.79
15	03:45	855.0	1033.0	670.0	450.0	3.44	4.12	2.74	1.79
16	04:00	852.0	1030.0	668.0	407.0	3.41	4.09	2.72	1.61
17	04:15	834.0	1023.0	665.0	413.0	3.34	4.07	2.70	1.63
18	04:30	826.0	1015.0	659.0	408.0	3.30	4.03	2.67	1.60
19	04:45	831.0	1004.0	664.0	410.0	3.30	3.96	2.68	1.60
20	05:00	851.0	1025.0	684.0	398.0	3.32	3.97	2.71	1.51
21	05:15	874.0	1032.0	698.0	392.0	3.38	3.96	2.75	1.47
22	05:30	914.0	1046.0	729.0	413.0	3.50	3.97	2.84	1.53
23	05:45	980.0	1099.0	778.0	451.0	3.69	4.10	3.00	1.65
24	06:00	1140.0	1271.0	903.0	631.0	4.10	4.52	3.35	2.11
25	06:15	1275.0	1386.0	1019.0	678.0	4.48	4.81	3.69	2.18
26	06:30	1393.0	1481.0	1100.0	626.0	4.78	5.02	3.90	1.98
27	06:45	1511.0	1586.0	1190.0	596.0	5.05	5.24	4.12	1.84
28	07:00	1695.0	1751.0	1334.0	810.0	5.43	5.55	4.44	2.32
29	07:15	1815.0	1858.0	1433.0	768.0	5.66	5.73	4.64	2.15
30	07:30	1931.0	1959.0	1499.0	732.0	5.87	5.89	4.75	2.01
31	07:45	2018.0	2041.0	1559.0	665.0	5.99	6.00	4.83	1.81
32	08:00	2160.0	2154.0	1653.0	802.0	6.21	6.12	4.95	2.08
33	08:15	2248.0	2225.0	1699.0	752.0	6.30	6.17	4.96	1.93
34	08:30	2305.0	2291.0	1728.0	712.0	6.33	6.22	4.95	1.82
35	08:45	2335.0	2336.0	1736.0	684.0	6.31	6.24	4.88	1.73
36	09:00	2312.0	2349.0	1730.0	684.0	6.15	6.18	4.78	1.72
37	09:15	2346.0	2385.0	1754.0	709.0	6.17	6.20	4.78	1.77
38	09:30	2339.0	2393.0	1746.0	729.0	6.12	6.17	4.72	1.80
39	09:45	2318.0	2394.0	1722.0	756.0	6.03	6.14	4.64	1.87
40	10:00	2296.0	2389.0	1719.0	778.0	5.95	6.09	4.60	1.92
41	10:15	2301.0	2392.0	1718.0	811.0	5.94	6.08	4.58	2.00
42	10:30	2295.0	2400.0	1737.0	845.0	5.89	6.07	4.61	2.08
43	10:45	2303.0	2395.0	1733.0	887.0	5.90	6.06	4.59	2.19
44	11:00	2309.0	2397.0	1737.0	903.0	5.91	6.05	4.58	2.22
45	11:15	2293.0	2375.0	1736.0	949.0	5.84	5.98	4.56	2.33
46	11:30	2314.0	2399.0	1759.0	939.0	5.88	6.03	4.61	2.31
47	11:45	2298.0	2404.0	1743.0	940.0	5.84	6.04	4.56	2.32
48	12:00	2189.0	2295.0	1676.0	1015.0	5.61	5.83	4.40	2.52
49	12:15	2160.0	2303.0	1662.0	1020.0	5.56	5.87	4.38	2.56

Slot	Time	MAE A.	MAE G.	MAE E.	MAE T.	MAPE A.	MAPE G.	MAPE E.	MAPE T.
50	12:30	2139.0	2301.0	1658.0	1018.0	5.53	5.90	4.39	2.57
51	12:45	2191.0	2329.0	1682.0	1001.0	5.71	6.01	4.50	2.55
52	13:00	2245.0	2366.0	1710.0	943.0	5.88	6.13	4.60	2.42
53	13:15	2275.0	2388.0	1726.0	961.0	5.99	6.23	4.68	2.49
54	13:30	2332.0	2464.0	1783.0	914.0	6.17	6.44	4.86	2.37
55	13:45	2377.0	2486.0	1800.0	912.0	6.34	6.55	4.94	2.37
56	14:00	2401.0	2511.0	1821.0	855.0	6.40	6.62	5.00	2.21
57	14:15	2406.0	2515.0	1836.0	852.0	6.42	6.64	5.07	2.20
58	14:30	2410.0	2495.0	1834.0	843.0	6.44	6.58	5.07	2.17
59	14:45	2421.0	2495.0	1835.0	841.0	6.48	6.61	5.09	2.16
60	15:00	2425.0	2485.0	1838.0	822.0	6.51	6.59	5.10	2.11
61	15:15	2401.0	2460.0	1826.0	837.0	6.46	6.54	5.09	2.15
62	15:30	2404.0	2454.0	1810.0	810.0	6.48	6.53	5.06	2.08
63	15:45	2356.0	2424.0	1770.0	786.0	6.34	6.45	4.94	2.02
64	16:00	2335.0	2421.0	1746.0	741.0	6.30	6.45	4.88	1.91
65	16:15	2340.0	2432.0	1752.0	736.0	6.31	6.48	4.88	1.90
66	16:30	2299.0	2426.0	1718.0	731.0	6.19	6.45	4.77	1.89
67	16:45	2243.0	2382.0	1688.0	744.0	6.01	6.30	4.65	1.91
68	17:00	2159.0	2321.0	1628.0	788.0	5.77	6.13	4.46	2.02
69	17:15	2145.0	2327.0	1611.0	768.0	5.70	6.13	4.39	1.96
70	17:30	2086.0	2273.0	1563.0	763.0	5.54	5.99	4.24	1.94
71	17:45	2018.0	2204.0	1521.0	732.0	5.35	5.78	4.12	1.86
72	18:00	1945.0	2151.0	1467.0	740.0	5.15	5.65	3.96	1.87
73	18:15	1900.0	2125.0	1443.0	715.0	5.01	5.57	3.88	1.80
74	18:30	1870.0	2097.0	1411.0	720.0	4.90	5.49	3.78	1.81
75	18:45	1838.0	2062.0	1370.0	731.0	4.78	5.36	3.64	1.82
76	19:00	1783.0	2005.0	1330.0	724.0	4.62	5.20	3.52	1.80
77	19:15	1757.0	1962.0	1312.0	690.0	4.54	5.07	3.46	1.71
78	19:30	1712.0	1906.0	1262.0	656.0	4.42	4.92	3.33	1.62
79	19:45	1672.0	1871.0	1225.0	632.0	4.32	4.83	3.23	1.56
80	20:00	1627.0	1816.0	1189.0	644.0	4.21	4.70	3.15	1.60
81	20:15	1589.0	1774.0	1153.0	637.0	4.14	4.62	3.08	1.59
82	20:30	1559.0	1746.0	1138.0	623.0	4.10	4.58	3.06	1.58
83	20:45	1521.0	1699.0	1118.0	619.0	4.04	4.50	3.04	1.58
84	21:00	1479.0	1664.0	1090.0	646.0	3.99	4.47	3.00	1.68
85	21:15	1445.0	1642.0	1063.0	640.0	3.96	4.47	2.97	1.69
86	21:30	1402.0	1600.0	1038.0	642.0	3.91	4.43	2.95	1.73
87	21:45	1325.0	1534.0	984.0	720.0	3.78	4.34	2.86	1.99
88	22:00	1261.0	1480.0	942.0	767.0	3.67	4.28	2.79	2.16
89	22:15	1230.0	1455.0	922.0	702.0	3.66	4.29	2.79	2.03
90	22:30	1202.0	1424.0	904.0	665.0	3.65	4.29	2.79	1.97
91	22:45	1159.0	1379.0	878.0	653.0	3.60	4.24	2.77	1.98
92	23:00	1124.0	1372.0	857.0	639.0	3.56	4.31	2.76	1.98
93	23:15	1104.0	1342.0	848.0	627.0	3.58	4.31	2.79	1.99
94	23:30	1085.0	1302.0	834.0	646.0	3.60	4.28	2.82	2.10
95	23:45	1058.0	1285.0	821.0	647.0	3.60	4.33	2.84	2.16

Table 9.5: Sardinia load: MAE and MAPE by fifteen minute slot

Slot	**Time**	**MAE A.**	**MAE G.**	**MAE E.**	**MAE T.**	**MAPE A.**	**MAPE G.**	**MAPE E.**	**MAPE T.**
0	00:00	24.9	35.5	23.2	19.4	3.01	4.19	2.80	2.24
1	00:15	22.5	32.3	20.9	18.3	2.73	3.84	2.52	2.16
2	00:30	21.5	31.2	19.7	19.1	2.66	3.77	2.41	2.30
3	00:45	21.3	30.8	19.5	19.2	2.69	3.81	2.44	2.34
4	01:00	21.1	30.2	19.7	18.3	2.69	3.78	2.50	2.27
5	01:15	21.0	30.2	19.3	17.7	2.70	3.82	2.47	2.22
6	01:30	20.6	29.6	19.6	16.9	2.68	3.79	2.54	2.14
7	01:45	20.8	29.5	19.4	16.1	2.74	3.81	2.55	2.06
8	02:00	20.6	30.5	19.6	15.1	2.72	3.98	2.59	1.95
9	02:15	20.3	29.7	19.7	15.0	2.71	3.90	2.63	1.95
10	02:30	20.1	29.1	19.1	14.6	2.69	3.84	2.55	1.91
11	02:45	20.2	29.1	19.2	14.5	2.71	3.86	2.58	1.91
12	03:00	19.9	29.0	19.7	13.6	2.69	3.86	2.66	1.78
13	03:15	19.3	28.4	18.1	13.5	2.62	3.80	2.45	1.77
14	03:30	19.7	28.7	18.8	13.4	2.67	3.84	2.54	1.78
15	03:45	19.7	28.6	18.7	13.5	2.68	3.83	2.54	1.78
16	04:00	19.5	28.4	18.5	12.2	2.64	3.80	2.51	1.60
17	04:15	18.9	27.5	17.9	12.3	2.55	3.67	2.41	1.62
18	04:30	19.2	28.1	18.4	12.2	2.59	3.75	2.48	1.59
19	04:45	19.3	27.4	18.1	12.2	2.59	3.64	2.42	1.59
20	05:00	21.3	29.1	19.1	11.8	2.81	3.80	2.51	1.50
21	05:15	20.8	28.4	18.1	11.6	2.74	3.69	2.37	1.47
22	05:30	22.5	28.6	19.1	12.1	2.92	3.68	2.47	1.52
23	05:45	24.3	28.9	19.6	13.2	3.11	3.66	2.50	1.63
24	06:00	28.4	31.3	21.7	18.0	3.47	3.79	2.66	2.09
25	06:15	27.8	30.7	21.5	18.7	3.38	3.70	2.62	2.16
26	06:30	29.3	32.1	23.4	17.3	3.51	3.82	2.80	1.96
27	06:45	31.7	33.7	25.7	16.5	3.69	3.90	3.01	1.83
28	07:00	35.6	37.0	28.9	22.0	3.99	4.11	3.25	2.31
29	07:15	37.3	38.4	31.0	20.8	4.08	4.17	3.41	2.14
30	07:30	41.1	41.2	34.1	20.0	4.39	4.36	3.67	2.01
31	07:45	43.8	43.4	37.0	18.4	4.58	4.50	3.89	1.81
32	08:00	49.4	48.1	41.2	21.9	4.99	4.82	4.18	2.07
33	08:15	50.9	49.0	39.4	20.5	5.03	4.81	3.89	1.93
34	08:30	55.3	54.4	43.1	19.5	5.37	5.25	4.20	1.82
35	08:45	58.3	58.0	46.1	18.9	5.58	5.52	4.43	1.74
36	09:00	60.0	60.9	47.8	19.0	5.66	5.70	4.51	1.73
37	09:15	62.2	63.2	51.2	19.6	5.79	5.85	4.78	1.77
38	09:30	63.9	66.7	52.6	20.1	5.92	6.14	4.89	1.81
39	09:45	60.8	65.3	52.5	20.8	5.63	6.01	4.88	1.87
40	10:00	61.4	68.1	53.8	21.3	5.67	6.26	4.99	1.92
41	10:15	63.5	70.5	56.0	22.1	5.84	6.48	5.17	2.00
42	10:30	62.3	69.1	55.3	23.0	5.78	6.37	5.16	2.08
43	10:45	59.7	66.8	53.9	24.0	5.57	6.20	5.04	2.19
44	11:00	59.6	68.5	53.2	24.3	5.58	6.39	4.99	2.22
45	11:15	60.2	68.9	54.7	25.4	5.62	6.41	5.13	2.33
46	11:30	61.4	72.7	56.7	25.1	5.76	6.80	5.33	2.31
47	11:45	58.4	70.5	54.3	25.1	5.48	6.59	5.12	2.32
48	12:00	57.0	70.2	55.5	26.9	5.41	6.63	5.29	2.52
49	12:15	60.6	73.5	58.3	27.4	5.70	6.86	5.49	2.56
50	12:30	62.0	76.6	58.6	27.6	5.83	7.17	5.53	2.57
51	12:45	63.3	77.5	58.2	27.5	5.91	7.19	5.44	2.55

Slot	Time	MAE A.	MAE G.	MAE E.	MAE T.	MAPE A.	MAPE G.	MAPE E.	MAPE T.
52	13:00	63.4	77.8	57.0	26.2	5.94	7.22	5.34	2.42
53	13:15	64.6	78.8	58.6	26.4	6.14	7.41	5.58	2.49
54	13:30	67.4	82.7	60.3	25.1	6.48	7.85	5.80	2.38
55	13:45	66.1	80.7	58.9	24.9	6.39	7.72	5.71	2.38
56	14:00	67.2	79.3	59.5	23.3	6.54	7.62	5.80	2.22
57	14:15	66.9	81.0	59.6	23.0	6.57	7.86	5.86	2.21
58	14:30	68.4	83.0	60.7	22.6	6.74	8.07	6.00	2.18
59	14:45	69.3	82.3	62.5	22.5	6.85	8.03	6.19	2.17
60	15:00	68.8	80.1	60.8	21.9	6.81	7.83	6.03	2.12
61	15:15	68.5	78.8	59.5	22.2	6.84	7.77	5.94	2.17
62	15:30	66.1	77.3	56.9	21.6	6.62	7.64	5.73	2.10
63	15:45	65.8	77.5	55.9	20.8	6.60	7.69	5.62	2.03
64	16:00	62.8	74.4	53.6	19.7	6.28	7.37	5.38	1.92
65	16:15	61.8	73.0	53.1	19.6	6.21	7.25	5.34	1.91
66	16:30	59.2	70.4	50.8	19.5	5.93	6.98	5.10	1.90
67	16:45	59.2	69.5	50.2	19.9	5.89	6.83	4.98	1.92
68	17:00	55.2	66.4	47.9	21.0	5.51	6.55	4.79	2.03
69	17:15	54.9	66.4	47.2	20.7	5.44	6.48	4.68	1.97
70	17:30	52.9	63.9	45.1	20.6	5.19	6.18	4.43	1.95
71	17:45	50.7	61.7	43.2	19.9	4.91	5.91	4.20	1.86
72	18:00	48.1	59.4	41.2	20.2	4.67	5.69	4.01	1.87
73	18:15	46.5	57.2	39.1	19.6	4.47	5.42	3.77	1.80
74	18:30	44.9	55.2	38.0	19.7	4.29	5.21	3.64	1.80
75	18:45	43.0	53.8	35.8	20.0	4.05	5.03	3.39	1.81
76	19:00	40.3	50.9	33.4	19.8	3.79	4.73	3.14	1.79
77	19:15	39.6	49.6	32.3	19.0	3.69	4.57	3.02	1.70
78	19:30	38.4	47.8	31.0	18.2	3.56	4.39	2.88	1.61
79	19:45	37.6	47.1	30.3	17.6	3.45	4.29	2.79	1.55
80	20:00	35.7	45.8	28.9	18.1	3.27	4.17	2.65	1.59
81	20:15	34.3	45.1	27.8	18.0	3.15	4.10	2.54	1.58
82	20:30	34.1	44.5	27.8	17.8	3.13	4.05	2.55	1.57
83	20:45	32.9	43.5	26.7	17.7	3.03	3.97	2.46	1.57
84	21:00	31.3	42.7	26.1	18.5	2.92	3.95	2.43	1.67
85	21:15	31.0	42.1	25.7	18.4	2.93	3.95	2.42	1.68
86	21:30	29.5	41.0	24.8	18.5	2.84	3.91	2.38	1.72
87	21:45	27.5	39.8	23.7	20.6	2.72	3.90	2.33	1.97
88	22:00	26.6	38.7	23.5	22.0	2.68	3.85	2.35	2.14
89	22:15	27.7	39.3	23.5	20.4	2.83	3.97	2.39	2.02
90	22:30	26.7	38.8	22.9	19.4	2.79	4.01	2.37	1.95
91	22:45	26.8	38.2	23.8	19.1	2.86	4.02	2.54	1.96
92	23:00	25.0	36.9	21.9	18.7	2.75	3.98	2.40	1.96
93	23:15	23.8	35.3	20.6	18.4	2.67	3.88	2.30	1.98
94	23:30	23.6	34.5	20.9	18.9	2.72	3.89	2.39	2.08
95	23:45	22.7	33.6	20.1	18.9	2.69	3.88	2.36	2.14

Table 9.6: Sicily load: MAE and MAPE by fifteen minute slot

Slot	**Time**	**MAE A.**	**MAE G.**	**MAE E.**	**MAE T.**	**MAPE A.**	**MAPE G.**	**MAPE E.**	**MAPE T.**
0	00:00	73.3	99.5	67.6	43.3	3.88	5.13	3.58	2.24
1	00:15	69.1	94.1	61.1	40.5	3.74	4.95	3.30	2.16
2	00:30	69.0	92.8	61.0	42.2	3.82	5.00	3.37	2.30
3	00:45	68.8	91.8	60.1	42.1	3.90	5.04	3.41	2.34
4	01:00	69.1	91.8	61.0	40.1	3.98	5.11	3.51	2.27
5	01:15	69.1	90.6	60.7	38.4	4.05	5.15	3.56	2.22
6	01:30	67.8	89.4	60.3	36.6	4.03	5.15	3.58	2.14
7	01:45	67.4	89.0	59.6	34.8	4.06	5.20	3.59	2.06
8	02:00	67.0	89.6	60.0	32.7	4.06	5.28	3.64	1.95
9	02:15	65.8	86.5	59.5	32.2	4.04	5.17	3.66	1.95
10	02:30	65.7	86.7	58.7	31.3	4.06	5.22	3.63	1.91
11	02:45	65.7	86.8	59.7	31.0	4.09	5.26	3.72	1.91
12	03:00	64.9	85.7	58.8	28.9	4.04	5.21	3.68	1.78
13	03:15	64.8	85.0	58.4	28.6	4.08	5.21	3.68	1.77
14	03:30	64.0	84.6	57.8	28.5	4.03	5.20	3.65	1.78
15	03:45	63.2	83.2	57.5	28.5	3.99	5.12	3.64	1.78
16	04:00	61.7	81.2	56.1	25.8	3.88	4.98	3.53	1.60
17	04:15	63.7	81.5	57.7	26.1	4.02	5.03	3.65	1.62
18	04:30	63.4	81.7	57.0	25.7	4.00	5.05	3.60	1.59
19	04:45	64.3	79.9	56.1	25.8	4.03	4.91	3.52	1.59
20	05:00	67.3	83.9	58.2	24.7	4.15	5.07	3.59	1.50
21	05:15	68.1	82.7	57.3	24.2	4.17	4.98	3.50	1.47
22	05:30	71.2	84.2	59.1	25.5	4.32	5.02	3.58	1.52
23	05:45	75.0	84.8	61.9	27.7	4.48	4.99	3.70	1.63
24	06:00	86.1	92.2	67.2	37.2	4.93	5.20	3.85	2.09
25	06:15	92.0	96.5	72.7	38.7	5.26	5.43	4.15	2.16
26	06:30	98.2	101.0	78.3	35.6	5.52	5.61	4.39	1.96
27	06:45	103.0	103.0	82.1	33.9	5.65	5.60	4.49	1.83
28	07:00	115.0	112.0	90.7	45.1	6.03	5.82	4.77	2.31
29	07:15	122.0	117.0	97.6	42.7	6.23	5.93	5.01	2.14
30	07:30	131.0	126.0	109.0	41.5	6.47	6.15	5.38	2.01
31	07:45	142.0	134.0	114.0	38.5	6.74	6.33	5.43	1.81
32	08:00	151.0	142.0	121.0	46.3	6.88	6.43	5.55	2.08
33	08:15	156.0	149.0	124.0	43.9	6.95	6.56	5.52	1.93
34	08:30	161.0	155.0	130.0	42.3	6.98	6.67	5.62	1.82
35	08:45	164.0	162.0	133.0	41.3	6.97	6.82	5.65	1.74
36	09:00	165.0	164.0	136.0	41.8	6.86	6.76	5.67	1.73
37	09:15	164.0	165.0	138.0	43.4	6.73	6.74	5.69	1.77
38	09:30	158.0	163.0	136.0	44.7	6.45	6.60	5.58	1.81
39	09:45	153.0	158.0	134.0	46.2	6.22	6.36	5.46	1.87
40	10:00	145.0	155.0	128.0	47.4	5.89	6.23	5.23	1.92
41	10:15	142.0	153.0	125.0	49.3	5.75	6.14	5.11	2.00
42	10:30	138.0	152.0	126.0	51.3	5.59	6.10	5.15	2.08
43	10:45	141.0	154.0	126.0	53.5	5.76	6.24	5.19	2.19
44	11:00	140.0	153.0	125.0	54.2	5.68	6.18	5.15	2.22
45	11:15	137.0	151.0	124.0	56.6	5.57	6.10	5.11	2.33
46	11:30	133.0	149.0	124.0	55.9	5.44	6.05	5.13	2.31
47	11:45	133.0	153.0	125.0	55.9	5.44	6.21	5.13	2.32
48	12:00	128.0	151.0	125.0	59.8	5.31	6.23	5.20	2.52
49	12:15	136.0	157.0	127.0	60.9	5.58	6.42	5.25	2.56
50	12:30	138.0	161.0	128.0	61.5	5.66	6.57	5.28	2.57
51	12:45	147.0	170.0	131.0	61.1	6.04	6.96	5.40	2.55

Slot	Time	MAE A.	MAE G.	MAE E.	MAE T.	MAPE A.	MAPE G.	MAPE E.	MAPE T.
52	13:00	150.0	171.0	136.0	58.1	6.17	6.97	5.62	2.42
53	13:15	150.0	172.0	131.0	58.6	6.21	7.06	5.48	2.49
54	13:30	153.0	172.0	129.0	56.0	6.39	7.10	5.43	2.38
55	13:45	157.0	175.0	130.0	55.7	6.63	7.32	5.55	2.38
56	14:00	162.0	177.0	134.0	52.1	6.86	7.40	5.69	2.22
57	14:15	168.0	182.0	137.0	51.7	7.14	7.61	5.87	2.22
58	14:30	167.0	181.0	136.0	51.2	7.10	7.59	5.81	2.18
59	14:45	170.0	183.0	137.0	50.9	7.25	7.68	5.86	2.17
60	15:00	173.0	184.0	140.0	49.8	7.42	7.79	6.06	2.12
61	15:15	171.0	182.0	141.0	50.6	7.34	7.70	6.08	2.17
62	15:30	169.0	181.0	140.0	49.2	7.27	7.67	6.06	2.10
63	15:45	166.0	176.0	137.0	47.7	7.17	7.45	5.91	2.03
64	16:00	167.0	178.0	134.0	45.2	7.14	7.49	5.80	1.92
65	16:15	168.0	177.0	133.0	45.1	7.18	7.41	5.71	1.91
66	16:30	167.0	177.0	133.0	45.1	7.09	7.39	5.68	1.90
67	16:45	161.0	176.0	129.0	46.0	6.83	7.31	5.49	1.92
68	17:00	154.0	168.0	125.0	48.9	6.49	6.99	5.32	2.03
69	17:15	153.0	168.0	123.0	48.0	6.39	6.89	5.18	1.97
70	17:30	151.0	168.0	121.0	48.0	6.27	6.90	5.07	1.95
71	17:45	145.0	162.0	117.0	46.2	6.03	6.64	4.90	1.86
72	18:00	141.0	158.0	115.0	46.6	5.86	6.46	4.78	1.87
73	18:15	138.0	154.0	112.0	45.1	5.72	6.29	4.65	1.80
74	18:30	132.0	147.0	106.0	45.3	5.42	5.99	4.38	1.80
75	18:45	126.0	141.0	101.0	45.6	5.17	5.73	4.18	1.81
76	19:00	125.0	138.0	97.4	45.0	5.11	5.60	4.01	1.79
77	19:15	126.0	137.0	98.8	43.0	5.14	5.54	4.05	1.70
78	19:30	125.0	139.0	97.5	41.1	5.07	5.57	3.98	1.61
79	19:45	123.0	137.0	95.5	40.0	4.96	5.44	3.86	1.55
80	20:00	121.0	134.0	93.6	41.3	4.83	5.31	3.76	1.59
81	20:15	118.0	132.0	92.1	41.2	4.69	5.22	3.69	1.58
82	20:30	116.0	133.0	89.7	40.6	4.63	5.24	3.60	1.57
83	20:45	111.0	130.0	88.5	40.4	4.47	5.17	3.59	1.57
84	21:00	106.0	126.0	86.0	42.2	4.35	5.11	3.54	1.67
85	21:15	103.0	125.0	83.5	41.7	4.31	5.18	3.51	1.68
86	21:30	100.0	123.0	81.6	41.8	4.26	5.18	3.49	1.72
87	21:45	96.3	122.0	79.8	46.8	4.19	5.24	3.48	1.97
88	22:00	92.4	118.0	75.7	50.1	4.08	5.13	3.36	2.14
89	22:15	92.7	118.0	75.8	46.3	4.17	5.23	3.43	2.02
90	22:30	89.5	116.0	74.2	43.9	4.11	5.27	3.42	1.95
91	22:45	85.1	112.0	72.7	43.1	4.00	5.19	3.43	1.96
92	23:00	83.3	111.0	71.9	42.1	4.00	5.24	3.46	1.96
93	23:15	78.6	105.0	67.3	41.3	3.88	5.09	3.33	1.98
94	23:30	75.3	101.0	65.0	42.4	3.80	5.01	3.28	2.08
95	23:45	73.3	99.5	64.3	42.5	3.79	5.03	3.32	2.14

Bibliography

[1] Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 1986.

[2] George E. P. Box and David R. Cox. An analysis of transformations. *Journal of the Royal Statistical Society: Series B*, 1964.

[3] Adam E. Clements, A. Stan Hurn, and Zili Li. Forecasting day-ahead electricity load using a multiple equation time series approach. *European Journal of Operational Research*, 2016.

[4] Robert B. Cleveland, William S. Cleveland, Jean E. McRae, and Irma Terpenning. Stl: A seasonal-trend decomposition procedure based on loess. *Journal of Official Statistics*, 1990.

[5] Javier Contreras, Rosario Espínola, Francisco J. Nogales, and Antonio J. Conejo. Arima models to predict next-day electricity prices. *IEEE Transactions on Power Systems*, 2003.

[6] Virginie Dordonnat, Siem Jan Koopman, Marius Ooms, Alain Dessertaine, and Jérôme Collet. An hourly periodic state space model for modelling french national electricity load. *International Journal of Forecasting*, 2008.

[7] Robert F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 1982.

[8] Victor M. Guerrero. Time-series analysis supported by power transformations. *Journal of Forecasting*, 1993.

[9] Tao Hong and Shu Fan. Probabilistic electric load forecasting: A tutorial review. *International Journal of Forecasting*, 2016.

[10] Tao Hong, Pierre Pinson, and Shu Fan. Global energy forecasting competition 2012. *International Journal of Forecasting*, 2014.

[11] Rob J. Hyndman and George Athanasopoulos. *Forecasting: Principles and Practice.* OTexts, Melbourne, Australia, 2021.

[12] Rob J. Hyndman, Anne B. Koehler, Ralph D. Snyder, and Simone Grose. A state space framework for automatic forecasting using exponential smoothing methods. *International Journal of Forecasting*, 2002.

[13] Jesus Lago, Fjo De Ridder, Peter Vrancx, and Bart De Schutter. Forecasting day-ahead electricity prices in europe: The importance of considering market integration. *Applied Energy*, 2018.

[14] Jesus Lago, Grzegorz Marcjasz, Bart De Schutter, and Rafał Weron. Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy*, 2021.

[15] Lacir J. Soares and Marcelo C. Medeiros. Modeling and forecasting short-term electricity load: A comparison of methods with an application to brazilian data. *International Journal of Forecasting*, 2008.

[16] James W. Taylor. Triple seasonal methods for short-term electricity demand forecasting. *European Journal of Operational Research*, 2010.

[17] James W. Taylor and Patrick E. McSharry. Short-term load forecasting methods: An evaluation based on european data. *IEEE Transactions on Power Systems*, 2007.

[18] Shadi Tehrani, Jesús Juan, and Eduardo Caro. Electricity spot price modeling and forecasting in european markets. *Energies*, 2022.

[19] Rafał Weron. Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, 2014.

[20] Zhongyang Zhao, Caisheng Wang, Matthew Nokleby, and Carol Miller. Improving short-term electricity price forecasting using day-ahead lmp with arima models. In *2017 IEEE Power & Energy Society General Meeting*, 2017.

[21] Florian Ziel and Rafał Weron. Day-ahead electricity price forecasting with high-dimensional structures: Univariate vs. multivariate modeling frameworks. *Energy Economics*, 2018.