



GRADO EN INGENIERÍA EN TECNOLOGÍAS INDUSTRIALES

TRABAJO DE FIN DE GRADO

Strategy optimization for arbitrage between Spanish Energy Markets

Autor: Teresa Martínez Serrano

Director: Christian de Peretti

Lyon
Julio 2026

Declaration of originality

I declare under my responsibility that the Project presented with the title **STRATEGY OPTIMIZATION FOR ARBITRAGE BETWEEN SPANISH ENERGY MARKETS** at the ICAI School of Engineering of the Comillas Pontifical University in the academic year **2025/2026** is of my authorship and has not been presented previously for other purposes. The Project is not plagiarised from any other, either totally or partially, and the information that has been taken from other documents is duly referenced.

Use of Artificial Intelligence

I declare under my responsibility that (indicate the correct option):

☐ I have not used Artificial Intelligence in the preparation of this document.

☒ I have used Artificial Intelligence in the preparation of this document and/or Annex B under the conditions allowed by Comillas Pontifical University, i.e. applying Level 2 of the Perkins et al. (2024) Assessment Scale: "AI can be used for pre-task activities such as brainstorming, description and initial research. This level focuses on the use of AI for planning, synthesising and generating ideas, but assessments should emphasise the ability to develop and refine these ideas independently". Specifically, Artificial Intelligence has been used to:

(indicate here the concrete use that has been made of Artificial Intelligence)
It has been used for initial research and identifying relevant bibliographic sources.

Signature (student): Martínez Serrano, Teresa

Date: 14/07/2026



This declaration refers to the use of generative Artificial Intelligence to carry out the Project documents (Annex B and Memory). It does not apply to Projects where, by their nature, artificial intelligence must be used as part of them (application of machine learning techniques, neural networks, data analysis...).

Authorisation for Project delivery

Thesis supervisor	Thesis deputy supervisor (if any)
Christian de Peretti Associate Professor fo Finance Ecole Centrale de Lyon 	
Signature:	Signature:
Date: 14/07/2026 Lyon, France	Date:



GRADO EN INGENIERÍA EN TECNOLOGÍAS INDUSTRIALES

TRABAJO DE FIN DE GRADO

Strategy optimization for arbitrage between Spanish Energy Markets

Autor: Teresa Martínez Serrano

Director: Christian de Peretti

Lyon
Julio 2026

OPTIMIZACIÓN DE ESTRATEGIAS DE ARBITRAJE ENTRE MERCADOS ELÉCTRICOS ESPAÑOLES

Autor: Martínez Serrano, Teresa.

Director: de Peretti, Christian.

Entidad Colaboradora: ENGIE.

RESUMEN

En este trabajo se estudian tres modelos de predicción de precios para el mercado diario e intradiario 1 del Mercado Eléctrico Español, con el objetivo de desarrollar una estrategia de arbitraje entre ambos mercados.

Palabras clave: Regresión Lineal, SARIMAX, LightGBM, Mercado Eléctrico, Arbitraje.

1. Introducción

Los precios de los mercados eléctricos cambian continuamente ya que dependen de numerosos factores, entre ellos las condiciones meteorológicas y los precios del combustible. Además, hay varios mercados eléctricos dentro del sistema eléctrico español y cada uno tiene su propio precio. Estos mercados son necesarios para asegurar la estabilidad de la red eléctrica.

En este contexto, aparecen oportunidades de arbitraje. Una empresa puede comprar en uno de los mercados eléctricos españoles y vender en otro. Este tipo de estrategias requieren predicciones precisas de los precios. Estas predicciones no son sencillas ya que los precios dependen de numerosos factores y pueden experimentar importantes fluctuaciones. Las estrategias deben ser estudiadas cuidadosamente para encontrar un balance entre los beneficios y riesgos de esta.

El principal objetivo de este proyecto es proporcionar a la empresa ENGIE, unas bases para trabajo futuro. Como estudiar modelos predictivos y diferentes estrategias lleva tiempo, ENGIE quiere tener un análisis preliminar en el que varios modelos predictivos y estrategias de arbitraje se evalúan para tener un punto de partida desde el cuál puedan mejorar tanto los modelos como la estrategia. Este trabajo se centra en los mercados diario e intradiario 1.

En una primera parte se realiza un análisis de las variables para entenderlas y para hacer una primera selección. Después, se construyen las predicciones ajustando sus parámetros y seleccionando las exógenas. Finalmente se construyen estrategias de arbitraje sencillas sin utilizar las predicciones y otras implementándolas.

2. Metodología

Para las predicciones de precios se han usado tres modelos: regresión lineal, SARIMAX y LightGBM. Finalmente se ha determinado que el más adecuado era LightGBM.

Los residuos de la regresión lineal tienen autocorrelación y heterocedasticidad, haciendo que los intervalos de confianza no sean fiables. La autocorrelación muestra que hay información en los residuos que no se está aprovechando. Estos dos puntos hacen que este modelo se descarte. Solo se usó para eliminar alguna de las variables exógenas fijándose en las métricas de evaluación del modelo.

El modelo SARIMAX permite incorporar las dependencias temporales de las series. Este modelo debería eliminar o disminuir la autocorrelación de los residuos, sin embargo, los resultados obtenidos no son nada positivos, con métricas significativamente peores.

LightGBM es un modelo de árbol, el conjunto de datos se divide iterativamente en grupos llamados hojas a medida que se van imponiendo condiciones para estar en un grupo u otro. Este modelo tiene en sus residuos autocorrelación, no se ha podido encontrar un método sencillo para poder capturarla en el modelo, sin embargo, esta no invalida el modelo. Esto ha hecho que sea elegido para incorporarlo en la estrategia. La heterocedasticidad presente en sus residuos no es sorprendente debido a los picos presentes en los precios eléctricos.

Todas las estrategias desarrolladas se basan en determinar el signo de la diferencia de precio entre los dos mercados. Esto se usa para determinar en cuál de los dos se compra y en cuál se vende.

Para poder ver el impacto de utilizar la predicción en la estrategia, se elaboraron estrategias sencillas. Una se basa en vender siempre en el mercado diario mientras que otra utiliza una media móvil de la diferencia de precios entre los mercados diario e intradiario. Estas dos estrategias dan beneficios hasta octubre, después de este mes empiezan a generar pérdidas. Una posible explicación son las reformas del mercado eléctrico en este mes.

A esta estrategia sencilla se le añadieron umbrales a partir de los cuales no se operaba. Se observó que estos disminuyen el riesgo de la estrategia, pero también el beneficio. Esto significa que no mejoran el rendimiento de la estrategia, simplemente se pueden usar para balancear riesgo y beneficio.

La estrategia incorporando las predicciones da resultados ligeramente mejores. En esta

estrategia se calcula la diferencia de precios entre los dos mercados usando las predicciones y según el signo se elige en cual comprar o vender.

3. Resultados

La Figura 1 muestra el PnL diario y el acumulado de esta estrategia. Observando las gráficas de beneficio medio por cuarto de hora se decidió no operar en los cuartos de hora que generaban pérdidas. Esto mejoró también ligeramente los resultados, pero puede sobreajustar la estrategia a ese periodo concreto.




Figura 1: Evolución de las ganancias por cuarto de hora y acumuladas para la estrategia usando las predicciones de LightGBM.

En la Tabla 1, se muestran las principales métricas diarias de la estrategia basada en la predicción LightGBM.

PnL medio (€/MWh)	Volatilidad (€/MWh)	Ratio de Sharpe	Tasa de aciertos
201.32	242.31	0.831	0.840

Tabla 1: Métricas diarias de la estrategia basada en la predicción.

4. Conclusiones

Las predicciones obtenidas tienen margen de mejora. Dos de ellas no son adecuadas para este trabajo, mientras que la tercera ha dado resultados aceptables que aún podrían mejorarse.

En cuanto a las estrategias, sin haber desarrollado ninguna especialmente compleja, se ha visto que generan ganancias. Las predicciones ayudaron a mejorar la estrategia y a que las reformas de mercado afectaran menos a su rendimiento.

Como trabajo futuro, se debería estudiar el volumen máximo que se puede comprar sin modificar los precios de mercado. También hay margen de mejora para la estrategia si se estudia cómo hacer que el volumen de energía comprado y vendido cambie para optimizar las ganancias.

STRATEGY OPTIMIZATION FOR ARBITRAGE BETWEEN SPANISH ENERGY MARKETS

Author: Martínez Serrano, Teresa.

Supervisor: de Peretti, Christian.

Collaborating Institution: ENGIE.

This thesis analyses three electricity price forecasting models for the Spanish day-ahead market and intraday 1 market, with the aim of developing an arbitrage strategy between these two markets.

Keywords: Linear Regression, SARIMAX, LightGBM, Electricity Market, Arbitrage.

ABSTRACT

1. Introduction

Electricity market prices are changing continuously as they depend on numerous factors, like the weather conditions and fuel prices. In addition, there are several electricity markets and the prices at which the energy companies sell their electricity vary from one market to another. These markets are necessary to ensure the stability of the electricity grid.

In this context, arbitrage opportunities may arise. A company can buy in one of the Spanish electricity markets and sell in another, or vice versa. These types of strategies require accurate electricity price forecasts. These forecasts are not simple because, as mentioned previously, the electricity prices depend on a great number of factors and may experience significant fluctuations. The possible strategies have to be also carefully studied to achieve a balance between profitability and risk.

The main objective of this project is to provide ENGIE with a foundation for future work. Since studying predictive models and different strategies is time-consuming, ENGIE wants to have a preliminary analysis in which several forecasting methods and arbitrage strategies are evaluated in order to have a starting point from which they can improve the forecasting models and the trading strategies. The company wants to focus on the day-ahead market and the first intraday market.

The first task is to analyse the data to understand it and to make a first selection. Then, the models are built adjusting their parameters and selecting the exogenous variables. Finally, the arbitrage strategies are built, starting by simple ones that do not use the forecasts and then others that use it.

2. Methodology

Three forecasting models were evaluated for electricity price prediction: Linear Regression, SARIMAX, and LightGBM. LightGBM was the one with the best results.

The residuals of the Linear Regression model show autocorrelation and heteroscedasticity, making the confidence intervals unreliable. The presence of autocorrelation indicates that part of the information contained in the residuals is not being captured by the model. Consequently, this model was discarded. It was only used to select variables based on the model metrics.

The SARIMAX model allows temporal dependencies to be incorporated into the forecasting process. In principle, this should reduce or eliminate the autocorrelation present in the residuals. However, the obtained results were not satisfactory, with evaluation metrics significantly worse than those of the Linear Regression model.

LightGBM is a tree-based model in which the dataset is iteratively divided into groups, known as leaves, by applying a sequence of splitting conditions. The residuals of this model still present autocorrelation, and no simple method was found to capture it within the model. This does not invalidate the LightGBM model. As a result, it was selected as the forecasting model used in the arbitrage strategy. The heteroscedasticity observed in its residuals is not unexpected due to the price spikes that frequently occur in electricity markets.

All the developed strategies are based on determining the sign of the price difference between the two markets. This information is used to decide in which market electricity should be bought and in which it should be sold.

To evaluate the impact of incorporating price forecasts into the strategy, two simple strategies were developed. The first always sells electricity in the day-ahead market, while the second uses a moving average of the price difference between the day-ahead and intraday markets. Both strategies generate profits until October, after they begin to produce losses. A possible explanation is the electricity market reforms introduced in that month.

The forecast-based strategy gave better results than the other ones. This strategy determines whether electricity should be bought in the day-ahead or the intraday 1 market using the price forecasts.

Thresholds were then added to these simple strategies so that no trades were executed when the expected price difference was below a given value. It was observed that increasing the threshold reduces the risk of the strategy, but also decreases its profitability. Therefore, thresholds do not improve the overall performance of the strategy, but rather

provide a mechanism for balancing risk and return.

3. Results

Figure 2 shows the daily PnL and the accumulated PnL. Based on the average PnL per quarter-hour, it was also tried to not operate in the quarter-hours where the average PnL was negative, but this may result in overfitting the strategy to the test period.




Figure 2: Evolution of the accumulated PnL and the daily PnL for the forecast-based strategy.

Table 2 contains the main daily metrics of the strategy based on the LightGBM forecast.

Average PnL (€/MWh)	Volatility(€/MWh)	Sharpe ratio	Win rate
201.32	242.31	0.831	0.840

Table 2: Daily metrics of the forecast-based strategy.

4. Conclusions

The forecasting models presented in this work can still be improved. Two of the evaluated models proved unsuitable for this application, while the third produced acceptable results that could still be improved.

Regarding the arbitrage strategies, it has been shown that profits can be achieved even without developing complex strategies. Furthermore, incorporating price forecasts im-

proved the strategy and reduced the impact of the electricity market reforms on its performance.

Future work should investigate the maximum volume of electricity that can be traded without affecting market prices. In addition, the strategy could be further improved by optimising the amount of electricity bought and sold in each trading period to maximise profits.

Acknowledgments

I would like to express my sincere gratitude to everyone who has supported me throughout the completion of this project.

First, I would like to thank my tutor, Christian de Peretti, for his guidance and support throughout the project. His advice and knowledge have been fundamental to its successful completion.

I would also like to thank my family and friends for their constant encouragement throughout my degree. Their patience, confidence, and support over the past four years have made this achievement possible.

Contents

1	Introduction	1
1.1	Related Work	2
1.2	Structure of the project	3
1.3	Alignment with the Sustainable Development Goals (SDGs)	3
2	The Spanish Electricity Market	5
2.1	Wholesale Electricity Market	6
2.1.1	Day-ahead Market	6
2.1.2	Intraday Market	7
2.2	Adjustment and Balancing Services	8
3	The dataset	11
3.1	Weather information	13
3.1.1	Weighting coefficients	13
3.1.2	Weather forecast	13
4	Data analysis	15
4.1	Descriptive statistics	15
4.2	Time series visualization	16
4.2.1	Dispatchable and non-dispatchable energy sources	17
4.2.2	Comparison of dispatchable and non-dispatchable energy sources	20
4.2.3	Intraday 1 market and the actual generation mix	22
4.3	Distribution analysis	23
4.4	Stationary analysis	25
4.5	Autocorrelation analysis	27
4.6	Correlation analysis	31
5	Electricity Price Forecasting	37
5.1	Evaluation Metrics	37

5.2	Linear Regression	38
5.2.1	Forecast	39
5.3	SARIMAX	43
5.3.1	SARIMAX model	43
5.3.2	SARIMAX parameters	44
5.4	LightGBM	46
5.4.1	Tree-based methods	47
5.4.2	Gradient Boosting	48
5.4.3	LightGBM	49
5.4.4	Forecast	50
6	Arbitrage strategy	57
6.1	The Naive strategy	59
6.2	The Moving Average strategy	60
6.2.1	The Moving Average with threshold	62
6.3	The Sortino Ratio strategy	63
6.4	Forecast-based strategy	64
7	Conclusion	69
7.1	Future work	69
	Bibliography	71
8	Appendices	75
8.1	Data-set	75
8.1.1	Weather information	75
8.2	Electricity Price Forecasting	78
8.2.1	Linear Regression	78
8.2.2	LightGBM	82

List of Figures

1	Evolución de las ganancias por cuarto de hora y acumuladas para la estrategia usando las predicciones de LightGBM.	v
2	Evolution of the accumulated PnL and the daily PnL for the forecast-based strategy.	ix
2.1	Aggregated supply and demand curves for the 27/05/2026 and period H1Q1. Source: [27], translated into English.	7
4.1	Day-ahead and Intraday 1 market prices.	17
4.2	Non-dispatchable energy production, total demand and day-ahead market price	18
4.3	Percentage of non-dispatchable energy sources in the day-ahead market	18
4.4	Dispatchable energy production and day-ahead market price.	19
4.5	Share of dispatchable energy sources in total day-ahead electricity generation.	20
4.6	Dispatchable and non-dispatchable electricity generation and day-ahead market price.	21
4.7	Share of dispatchable and non-dispatchable electricity generation and day-ahead market price.	21
4.8	Non-dispatchable energy production and intraday 1 market price.	22
4.9	Dispatchable energy production and intraday 1 market price.	23
4.10	Day-ahead market price histogram.	24
4.11	Intraday 1 market price histogram.	24
4.12	Autocorrelation Function of the day-ahead market price	28
4.13	Autocorrelation Function of the day-ahead market price after seasonally differencing.	28
4.14	Autocorrelation Function of the intraday 1 market price.	29
4.15	Autocorrelation Function of the intraday 1 market price after seasonally differencing.	30
4.16	Correlation matrix for the day-ahead market.	32

4.17	Correlation matrix for the intraday 1 market.	33
5.1	Example of a tree. [17]	47
5.2	Division of the predictor following the example tree from Figure 5.1	48
5.3	Importance of the variables for the day ahead market	51
5.4	Importance of the variables for the day ahead market after elimi- nating.	51
5.5	Importance of the variables for the intraday 1 market	52
5.6	Importance of the variables for the intraday 1 market after elimi- nating.	52
5.7	LightGBM forecast compared to the real day-ahead price.	54
5.8	LightGBM forecast compared to the real intraday 1 price.	55
6.1	Evolution of the accumulated PnL and the daily PnL for the naive strategy.	59
6.2	Evolution of the accumulated PnL and the daily PnL for the moving average strategy.	61
6.3	Evolution of the accumulated PnL and the daily PnL for the forecast- based strategy.	64
6.4	Average PnL per quarter-hour of the forecast-based strategy.	66
6.5	Win rate of the forecast-based strategy.	66
8.1	Autocorrelation function of the residuals of the linear regression model for the day-ahead market.	78
8.2	Autocorrelation function of the residuals of the linear regression model for the intraday 1 market.	79
8.3	Partial autocorrelation function of the residuals of the linear regres- sion model for the day-ahead market.	79
8.4	Partial autocorrelation function of the residuals of the linear regres- sion model for the intraday 1 market.	80
8.5	Residuals over time for the linear regression model of the day-ahead market.	80
8.6	Residuals over time for the linear regression model of the intraday 1 market.	81
8.7	Residuals versus fitted values for the linear regression model for the day-ahead market.	81
8.8	Residuals versus fitted values for the linear regression model for the intraday 1 market.	82
8.9	Autocorrelation function of the residuals of the LightGBM for the day-ahead market	83

8.10	Autocorrelation function of the residuals of the LightGBM for the intraday 1 market	83
8.11	Partial autocorrelation function of the residuals of the LightGBM model for the day-ahead market.	84
8.12	Partial autocorrelation function of the residuals of the LightGBM model for the intraday1 market.	84
8.13	Residuals versus fitted values for the LightGBM model for the day-ahead market.	85
8.14	Residuals versus fitted values for the LightGBM model for the intraday 1 market.	85
8.15	Residuals over time for the LightGBM model of the day-ahead market.	86
8.16	Residuals over time for the LightGBM model of the intraday 1 market.	86

List of Tables

1	Métricas diarias de la estrategia basada en la predicción.	v
2	Daily metrics of the forecast-based strategy.	ix
2.1	Schedule of the Spanish intraday auction market sessions. Table translated from [31].	8
3.1	Dataset information.	11
3.2	Suffix reference table	12
3.3	Average weighting coefficients of installed solar and wind power in the peninsular autonomous communities.	14
4.1	Descriptive statistics of the variables used in the study.	16
4.2	Augmented Dickey-Fuller test results for the DA and the IDA1 markets.	25
4.3	KPSS Test Results for the DA and the IDA1 markets.	25
4.4	Augmented Dickey-Fuller and KPSS test results for the exogenous variables.	26
4.5	Augmented Dickey-Fuller and KPSS test results for the seasonal differenced exogenous variables.	26
4.6	Autocorrelation results for the day-ahead market	27
4.7	Autocorrelation results for the day-ahead market price after seasonally differencing.	27
4.8	Autocorrelation results for the intraday 1 market price.	29
4.9	Autocorrelation results for the intraday 1 market after seasonally differencing.	29
4.10	Partial Autocorrelation Function of the day-ahead market price. . .	30
4.11	Partial Autocorrelation Function of the intraday 1 market price. . .	31
4.12	VIF for the day-ahead market variables before removing collinear variables.	34
4.13	VIF for the intraday market variables before removing collinear variables.	35

4.14	VIF for the day-ahead market variables after removing collinear variables.	35
4.15	VIF for the intraday 1 market variables after removing collinear variables.	36
5.1	Jarque-Bera test of the residuals of the regression model.	39
5.2	ADF and KPSS tests of the residuals of the regression model. . . .	40
5.3	Breusch-Pagan test of the residuals of the regression model.	40
5.4	Ljung-Box of the residuals of the DA and IDA1 regression models. .	40
5.5	Regression coefficients for the day-ahead market.	41
5.6	Regression Coefficients for the day-ahead market after variable selection.	41
5.7	Linear Regression results for the day-ahead market.	41
5.8	Regression coefficients for the intraday 1 market.	42
5.9	Regression Coefficients for the intraday 1 after variable selection. .	42
5.10	Linear Regression results for the intraday 1 market.	42
5.11	SARIMAX possible parameters for the day-ahead and the intraday 1 markets.	46
5.12	Results of the SARIMAX model of the day-ahead market with different parameters.	46
5.13	Results of the SARIMAX model of the intraday 1 market with different parameters.	46
5.14	Jarque-Bera test of the residuals of the lightGBM model.	53
5.15	ADF and KPSS tests of the residuals of the lightGBM model. . . .	53
5.16	Breusch-Pagan test of the residuals of the lightGBM model.	53
5.17	Ljung-Box of the residuals of the DA and IDA1 lightGBM models. .	53
5.18	Results LightGBM	54
6.1	Metrics of the naive strategy (15 min).	59
6.2	Metrics of the naive strategy (24 h).	60
6.3	Metrics of the moving average strategy (15 min).	61
6.4	Metrics of the moving average strategy (24 h).	62
6.5	Metrics of the moving average strategy with threshold (15 min). . .	62
6.6	Metrics of the moving average strategy with threshold (24 h). . . .	63
6.7	Metrics of the Sortino strategy (15 min).	63
6.8	Metrics of the Sortino strategy (24 h).	64
6.9	Metrics of the forecast-based strategy (15 min).	65
6.10	Metrics of the forecast-based strategy (24 h).	65
6.11	Metrics of the forecast-based strategy after eliminating quarter-hours (15 min).	67

6.12	Metrics of the forecast-based strategy after eliminating quarter-hours (24 h).	67
8.1	Monthly photovoltaic weighting coefficients, part I.	75
8.2	Monthly photovoltaic weighting coefficients, part II.	76
8.3	Monthly wind weighting coefficients, part I.	76
8.4	Monthly wind weighting coefficients, part II.	77
8.5	Latitude and longitude of the cloud cover and wind data points . .	77

Chapter 1

Introduction

This project was carried out as part of the General Engineering program at École Centrale de Lyon. During the second year of the program, the students work on a project in collaboration with a company. During the project life cycle, they are required to address the company's needs and have several meetings to show their progress. In addition, a university teacher also supervises and evaluates the project to ensure that it meets the academic requirements.

The academic objective of this project is to expose students to a professional environment and the development of an engineering project. The project topic is chosen by the student, which allows them to gain in-depth knowledge of a specific engineering field.

ENGIE is the company that proposed this project. An alumnus from École Centrale de Lyon who is currently working at ENGIE España has been the supervisor.

ENGIE is a French multinational energy company. Its history starts in the nineteenth century with the *Compagnie Universelle du Canal Maritime de Suez*. In 2008, SUEZ and Gaz de France merged to form GDF SUEZ and in 2015 it changed its name to ENGIE [3].

Today, ENGIE is involved in a wide range of energy-related activities, including renewable energy projects like wind farms and photovoltaic solar plants. ENGIE also develops battery energy storage systems in India, Chile and Belgium, district heating and cooling networks, and energy supply solutions for data centers. Although ENGIE is committed to the energy transition, it continues to work on gas projects. The company is also working on electrical transmission lines in Brazil ([5], pp. 8-9). All these activities show the wide variety of services that ENGIE provides within the energy sector.

Particularly, in Spain, ENGIE has developed renewable energy facilities such as wind farms, photovoltaic solar plants, and biomass power plants. The company also operates biogas production facilities and district cooling and heating.[4].

Electricity market prices are changing continuously as they depend on numerous factors, like the weather conditions and fuel prices. In addition, there are several electricity markets and the prices at which the energy companies sell their electricity vary from one market to another. These markets are necessary to ensure the stability of the electrical grid.

In this context, arbitrage opportunities may arise. A company can buy in one of the Spanish electricity markets and sell in another, or vice versa. These type of strategies require accurate electricity price forecasts. These forecast are not simple because, as mentioned previously, the electricity prices depend on a great number of factors and may experience significant fluctuations. The possible strategies have to be also carefully studied to achieve a balance between profitability and risk.

The main objective of this project is to provide ENGIE with a foundation for future work. Since studying predictive models and different strategies is time-consuming, ENGIE wants to have a preliminary analysis in which several forecasting methods and arbitrage strategies are evaluated in order to have a starting point from which they can improve the forecasting models and the trading strategies. The company wants to focus on the day-ahead market and the first intraday market. The operation of both markets is described in Chapter 2.

1.1 Related Work

For the prediction of prices, these types of projects usually start by studying the results of different statistical prediction methods such as SARIMAX. They also use neural network models such as Long Short-Term Memory networks (LSTM) or Convolutional Neural Networks (CNNs), and they also study transformers [22, 35].

It is common to start by developing SARIMAX models because they capture the temporal dependencies and the influence of exogenous variables. However, the reviewed studies conclude that this model is not able to capture to the prices variations. Neural networks and Transformers give better results. The problem of Transformers is that they require longer training time than the other models [35].

1.2 Structure of the project

The project begins with a study of the Spanish electricity market. This is an important part as it provides the foundation for the rest of the work. Without an understanding of the market rules, there is a risk of developing a strategy that uses information that is not available at the time the forecast is made, for example.

Next, the forecasting models are developed. They incorporate temporal dependencies as well as the exogenous variables that can influence the prices. They are evaluated in order to choose the most suitable model for the arbitrage strategy.

Finally, the trading strategies are developed, starting from simple ones and then incorporating more sophisticated ones. This makes it possible to compare the simple strategies that do not rely on price forecasts with more complex ones which include them, thereby evaluating the contribution of the forecasts.

1.3 Alignment with the Sustainable Development Goals (SDGs)

This project will contribute to the 9th SDG, *Industry, Innovation and Infrastructure* and the 13th SDG, *Climate Action* as it will help ENGIE, a company involved in the energy transition and promoting a low-carbon electricity system, to gain a deeper understanding of the Spanish electricity market.

The project could also contribute to the 7th SDG, *Affordable and Clean Energy*. Electricity market arbitrage seeks to buy energy in the markets with lower prices and to sell it in the markets with higher prices. Since renewable electricity, such as wind or solar, is often sold at lower prices, this project could reduce the use of more polluting energy sources.

Chapter 2

The Spanish Electricity Market

It is fundamental to balance electricity supply and demand. The power consumed must be equal to the power generated. To achieve this, the wholesale electricity market is composed of several markets ([30], p. 17). These different markets have different closing times and market mechanisms, which help guarantee the continuous balancing of the electricity grid.

It is important to differentiate between the different activities within the electricity sector. Some of them are managed by a single company, while others are shared among several. There are four main activities that have to be distinguished in the Spanish electricity sector [32].

- Generation: electricity production.
- Transport: transport of electricity, managed by Red Eléctrica.
- Distribution: delivery of electricity from the transmission network to end consumers.
- Supply: sale of electricity to consumers.

The participants in the electricity markets include electricity producers, suppliers, direct consumers, energy storage facility owners, and representatives ([30], p19). The market operator, the electricity system operator and the regulators are responsible for ensuring the proper operation of the market and the electricity grid.

OMIE (Operador del Mercado Ibérico de Energía) is the operator of the Iberian electricity market. Its main function is to receive market bids and match electricity supply and demand [29].

Red Eléctrica (REE) is the operator of the Spanish electricity system. Its main

mission is to transport electricity from generation plants to the distribution network while guaranteeing the security and reliability of electricity supply.

2.1 Wholesale Electricity Market

The wholesale electricity market is composed of several submarkets in which the participants can buy or sell electricity. These include the day-ahead market, the intraday market, the forward market, adjustment and balancing services and non-organised markets [30].

2.1.1 Day-ahead Market

The day-ahead market is interconnected with the Portuguese and French electricity markets [37]. As its name indicates, this market closes the day before the day in which the production and consumption will take place. It closes at 12:00, at which point OMIE matches electricity supply and demand ([28], p. 3).

Formation of prices

Electricity prices are determined by the supply and demand curves. The intersection point defines the bids and offers that will be accepted. All accepted participants pay or are paid the same market clearing price. Therefore, the day-ahead market is a pay-as-clear market. Participants do not pay or receive the price that they offered as it occurs in the pay-as-bid markets [37].

Bids are made for each quarter-hour of the next day, so there are 96 prices. The selling bids are ordered in ascending price order and the purchase bids in descending price order in order to match supply and demand for 24 hours ([28], p. 4).

The algorithm used to clear the market is called *Euphemia* ([28], p5). Red Eléctrica receives the results of *Euphemia* in order to validate their technical feasibility. It is possible that some adjustments are necessary to guarantee the respect of technical restrictions ([28], p. 5). Figure 2.1 shows the supply and demand curves for 27 May 2026 for the first quarter-hour of the first hour.

Electricity generation bids are made based on opportunity cost. This means that their price depends on, for example, at what price they could sell their fuel if it were used for another purpose instead of electricity generation. For hydroelectric generation, storing or using water during specific periods of the day may provide higher economic value than during others, this determines their opportunity cost [37].




Figure 2.1: Aggregated supply and demand curves for the 27/05/2026 and period H1Q1. Source: [27], translated into English.

2.1.2 Intraday Market

The intraday market is divided into the auction markets and the continuous market. There are six intraday auction markets within the Iberian Electricity Market (MIBEL) and a European continuous market. The intraday auctions allow participants to adjust electricity purchases and sales up to four hours before the delivery period. Prices in the intraday market are usually close to the day-ahead market prices. Table 2.1 shows when each of the intraday auction markets open, close, clear, publish the results and the delivery horizon they cover [31].

Bids are divided into simple and complex. Simple bids are those in which energy is offered or demanded through one to five bid blocks for each trading period, in each block a price and a quantity of energy are specified. Complex bids, are those that include additional complex conditions such as load gradient conditions [31].

OMIE is responsible for the economic management of the day-ahead and the intraday markets. It is the entity that receives all purchase and sale bids and matches supply and demand after the closure of the market. It is also the entity that publishes the results [31].

The closure times of the six intraday auction sessions are different. The first one closes at 15:00 on the day before electricity delivery [31].

	Session 1	Session 2	Session 3	Session 4	Session 5	Session 6
Session Opening	14:00	17:00	21:00	1:00	4:00	9:00
Session Closure	15:00	17:50	21:50	1:50	4:50	9:50
Market Clearing	15:00	17:50	21:50	1:50	4:50	9:50
Publication of Aggregated Schedule (PIBCA)	15:07	17:57	21:57	1:57	4:57	9:57
Publication of Final Hourly Schedules (PHF)	16:20	18:20	22:20	2:20	5:20	10:20
Programming Horizon	24 h	28 h	24 h	20 h	17 h	12 h
Hourly Periods	1-24 D+1	21-24 and 1-24 D+1	1-24 D+1	5-24	8-24	13-24

Table 2.1: Schedule of the Spanish intraday auction market sessions. Table translated from [31].

The continuous intraday market (MIC) aims to allow continuous electricity trading in the European zones in order to improve the overall efficiency. Electricity bids from one country can be matched with participants from other countries if there is sufficient capacity of transportation available [31].

In this project, only the day-ahead (DA) market and the first intraday auction market (IDA) are considered in order to simplify the trading strategy. This reduces the number of forecasts that are made. The trading strategy is simpler because it involves only two markets. Naturally, using all the markets would be the ideal case and would likely improve the trading performance as long as the forecasts are good enough.

2.2 Adjustment and Balancing Services

Electricity cannot be stored in large quantities, this is why Red Eléctrica uses adjustment and balancing services, to guarantee that electricity supply always matches demand, even when real time conditions are different from the forecasts. These services are necessary to maintain the stability and security of the electrical grid [38].

According to [38], there are three types of interventions:

- Management of technical restrictions: to solve the network congestion for the following day or in real-time.
- Management of additional services: control of power and voltage.
- Imbalance management: solves in real time the imbalances between supply and demand.

Market participants providing balancing services are paid when they have to increase their electricity generation. If they are required to generate less, they will

have to pay for the energy not generated, which is always lower than the day-ahead market price [30].

Chapter 3

The dataset

The dataset provided by ENGIE contains historical information from the Spanish electricity market for previous days. The dataset covers the period from 15 June 2024 to 14 December 2025. The information contained in ENGIE's dataset is summarised in Table 3.1.

Variable	Unit	Description
price_da_es	€/MWh	Spanish day-ahead prices
price_da_fr	€/MWh	French day-ahead prices
price_da_pt	€/MWh	Portuguese day-ahead prices
price_ida1	€/MWh	Intraday prices
price_ida2	€/MWh	Intraday prices
price_ida3	€/MWh	Intraday prices
price_rrtt_faseI_up	€/MWh	Real-time reserves
pbf_X	MWh	Base dispatch schedule
p48_X	MWh	Actual production plan
pvp_X	MWh	Provisional feasible schedule
price_CO ₂	€/tCO ₂	CO ₂ price
mibgas	€/MWh	Gas market

Table 3.1: Dataset information.

All the data is provided at a 15-minute frequency. The day-ahead prices for Portugal and France are not used, since the project focuses exclusively on the Spanish electricity market. The variables that will be forecast are the Spanish day-ahead price and the intraday 1 price.

It is important to emphasise that all the data corresponding to the delivery day

is not available when the forecast is made. This is why the explanatory variables must be lagged in order to guarantee a model that is valid and does not use unavailable information. The forecasts for the day-ahead and intraday 1 markets for day D are made on day $D-1$ because the day-ahead market closes the day before the delivery date. Consequently, the day-ahead forecast cannot be made with any of the data from the day D . Since the objective of the project is to develop an arbitrage strategy between the two markets, the intraday 1 prediction has to also be available on day $D-1$ before the day-ahead market closes in order to be able to determine in which market it is better to sell or buy, or even if it is not profitable to trade.

The variable *price_rrtt_faseI_up* refers to the balancing market price. This market helps ensure the balance between supply and demand in real time.

The variable *pbf_X* corresponds to the base dispatch schedule. It is an image of the generation of each source after the day-ahead market.

The variable *pvp_X* corresponds to the provisional feasible schedule. It represents the scheduled generation of each energy source after technical constraints have been applied by the grid operator, Red Eléctrica.

The variable *p48_X* corresponds to the actual production plan. It is the final production of each energy source.

Table 3.2 contains the possible suffixes for the variables *pbf*, *p48* and *pvp*. These suffixes correspond to the different electricity generation technologies.

Suffix	Description	Example
<i>_wind</i>	Wind production	<i>pbf_wind</i>
<i>_solarpv</i>	Solar PV production	<i>pbf_solarpv</i>
<i>_solart</i>	Solar thermal production	<i>pvp_solart</i>
<i>_dam</i>	Hydro dam generation	<i>pvp_dam</i>
<i>_pump</i>	Pumped hydro generation	<i>pbf_pump</i>
<i>_nuclear</i>	Nuclear generation	<i>p48_nuclear</i>
<i>_gas</i>	Gas generation	<i>pvp_gas</i>
<i>_coal</i>	Coal generation	<i>pbf_coal</i>
<i>_ror</i>	Run of river generation	<i>p48_ror</i>
<i>_interco_*</i>	Cross-border flow	<i>pbf_interco_frsp</i>
<i>_demand</i>	Total demand	<i>pbf_demand</i>

Table 3.2: Suffix reference table

The suffix *interco* denotes the interconnection between Spain and France or be-

tween Spain and Portugal in either direction of flow. There are four possible interconnection variables.

For the forecasting models, the Spanish day-ahead and intraday 1 prices were used as target variables. For the day-ahead market forecast the *pbf* production variables were used, whereas the *pvp* production variables were used for the intraday 1.

3.1 Weather information

It is necessary to include weather forecasts for the delivery day. This is because in the dataset provided by ENGIE, none of the variables can be used without applying a lag. Otherwise, the model would be trained using information that is not available at the time of the prediction. However, the weather forecasts for day D issued the day before the delivery date is data that is available and therefore can be used.

REE (Red Eléctrica de España) provides information on installed generation capacity by technology and autonomous community on its website. This information is given for every month of each year since 2015.

3.1.1 Weighting coefficients

To simplify, instead of using the weather data from each energy generation plant, weighting coefficients were computed using the installed capacity of wind and solar photovoltaic in each autonomous community. The data is taken from [7]. It has been organised in an Excel file to obtain the proportion of installed power of each of these energy sources for every month from June 2024 to December 2025.

Only the peninsular autonomous communities are considered because these are the ones that participate in the day-ahead and intraday 1 markets.

In Table 3.3, the average proportions obtained for solar and wind generation are shown. In the appendices, the Tables 8.1, 8.2, 8.3 and 8.4 contain all the proportions for each month.

3.1.2 Weather forecast

The weather forecasts were obtained from [24]. The percentage of cloud cover and the wind speed (km/h) at $100m$ have been retained. This weather forecast is given two days before, which guarantees that this information will be available before the closure of the markets.

The latitude and longitude in each autonomous community were selected by observing the maps from [6], which contain the distribution of installed power in every municipality. For cloud cover, the installed power of photovoltaic energy has been observed and for the wind speed the installed power of wind energy. In the appendices, Table 8.5 contains the coordinates chosen for each weather location. The weather data covers the period from June 2024 to March 2026 with an hourly data point.

As the weather data is given hourly, it was linearly interpolated to obtain it with a 15-minute frequency.

Each weather variable has been multiplied by the corresponding monthly weighting coefficient. Cloud cover was weighted according to the installed photovoltaic capacity and wind speed by the installed wind capacity.

The day of the blackout, 28 April 2025, was removed as it is an extraordinary event.

Autonomous Community	Photovoltaic	Wind
Andalusia	0.241	0.116
Aragon	0.082	0.182
Asturias	0.001	0.022
Cantabria	0.001	0.001
Castile-La Mancha	0.200	0.156
Castile and León	0.096	0.232
Catalonia	0.039	0.044
Community of Madrid	0.018	0.000
Valencian Community	0.038	0.039
Extremadura	0.197	0.003
Galicia	0.006	0.125
La Rioja	0.005	0.015
Murcia	0.056	0.008
Navarre	0.011	0.051
Basque Country	0.007	0.005

Table 3.3: Average weighting coefficients of installed solar and wind power in the peninsular autonomous communities.

Chapter 4

Data analysis

To gain a better understanding of the dataset, an initial analysis of the data has been done. This will help to identify the main characteristics of the variables and therefore, help to choose the model parameters and to study the behaviour of each variable.

4.1 Descriptive statistics

The descriptive analysis provides information about the central tendency, variability, extreme values and distribution of the variables.

The mean, median, maximum and minimum values, standard deviation, skewness, and kurtosis were calculated for all variables. Skewness measures the asymmetry of a distribution, while kurtosis describes the tails relative to a normal distribution.

Table 4.1 contains the descriptive statistics for all the variables. They have been calculated using the entire dataset.

The day-ahead and intraday 1 prices have a skewness values close to zero, indicating that they have a symmetric distribution. Their kurtosis values are below 3, meaning that they have less pronounced tails than a normal distribution. The large difference between the mean and the median of solar photovoltaic generation shows the strong daily pattern of this type of energy, which is logical as it produces during daylight hours, and it cannot produce during the night. Pump and gas generation have high kurtosis values, indicating that they have a greater tendency to produce extreme values.

Variable	Mean	Median	Std	Min	Max	Skewness	Kurtosis
price_da_es	71.42	77.50	46.64	-15.00	250.00	-0.02	2.11
price_ida1_es	69.36	77.40	46.74	-42.90	244.32	-0.02	2.20
pbf_demand	26819.16	27010.6	4288.79	16617.3	39366.0	0.06	2.20
pvp_demand	26816.62	27006.20	4289.64	16617.30	39366.10	0.06	2.20
pbf_wind	7976.14	7118.10	4452.79	461.70	22622.50	0.72	2.88
pvp_wind	7197.25	6360.30	4180.75	268.30	21453.40	0.76	2.96
pbf_solarpv	6300.20	1210.10	7546.30	0.10	25620.40	0.77	2.08
pvp_solarpv	5850.35	809.00	7098.12	0.00	63186.80	0.80	2.35
pbf_solar	736.07	512.80	564.90	2.10	2192.70	1.19	3.16
pvp_solar	537.07	344.40	542.17	0.10	4965.70	1.12	3.37
pbf_nuclear	5823.21	6111.80	1416.96	50.00	8160.60	-1.28	4.15
pvp_nuclear	6003.01	6286.20	1185.30	675.00	8160.60	-1.46	5.80
pbf_dam	3695.93	3339.30	2231.23	352.30	10414.90	0.59	2.52
pvp_dam	3355.75	3008.10	2120.36	238.70	9531.40	0.65	2.57
pbf_pump	-1541.30	-1446.00	1171.02	-4379.70	-0.80	-0.35	1.97
pvp_pump	-1579.84	-1429.40	1234.83	-6743.90	-0.10	-0.34	1.85
pbf_gas	2191.13	1432.90	2333.86	0.80	14321.80	1.68	6.45
pvp_gas	2572.86	2373.10	1689.99	7.00	14123.80	1.95	9.41
pbf_coal	366.45	243.60	235.54	30.00	1310.00	1.83	6.17
pvp_coal	297.03	243.00	141.08	36.60	1330.00	2.54	13.04
pbf_ror	682.13	697.60	228.87	174.80	1219.80	0.01	1.95
pvp_ror	644.13	666.20	225.17	174.80	1174.90	-0.07	1.91
pbf_interco_spfr	-1874.60	-1913.90	972.51	-4024.00	-0.70	0.16	2.04
pvp_interco_spfr	-1873.91	-1908.80	973.77	-4024.00	-0.70	0.16	2.04
pbf_interco_frsp	2004.82	2249.30	910.39	1.50	3931.00	-0.34	2.34
pvp_interco_frsp	2004.59	2248.50	913.19	1.50	7373.20	-0.31	2.46
pbf_interco_sppt	-1865.47	-1784.50	1101.12	-7871.30	-0.20	-0.23	2.18
pvp_interco_sppt	-1870.71	-1789.50	1111.32	-11847.50	-0.20	-0.35	3.20
pbf_interco_ptsp	941.22	754.05	781.71	0.40	4410.00	1.10	3.88
pvp_interco_ptsp	943.10	755.90	782.68	0.40	4410.00	1.10	3.88
mibgas	37.06	35.03	6.71	26.56	57.05	0.80	2.80
price_co2	72.10	70.41	6.52	59.87	90.16	0.63	2.66
wind_weighted	20.44	19.29	8.36	3.59	57.81	0.74	3.42
solar_weighted	50.04	49.98	30.00	0.00	100.00	-0.01	1.71

Table 4.1: Descriptive statistics of the variables used in the study.

4.2 Time series visualization

Visualizing variables helps to understand their dynamics and to identify possible seasonal patterns. Figure 4.1 shows that the day-ahead and the intraday 1 prices are very similar. The figure covers the period from 1 April 2025 to 30 June 2025. The entire dataset is not displayed as this would give a figure that is difficult to read.




Figure 4.1: Day-ahead and Intraday 1 market prices.

4.2.1 Dispatchable and non-dispatchable energy sources

Energy sources can be classified as dispatchable and non-dispatchable depending on whether it is possible to control the production depending on the consumption needs.

Dispatchable energy sources are those whose production can be adjusted depending on the electric system needs. The time needed to change their production is short. For example, gas power plants are a dispatchable form of energy [9, 33].

Non-dispatchable energy sources are those whose production cannot be rapidly adjusted. However, it is possible to stop their production completely. In the case of renewable energy sources, their maximum production level depends on meteorological conditions, and in consequence, it is not fully controllable [21]. Renewable energies have low operating costs as they do not need fuel, that is why they produce to the maximum capacity while conditions allow it.

The impact of these two types of energy sources on the day-ahead and intraday 1 electricity market prices is analysed.

Non-dispatchable energies

The production of the non-dispatchable energy sources included in the dataset is shown in Figure 4.2. Demand is higher during the day than during the night. This

is consistent with consumption patterns. It also shows that the day-ahead market price tends to be lower during the day than during the night.




Figure 4.2: Non-dispatchable energy production, total demand and day-ahead market price




Figure 4.3: Percentage of non-dispatchable energy sources in the day-ahead market

Figure 4.2 also shows that solar photovoltaic production is higher during the day

and that solar thermal production follows the same pattern, but with lower variations. During the night, the production of these energy sources is zero, as expected. For wind production, it is generally higher during the night. These observations are general and strongly depend on meteorological conditions. They have also been observed in other days apart from Figure 4.2.

Figure 4.3 presents the share of each non-dispatchable energy source (the production of each energy source divided by the total production of the day-ahead market). It is observed that solar photovoltaic production represents the largest share during the day.

Dispatchable energy sources

Figure 4.4 shows the production of the dispatchable energy sources together with the evolution of the day-ahead market price.

The largest dispatchable energy source in terms of electricity generation is hydroelectric power. The negative values associated with pump generation correspond to the electricity consumed to pump the water into the reservoirs. It is also observed that hydroelectric generation tends to increase when market prices are higher.




Figure 4.4: Dispatchable energy production and day-ahead market price.

Figure 4.5 presents the share of each dispatchable energy source. It is observed that the shares of hydroelectric and nuclear generation tend to increase when prices are higher.




Figure 4.5: Share of dispatchable energy sources in total day-ahead electricity generation.

4.2.2 Comparison of dispatchable and non-dispatchable energy sources

The dispatchable and non-dispatchable electricity generation are compared to analyse their influence on electricity market prices. To analyse how their share of the production changes market prices.

Figure 4.6 shows that electricity generation from non-dispatchable energy sources is generally higher than that from dispatchable sources. Figure 4.7 shows that a decrease in the share of non-dispatchable generation is associated with an increase in the day-ahead market price. Conversely, when the share of non-dispatchable generation increases, the market price decreases. Higher shares of dispatchable generation tend to be associated with higher electricity prices, while higher shares of non-dispatchable generation are associated with lower prices.

When the sum of both percentages is different from 100%, this can be explained by electricity imports from or exports to Portugal and France. They have not been included in dispatchable or in non-dispatchable energy source because their generation mix is unknown and they are probably a mix of energy sources.




Figure 4.6: Dispatchable and non-dispatchable electricity generation and day-ahead market price.




Figure 4.7: Share of dispatchable and non-dispatchable electricity generation and day-ahead market price.

4.2.3 Intraday 1 market and the actual generation mix

Figures 4.8 and 4.9 show the generation of the dispatchable and non-dispatchable energy sources together with the evolution of the intraday 1 market price. The observed relationships are very similar to those found in the day-ahead market. The most important difference is that the data have a higher frequency. This is due to the market reform. This reform was applied to the day-ahead market on October 2025 and to the intraday 1 market on March 2025, introducing a 15-minute frequency.

As observed in the day-ahead market, wind production is higher during the night and solar production is higher during daylight hours. While solar follows a pattern with small changes in each period, wind generation exhibits greater variability. Solar photovoltaic production is the dominant non-dispatchable energy source during daylight hours.

An increase in dam generation is associated with higher prices. Peaks in hydropower generation tend to coincide with peaks in the electricity price. It is also observed that during night, gas generation increases.




Figure 4.8: Non-dispatchable energy production and intraday 1 market price.




Figure 4.9: Dispatchable energy production and intraday 1 market price.

4.3 Distribution analysis

The distribution analysis was used to observe if the day-ahead and the intraday 1 prices follow a normal distribution. Also, it was used to identify the most frequent price ranges

The histograms shown in Figures 4.10 and 4.11 indicate that there is no single dominant range of values. Both markets have a high frequency of prices close to 0 €/MWh. These low prices have also been observed in Figures from Section 4.2.1, particularly in Figures 4.2 and 4.7. These show that when there is a lot of solar production, which typically occurs during daylight hours and corresponds to a higher share of non-dispatchable energy production, the prices are close to 0 €/MWh.

The other concentration of prices is observed around 100 €/MWh. This probably occurs when the share of dispatchable energies increase.

This shows that the prices do not follow a normal distribution.




Figure 4.10: Day-ahead market price histogram.




Figure 4.11: Intraday 1 market price histogram.

4.4 Stationary analysis

The Augmented Dickey-Fuller test (ADF) determines if a series has a unit root. The test is based on Equation 4.1, where $\Delta y_{t-1} = y_{t-1} - y_{t-2}$ and ϵ_t is the error term. The coefficient ρ is the one evaluated to determine the presence of a unit root [8].

The test was performed using data from 15 June 2024 to 14 June 2025.

$$\Delta Y_t = \rho y_{t-1} + \zeta_1 \Delta y_{t-1} + \zeta_2 \Delta y_{t-2} + \cdots + \zeta_{p-1} \Delta y_{t-p+1} + \epsilon_t \quad (4.1)$$

In the ADF test, the null hypothesis is that the series is non-stationary [26]. As shown in Section 4.2, the prices do not have a clear tendency and their mean is different from zero. This is why a constant is added to the equation.

Table 4.2 presents the results of this test for both markets. For both of them, the statistics are below the critical values and the p-values are below 0.05, which indicates that both series can be considered stationary.

Series	ADF Statistic	p-value	CV 1%	CV 5%	CV 10%
price_da_es	-10.84	0.0000	-3.43	-2.86	-2.57
price_ida1_es	-11.17	0.0000	-3.43	-2.86	-2.57

Table 4.2: Augmented Dickey-Fuller test results for the DA and the IDA1 markets.

To complement the ADF test, the Kwiatkowski-Phillips-Schmidt-Shin test (KPSS) was also applied. This test inverts the hypothesis with respect to the ADF test, the null hypothesis assumes that the series is stationary [20]. Table 4.3 shows that the KPSS test rejects the null hypothesis of stationarity as the p-values are below 0.05. This could be due to seasonality, which indicates the possible need for seasonal differencing the series.

Series	KPSS Statistic	p-value	CV 1%	CV 5%	CV 10%
price_da_es	5.91	0.01	0.74	0.46	0.35
price_ida1_es	5.53	0.01	0.74	0.46	0.35

Table 4.3: KPSS Test Results for the DA and the IDA1 markets.

The ADF and KPSS tests were also applied to all the exogenous variables to study if they needed differencing. Table 4.4 contains the results of both tests. For all

the variables except for *wind_weighted*, the results of the tests lead to opposite conclusions. The variables were seasonally differenced. Table 4.5 contains the results of both tests after seasonal differencing. It shows that the differenced variables are stationary.

Variable	Stationary ADF	Stationary KPSS	Variable	Stationary ADF	Stationary KPSS
pbf_demand	Yes	No	pvp_demand	Yes	No
pbf_wind	Yes	No	pvp_wind	Yes	No
pbf_solarpv	Yes	No	pvp_solarpv	Yes	Nos
pbf_solart	Yes	No	pvp_solart	Yes	No
pbf_nuclear	Yes	No	pvp_nuclear	Yes	No
pbf_dam	Yes	No	pvp_dam	Yes	No
pbf_pump	Yes	No	pvp_pump	Yes	No
pbf_gas	Yes	No	pvp_gas	Yes	No
pbf_coal	Yes	No	pvp_coal	Yes	No
pbf_ror	Yes	No	pvp_ror	Yes	No
pbf_interco_spfr	Yes	No	pvp_interco_spfr	Yes	No
pbf_interco_frsp	Yes	No	pvp_interco_frsp	Yes	No
pbf_interco_sppt	Yes	No	pvp_interco_sppt	Yes	No
pbf_interco_ptsp	Yes	No	pvp_interco_ptsp	Yes	No
mibgas	No	No	price_co2	No	No
wind_weighted	Yes	Yes	solar_weighted	Yes	No

Table 4.4: Augmented Dickey-Fuller and KPSS test results for the exogenous variables.

Variable	Stationary ADF	Stationary KPSS	Variable	Stationary ADF	Stationary KPSS
pbf_demand	Yes	Yes	pvp_demand	Yes	No
pbf_wind	Yes	Yes	pvp_wind	Yes	No
pbf_solarpv	Yes	Yes	pvp_solarpv	Yes	No
pbf_solart	Yes	Yes	pvp_solart	Yes	No
pbf_nuclear	Yes	Yes	pvp_nuclear	Yes	No
pbf_dam	Yes	Yes	pvp_dam	Yes	Yes
pbf_pump	Yes	Yes	pvp_pump	Yes	Yes
pbf_gas	Yes	Yes	pvp_gas	Yes	Yes
pbf_coal	Yes	Yes	pvp_coal	Yes	Yes
pbf_ror	Yes	Yes	pvp_ror	Yes	Yes
pbf_interco_spfr	Yes	Yes	pvp_interco_spfr	Yes	Yes
pbf_interco_frsp	Yes	Yes	pvp_interco_frsp	Yes	Yes
pbf_interco_sppt	Yes	Yes	pvp_interco_sppt	Yes	Yes
pbf_interco_ptsp	Yes	Yes	pvp_interco_ptsp	Yes	Yes
mibgas	Yes	Yes	price_co2	Yes	Yes
wind_weighted	Yes	Yes	solar_weighted	Yes	Yes

Table 4.5: Augmented Dickey-Fuller and KPSS test results for the seasonal differenced exogenous variables.

4.5 Autocorrelation analysis

The autocorrelation analysis was performed for the day-ahead and intraday 1 market prices. As in the ADF test, the analysis was performed using data from 15 June 2024 to 14 June 2025. Figures 4.12 and 4.14 show the autocorrelation functions of both markets. Tables 4.6 and 4.8 contain the values of the autocorrelation for the 15-minutes, 30-minutes, 45-minutes, 1-day, 2-days and 3-day lag. A strong autocorrelation is observed at the 24 h lag in both markets, which suggest a daily seasonal period of $s=96$ observations.

This suggests the need for seasonally differencing the series. Tables 4.7 and 4.9 contain the autocorrelations of the seasonally differenced series. The results show that the differentiation reduces the seasonal autocorrelation.

Lag	Autocorrelation
15 min	0.987
30 min	0.973
45 min	0.960
24 h	0.752
48 h	0.589
72 h	0.505

Table 4.6: Autocorrelation results for the day-ahead market

Lag	Autocorrelation
15 min	0.971
30 min	0.943
45 min	0.914
24 h	-0.172
48 h	-0.159
72 h	-0.0402

Table 4.7: Autocorrelation results for the day-ahead market price after seasonally differencing.




Figure 4.12: Autocorrelation Function of the day-ahead market price




Figure 4.13: Autocorrelation Function of the day-ahead market price after seasonally differencing.

Lag	Autocorrelation
15 min	0.986
30 min	0.972
45 min	0.958
24 h	0.750
48 h	0.588
72 h	0.517

Table 4.8: Autocorrelation results for the intraday 1 market price.

Lag	Autocorrelation
15 min	0.969
30 min	0.938
45 min	0.910
24 h	-0.177
48 h	-0.181
72 h	-0.031

Table 4.9: Autocorrelation results for the intraday 1 market after seasonally differencing.




Figure 4.14: Autocorrelation Function of the intraday 1 market price.

The Partial Autocorrelation Function (PACF) measures the correlation between




Figure 4.15: Autocorrelation Function of the intraday 1 market price after seasonally differencing.

a series and its lagged values after eliminating the influence of intermediate lags. Tables 4.10 and 4.11 contain the results of the PACF for the seasonally differenced series. For both markets, the partial autocorrelation is high at the 15-minute lag and decreases for the next lags.

Lags	Partial Autocorrelation
15 min	0.971
30 min	-0.014
45 min	-0.015
24 h	-0.015
48 h	-0.009

Table 4.10: Partial Autocorrelation Function of the day-ahead market price.

Lags	Partial Autocorrelation
15 min	0.970
30 min	-0.022
45 min	0.020
24 h	-0.041
48 h	-0.016

Table 4.11: Partial Autocorrelation Function of the intraday 1 market price.

4.6 Correlation analysis

The Pearson correlation matrices were calculated for the variables used in both markets. Figure 4.16 and Figure 4.17 show the correlation coefficients between the variables. The first variable shown (top left corner) corresponds to the target variable, to the rest of the variables, a 24-hour lag was applied as their actual value will not be available at the forecasting time. Only the *wind_weighted* and *solar_weighted* variables were not lagged as they are obtained from weather forecasts.

As discussed in Section 4.5, a strong correlation was observed with the 24-hour lag. The correlation with the 15-minute, 30-minute and 45-minute is higher, but this data will not be available at the moment of the forecast. Therefore the variable *price_da_es_lag* has been added to the correlation matrix for the day-ahead market and the *price_ida1_es_lag* to the correlation matrix of the intraday 1 market.

The variable *price_da_es_lag* has a strong correlation with both target variables.




Figure 4.16: Correlation matrix for the day-ahead market.




Figure 4.17: Correlation matrix for the intraday 1 market.

As a combination of several variables can be correlated with another variable, the correlation matrix is not enough to eliminate the variables. In order to verify that there is no multicollinearity between the variables, the Variance Inflation Factor (VIF) was calculated. The Equation for the VIF of variable X_j is in Equation 4.2 obtained from [36]. The VIF measures how much the variance of an estimated regression coefficient of a certain variable is inflated due to multicollinearity with other explanatory variables [36].

$$VIF = \frac{1}{1 - R_j^2} \quad (4.2)$$

Tables 4.12, and 4.13 contain the VIF for the explanatory variables of each market. The variables with a VIF higher than 10 have been eliminated to avoid problems of multicollinearity. Tables 4.14 and 4.15 contain the VIF for the final set of selected variables. The variables have been eliminated one by one until the remaining variables had VIF values below 10.

Variable	VIF
<i>price_da_es_lag</i>	109.68
<i>price_ida1_es_lag</i>	57.23
<i>pbf_demand_lag</i>	13576.76
<i>pbf_wind_lag</i>	4177.01
<i>pbf_solarpv_lag</i>	276.23
<i>pbf_solart_lag</i>	32.06
<i>pbf_nuclear_lag</i>	947.13
<i>pbf_dam_lag</i>	139.64
<i>pbf_pump_lag</i>	41.27
<i>pbf_ror_lag</i>	320.37
<i>pbf_gas_lag</i>	137.07
<i>pbf_coal_lag</i>	89.39
<i>pbf_interco_spfr_lag</i>	33.63
<i>pbf_interco_frsp_lag</i>	98.71
<i>pbf_interco_sppt_lag</i>	81.12
<i>pbf_interco_ptsp_lag</i>	3.29
<i>mibgas_lag</i>	1155.17
<i>price_co2_lag</i>	1441.99
<i>wind_weighted</i>	30.62
<i>solar_weighted</i>	10.36

Table 4.12: VIF for the day-ahead market variables before removing collinear variables.

Variable	VIF
<i>price_da_es_lag</i>	71.52
<i>price_ida1_es_lag</i>	43.93
<i>pvp_demand_lag</i>	3686.64
<i>pvp_wind_lag</i>	935.06
<i>pvp_solarpv_lag</i>	22.81
<i>pvp_solart_lag</i>	4.61
<i>pvp_nuclear_lag</i>	261.26
<i>pvp_dam_lag</i>	142.93
<i>pvp_pump_lag</i>	13.09
<i>pvp_ror_lag</i>	99.87
<i>pvp_gas_lag</i>	80.38
<i>pvp_coal_lag</i>	22.32
<i>pvp_interco_spfr_lag</i>	22.34
<i>pvp_interco_frsp_lag</i>	23.33
<i>pvp_interco_sppt_lag</i>	27.11
<i>pvp_interco_ptsp_lag</i>	2.76
<i>mibgas_lag</i>	388.02
<i>price_co2_lag</i>	569.16
<i>wind_weighted</i>	16.26
<i>solar_weighted</i>	6.86

Table 4.13: VIF for the intraday market variables before removing collinear variables.

Variable	VIF
<i>price_da_es_lag</i>	8.53
<i>pbf_solart_lag</i>	2.95
<i>pbf_dam_lag</i>	4.29
<i>pbf_pump_lag</i>	2.54
<i>pbf_gas_lag</i>	2.25
<i>pbf_interco_spfr_lag</i>	2.93
<i>pbf_interco_frsp_lag</i>	5.14
<i>pbf_interco_ptsp_lag</i>	1.13
<i>wind_weighted</i>	6.24
<i>solar_weighted</i>	4.55

Table 4.14: VIF for the day-ahead market variables after removing collinear variables.

Variable	VIF
<i>price_ida1_es_lag</i>	3.56
<i>pvp_solart_lag</i>	4.16
<i>pvp_dam_lag</i>	4.99
<i>pvp_pump_lag</i>	5.68
<i>pvp_gas_lag</i>	3.14
<i>pvp_interco_spfr_lag</i>	3.56
<i>pvp_interco_frsp_lag</i>	3.12
<i>pvp_interco_sppt_lag</i>	6.49
<i>pvp_interco_ptsp_lag</i>	1.39
<i>wind_weighted</i>	6.33
<i>solar_weighted</i>	4.24

Table 4.15: VIF for the intraday 1 market variables after removing collinear variables.

Chapter 5

Electricity Price Forecasting

Three different models were studied for electricity price forecasting. These models were proposed by the project supervisor. Two of them are classical forecasting models, Linear Regression and SARIMAX, while the third one is a machine learning model, LightGBM.

5.1 Evaluation Metrics

To compare the performance of the three forecasting models, seven evaluation metrics were considered, the Mean Absolute Error (MAE), the Mean Absolute Percentage Error (MAPE), the symmetric Mean Absolute Percentage Error (sMAPE), the Root Mean Squared Error (RMSE), the coefficient of determination (R^2), the Bayesian Information Criterion (BIC) and the Akaike Information Criterion (AIC).

In equations 5.1, 5.4 and 5.5, y_i denotes the observed value, $\hat{y}_i$ is the predicted value, $\bar{y}$ the mean of the observed values and n is the number of observed values.

- MAE: The MAE assigns the same weight to all errors as it is the average magnitude of all the errors between the observed value and the prediction [11].

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (5.1)$$

- MAPE: The Mean Absolute Percentage error (MAPE) main advantage is that is scale-independent. The problem is that it is undefined when the real value is zero [11].

$$MAPE = \frac{100}{n} \sum_{i=1}^n \frac{|y_i - \hat{y}_i|}{y_i} \quad (5.2)$$

- The sMAPE is a modified version of the MAPE. [11]:

$$sMAPE = \frac{1}{n} \sum_{i=1}^n \frac{200|y_i - \hat{y}_i|}{|y_i| + |\hat{y}_i|} \quad (5.3)$$

- RMSE: The RMSE penalises big prediction errors more heavily [11].

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (5.4)$$

- R^2 : The coefficient of determination measures how well the model explains the variance ([17], p 79).

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (5.5)$$

- BIC: The BIC measures the quality of a model while penalising its complexity. It provides a trade-off between the number of parameters (model complexity) and the goodness of fit. A lower BIC indicates a better model. A difference in BIC between 5 and 10 is considered significant, and above 10 it is considered very significant [18].
- AIC: The AIC measures penalizes less heavily than the BIC the number of parameters [10].

5.2 Linear Regression

The linear regression models the relation between a dependent variable and one or more independent variables. This is a classical model in which the target variable is explained as a linear combination of the explanatory variables.

Equation 5.6 shows how Y is expressed as a combination of X_1 , X_2 and X_3 . The coefficients β are obtained by the least squares method, which consists of minimizing the sum of the squared residuals.

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 + \epsilon \quad (5.6)$$

In linear regression, it is positive having high correlation between the objective variable and the explanatory variables. However, having high correlation between the explanatory variables should be avoided. Multicollinearity can result in a difficulty to establish the coefficients.

5.2.1 Forecast

The models were trained using the dataset and evaluated over the period from 15 June 2025 to 13 December 2025. Tables 5.5 and 5.6 contain the coefficients and p-values of the linear regression models for the day-ahead market. Table 5.5 shows that *pbfgas_lag* has a p-value greater than 0.05, so this variable is eliminated.

Tables 5.8 and 5.9 contain the coefficients and p-values of the linear regression of the variables for the intraday 1 market. In Table 5.8 it is observed that *pvp_interco_spfr_lag* has a p-value greater than 0.05, so this variable is eliminated.

Before removing variables, the residuals were analysed. Several tests were performed. The Jarque-Bera test was applied to analyse if the residuals followed a normal distribution. The results shown in Table 5.1 indicate that they are not normal. The ADF and the KPSS tests were performed to evaluate if they are stationary, in Table 5.2 it is observed that they the results are inconclusive. The Breusch-Pagan test determines whether a series has heteroscedasticity or not, Table 5.3 contains the results, it shows that residuals do have heterocedasticity. This is also seen in Figures 8.7 and 8.8. The Ljung-Box test was applied to determine if there is autocorrelation in the residuals, Table 5.4 shows that there is autocorrelation. This is also seen in Figures 8.1, 8.3, 8.2 and 8.4. Figures 8.5 and 8.6 do not show a clear temporal trend in the residuals.

Market	p-value	Normal distribution
DA before	0.000	No
DA after	0.000	No
IDA1 before	0.000	No
IDA1 after	0.000	No

Table 5.1: Jarque-Bera test of the residuals of the regression model.

Market	ADF p-value	KPSS p-value	Stationary
DA before	0.000	0.01	Inconclusive
DA after	0.000	0.01	Inconclusive
IDA1 before	0.000	0.01	Inconclusive
IDA1 after	0.000	0.01	Inconclusive

Table 5.2: ADF and KPSS tests of the residuals of the regression model.

Market	p-value	Heteroscedasticity
DA before	0.000	Yes
DA after	0.000	Yes
IDA1 before	0.000	Yes
IDA1 after	0.000	Yes

Table 5.3: Breusch-Pagan test of the residuals of the regression model.

Market	lag	p-value	Autocorrelation
DA	1	0.0	Yes
DA	2	0.0	Yes
DA	24	0.0	Yes
DA	96	0.0	Yes
IDA1	1	0.0	Yes
IDA1	2	0.0	Yes
IDA1	24	0.0	Yes
IDA1	96	0.0	Yes

Table 5.4: Ljung-Box of the residuals of the DA and IDA1 regression models.

As there is autocorrelation and heteroscedasticity in the residuals, and they do not follow a normal distribution, the assumptions of linear regression are not verified. The p-values are not reliable. To remove the variables the metrics were observed. If removing a variable improved the model, it was removed, if not, it was retained. Tables 5.5 and 5.8 contain the initial variables, the ones selected in Section 4 based on the VIF. Tables 5.6 and 5.9 contain the final set of variables.

Tables 5.7 and 5.10 contain the metrics of the linear regression for the initial and the final set of variables for both markets. There is a slight improvement after removing some variables.

	Coef.	t	p_value
Constant	73.27	515.50	0.000
<i>price_da_es_lag</i>	39.66	163.31	0.000
<i>pbf_solar_lag</i>	-7.47	-41.16	0.000
<i>pbf_dam_lag</i>	-2.67	-14.15	0.000
<i>pbf_pump_lag</i>	-5.01	-22.38	0.000
<i>pbf_gas_lag</i>	0.35	2.01	0.052
<i>pbf_interco_spfr_lag</i>	-0.95	-4.68	0.000
<i>pbf_interco_frsp_lag</i>	-0.86	-4.15	0.000
<i>pbf_interco_ptsp_lag</i>	0.31	1.86	0.062
<i>solar_weighted</i>	-1.28	-8.19	0.000
<i>wind_weighted</i>	-11.11	-74.23	0.000

Table 5.5: Regression coefficients for the day-ahead market.

	Coef.	t	p_value
Constant	73.28	515.47	0.000
<i>price_da_es_lag</i>	39.91	187.00	0.000
<i>pbf_solar_lag</i>	-7.50	-41.39	0.000
<i>pbf_dam_lag</i>	-2.52	-14.34	0.000
<i>pbf_pump_lag</i>	-5.10	-23.26	0.000
<i>pbf_interco_spfr_lag</i>	-0.99	-4.87	0.000
<i>pbf_interco_frsp_lag</i>	-0.88	-4.28	0.000
<i>solar_weighted</i>	-1.28	-8.22	0.000
<i>wind_weighted</i>	-11.15	-75.03	0.000

Table 5.6: Regression Coefficients for the day-ahead market after variable selection.

Model	RMSE	MAE	MAPE	sMAPE	R^2	BIC	AIC
DA test, initial	25.92	19.48	inf	50.41	0.637	328,481.53	328,388.46
DA test, final	25.91	19.46	inf	50.38	0.637	328,467.25	328,391.10
DA training, final	26.57	19.89	inf	54.34	0.71		

Table 5.7: Linear Regression results for the day-ahead market.

	Coef.	t	p_value
Constant	70.15	413.17	0.000
<i>price_ida1_es_lag</i>	32.44	130.08	0.000
<i>pvp_solart_lag</i>	-5.60	-24.70	0.0000
<i>pvp_dam_lag</i>	-1.13	-4.63	0.000
<i>pvp_pump_lag</i>	1.80	6.74	0.000
<i>pvp_gas_lag</i>	0.61	3.28	0.001
<i>pvp_interco_spfr_lag</i>	-0.04	-0.15	0.883
<i>pvp_interco_frsp_lag</i>	0.74	2.98	0.003
<i>pvp_interco_sppt_lag</i>	-1.52	-6.01	0.000
<i>pvp_interco_ptsp_lag</i>	1.08	5.26	0.000
<i>solar_weighted</i>	0.49	2.65	0.008
<i>wind_weighted</i>	-11.22	-62.37	0.000

Table 5.8: Regression coefficients for the intraday 1 market.

	Coef.	t	p_value
Constant	70.15	413.18	0.000
<i>price_ida1_es_lag</i>	32.43	130.30	0.000
<i>pvp_solart_lag</i>	-5.60	-24.77	0.000
<i>pvp_dam_lag</i>	-1.12	-4.66	0.000
<i>pvp_pump_lag</i>	1.80	6.77	0.000
<i>pvp_gas_lag</i>	0.61	3.28	0.001
<i>pvp_interco_frsp_lag</i>	0.72	3.60	0.000
<i>pvp_interco_sppt_lag</i>	-1.52	-6.00	0.000
<i>pvp_interco_ptsp_lag</i>	1.08	5.29	0.000
<i>solar_weighted</i>	0.49	2.65	0.008
<i>wind_weighted</i>	-11.22	-62.40	0.000

Table 5.9: Regression Coefficients for the intraday 1 after variable selection.

Model	RMSE	MAE	MAPE	sMAPE	R^2	BIC	AIC
IDA1 test, initial	34.80	24.93	inf	64.59	0.384	340,907.55	340,806.02
IDA1 test, final	34.79	24.92	inf	64.58	0.384	340,897.11	340,804.04
IDA1 training, final	31.73	23.43	inf	63.26	0.599		

Table 5.10: Linear Regression results for the intraday 1 market.

5.3 SARIMAX

The SARIMAX model makes it possible to describe the evolution of prices by integrating the temporal dependencies of the price that is being predicted and the effect of exogenous variables. In our particular model, the endogenous variable corresponds to the price of the day-ahead and the intraday market while the exogenous variable corresponds to the productions of electricity of the different energy sources and the prices of gas and CO_2 emission.

5.3.1 SARIMAX model

The SARIMAX model is defined by the parameters $(p, d, q)(P, D, Q, s)$. Each of these parameters adds a term to the model equation. It is necessary to choose the parameters according to the dataset in order to create a model that adjusts to the time series without overfitting.

-The parameter p represents the order of the non seasonal autoregression component (AR) [2]. In other words, it indicates how many past values of the series are used in the prediction of the current value. In this project, it corresponds to the number of previous observations used for the prediction of the current market price. This part of the equation is written as in [23]:

$$\Phi_p(G) = 1 - \Phi_1 G - \Phi_2 G^2 - \dots - \Phi_p G^p \quad (5.7)$$

-The parameter d corresponds to the degree of non seasonal differencing (I) [2], which indicates the number of times the price series has to be differenced to become stationary.

-The parameter q represents the order of the not seasonal moving average (MA). It makes it possible to take into account the effect of random shocks, for example the impact of unexpected variations of the solar and wind production on the evolution of the energy price [2]. This part of the equation is written as in [23]:

$$\gamma_q(G) = 1 - \gamma_1 G - \gamma_2 G^2 - \dots - \gamma_q G^q \quad (5.8)$$

The ARIMA model can be written as in [23]:

$$\Phi_p(G)(1 - G)^d X_t = \gamma_q(G)e_t \quad (5.9)$$

-The parameter P corresponds to the order of the seasonal auto-regressive component(SAR) [23]. In this project, it introduces the influence of prices observed at

the same quarter hour of the day or of the days before. This makes it possible to capture the cyclical behaviour of the energy markets. This part of the equation is written as in [23]:

$$\Phi_P(G^s) = 1 - \Phi_1 G^s - \Phi_2 G^{2s} - \dots - \Phi_P G^{Ps} \quad (5.10)$$

-The parameter D is the degree of seasonal differencing, which eliminate the repetitive tendencies due to seasonality. For example, in this project, it is in relation with the daily consumption profiles.

-The parameter Q corresponds to the order of the seasonal moving-average component (SMA) [23]. It makes it possible to take into account the effect of random shocks in a cycle. This part of the equation is written as in [23]:

$$\gamma(G^s_Q) = 1 - \gamma_1 G^s - \gamma_2 G^{2s} - \dots - \gamma_Q G^{Qs} \quad (5.11)$$

-The parameter s corresponds to the length of the season. It indicates how many time steps a cycle has.

The SARIMA model is written as in [23]:

$$\Phi_p(G)\Phi_P(G^s)(1-G)^d(1-G^s)^D X_t = \gamma_q(G)\gamma_Q(G^s)e_t \quad (5.12)$$

When exogenous variables are added to the SARIMA model, it is called SARIMAX, the X corresponds to the variables objective variable. This model is written as in [23]:

$$\Phi_p(G)\Phi_P(G^s)(1-G)^d(1-G^s)^D X_t = \alpha_k y_{k,t} + \gamma_q(G)\gamma_Q(G^s)e_t \quad (5.13)$$

In the equation 5.13, $y_{k,t}$ corresponds to the exogenous component, k is the index of the exogenous variable and α_k corresponds to the coefficient associated with the *exogenous_k*. The variable t indicates the position of the series that is being predicted.

5.3.2 SARIMAX parameters

The determination of the SARIMAX parameters has been made based on the results of the Data Analysis Chapter 4, specially using the ADF test 4.4 and the ACF and PACF 4.5.

The choice of the parameter d has been made with the Augmented Dickey-Fuller test (ADF). Table 4.2 presents the results of this test in both markets. For both

of them, the statistics are negative and the p-values are lower than 0.05, which proving evidence that the series is stationary. This means that the parameter d should be 0.

Parameter D is chosen based on the observation of a seasonality characterised by repetitive peaks that can vary in amplitude depending on the day. Tables 4.6 and 4.8 contain the values of the autocorrelation for the 24-hour, 48-hour and 72-hour lags. It is observed that the values of the autocorrelation are important. This indicates that the series might need seasonal differencing and in consequence that the parameter D is not 0. Tables 4.7 and 4.9 contain the values of the autocorrelation after differentiating both series with a 24-hour lag. It is observed that the values of the autocorrelation decrease significantly. It can be concluded that the 24-hour lag difference eliminates the majority of the seasonal non-stationarity. This is why the parameter D is set to 1.

It remains to determine the parameters p , P , q and Q . The parameter d has been set to 0, D to 1 and s to 96 due to the daily seasonality of the series.

To determine the values of p and P , the Partial Autocorrelation Function (PACF) has been used. In Tables 4.10 and 4.11, it is observed that PACF is high at the first lag, but it drops for the second one for both day-ahead and intraday markets. This indicates that the parameter p should be set to 1 for both markets, it is enough to use the last past value.

The value of P will be determined using SARIMA models because the partial autocorrelation for the 24-hour lag is very low. Values 0 and 1 will be tried.

To determine the parameters q and Q , the Autocorrelation Function has been used. Tables 4.7 and 4.9 show the values of the autocorrelation for both differentiated series. The decrease of the autocorrelation for both markets is progressive for the first three lags. This progressive decrease is observed in Figures 4.13 and 4.15 for larger lags. This suggests that different values for q should be tried. For the 24-hour, 48-hour and 72-hour lags, the autocorrelation is very small, so values 1 and 0 have been tried for Q .

Table 5.11 contains the parameters that will initially be tried.

Parameter	DA	IDA1
p	1	1
d	0	0
q	0/1	0/1
P	0/1	0/1
D	1	1
Q	0/1	0/1
s	96	96

Table 5.11: SARIMAX possible parameters for the day-ahead and the intraday 1 markets.

Tables 5.12 and 5.13 present the results of the SARIMAX model for both markets using different parameter combinations. The results are not as good as expected and the models perform worse than the linear regression models. This could be due to increased complexity when adding the parameters that model the time dependence. Another possible explanation is an error in the model implementation. As it was not possible to improve these results, the SARIMAX model was discarded.

Parameters	Stationary	Autocorrelation	Normal	Heteroscedasticity	R^2	AIC	BIC
(1,0,0)(0,0,0,96)	Yes	Yes	No	Yes	0.4806	10092.127	10150.826
(0,0,0)(0,1,0,96)	Yes	Yes	No	Yes	-6.757	13452.294	13505.011
(1,0,0)(0,1,0,96)	Yes	Yes	No	Yes	-1.477	9807.825	9865.814
(1,0,1)(0,1,0,96)	Yes	Yes	No	Yes	-1.490	9803.421	9866.673

Table 5.12: Results of the SARIMAX model of the day-ahead market with different parameters.

Parameters	Stationary	Autocorrelation	Normal	Heteroscedasticity	R^2	AIC	BIC
(1,0,0)(0,0,0,96)	Yes	Yes	No	Yes	0.1839	10764.851	10790.665
(0,0,0)(0,1,0,96)	Yes	Yes	No	Yes	-21.4385	13883.720	13946.980
(1,0,0)(0,1,0,96)	Yes	Yes	No	Yes	-7.1428	10103.717	10172.249
(1,0,1)(0,1,0,96)	Yes	Yes	No	Yes	7.2899	10098.841	10172.635

Table 5.13: Results of the SARIMAX model of the intraday 1 market with different parameters.

5.4 LightGBM

Light GBM (Lightweight Gradient Boosting Machines) is an algorithm developed by Microsoft. This algorithm is based on decision trees.

5.4.1 Tree-based methods

Tree-based methods consist of partitioning the predictor space into smaller regions. The predictor space contains all the variables (in this project the prices and volume of production of each energy source). For a given observation, the forecast is made with the mode or the mean of the specific region to which it belongs to. For example, if for a specific quarter-hour, the volume of wind energy production is greater than a value X , the algorithm will search in the corresponding region. [17]

All the rules to divide the space of variables can be represented with a tree [17]. Each terminal region is a tree leaf and all leaves form a partition of the predictor space. A leaf is a region that does not divide into more regions. Figure 5.1 is an example of a tree, there are two variables (X_1 and X_2) and four conditions. Each condition divides each region of the variables space into two other regions. Once all the rules have been applied, there are five leaves, as it can be observed in Figure 5.2, these regions are disjoint and their union is the complete predictor space [17]. The shapes of the rules can be different from rectangles, to maintain a simple example, the rectangular regions have been chosen.




Figure 5.1: Example of a tree. [17]




Figure 5.2: Division of the predictor following the example tree from Figure 5.1

5.4.2 Gradient Boosting

Ensemble learning consists of using several simple models which are called *weak models* and combine them all to obtain better predictions. There are different types of ensemble methods, such as bagging, boosting and Bayesian additive regression trees [17].

The problem of trees is that they have a big variance, small changes in the dataset may change significantly the forecasts. Bagging consists of dividing the data set in order to create several trees and combine their predictions to reduce the variance. The combination of the predictions is done doing the average. It can happen that the dataset is not big enough to divide it, in this cases the datasets generated can repeat the samples[17].

Random forest is similar to bagging, but instead of using all the predictors in each split, it uses only some of them. This modification regarding bagging, reduces the correlation between trees. If there is a very strong predictor, all of the trees will be very similar and therefore, averaging their predictions will not reduce as much the variance. In random forest, as not all the splits take into account all the variables, not all the splits will be determined by the strongest predictor and the trees will be less correlated [17].

Boosting consists of growing each tree in an ordered way instead of growing them all at the same time and independently as in bagging and in random forest. Each tree will grow depending on the previous ones. The trees learn from the residuals of the other trees instead of learning from the prediction. This means that boosting is a slow learning method in which each tree corrects the mistakes of the previous ones. It uses several parameters like the number of trees, the number of leaves and the learning rate of the trees. Having a very small learning rate requires a larger number of trees to achieve a good performance. Having too many trees may lead to overfitting [17].

Bayesian additive regression trees combines them all in one construction of trees of bagging and random forest with the growth of the trees depending on the residuals of boosting. This method updates all trees during the iterations. They start having single leaf and depending on the signal that has not been captured, it modifies the trees [17].

5.4.3 LightGBM

LightGBM is a tree-based gradient boosting method, a Gradient Boost Decision Tree (GBDT). It uses the boosting method to build the trees, particularly it uses Gradient-based One-side Sampling (GOSS) [19].

Every instance is associated with a gradient, a smaller gradient means that it is already well trained and a larger gradient means that it needs more training to improve accuracy. The instances with smaller gradient cannot be deleted because this would change the data and the model would be biased. GOSS selects the $a \times 100\%$ instances with the larger gradient and then randomly selects the $b \times 100\%$ with smaller gradient. The gradient of the smaller gradient instances that are left are multiplied by $\frac{a-1}{b}$ [19].

To split the trees, GBDT searches to maximise the gain, which is expressed in Equation 5.14 from [19]. In this equation n_O is the number of instances of the node that is split, x_i the instance vector, $n_{r|O}^j(d)$ and $n_{l|O}^j(d)$ the right and left nodes resulting from the split and g_i the gradients [19].

$$[H]V_{j|O}(d) = \frac{1}{n_O} \left(\frac{\left(\sum_{\{x_i \in O: x_{ij} \leq d\}} g_i \right)^2}{n_{l|O}^j(d)} + \frac{\left(\sum_{\{x_i \in O: x_{ij} > d\}} g_i \right)^2}{n_{r|O}^j(d)} \right) \quad (5.14)$$

Parameters

LightGBM has many parameters, only the most important ones will be explained in this list (they have been selected from [25]):

1. `n_estimators`: the number of trees.
2. `learning_rate`: It determines how much each tree contributes to the prediction. It set by default to 0.1.
3. `num_leaves`: The maximum number of leaves that each tree can have. It is set by default to 31.
4. `max_depth`: The maximum depth of each tree. The depth is the distance between the initial node and the farthest leaf. It is set to 1 by default.
5. `min_data_in_leaf`: The minimum data in each leaf. This is used to control overfitting. It is set to 1 by default.
6. `bagging_fraction`: The percentage of the data used in each tree. To use it, `bagging_freq` should be different from 0.
7. `bagging_freq`: Every how many iterations the part of the dataset used is changed.
8. `reg_alpha`: Applies L1 regularization.
9. `reg_lambda`: Applies L2 regularization.
10. `metric`: The metric used to evaluate the model. There are many options like the mean absolute error or the mean squared error.

5.4.4 Forecast

To do the forecast it is necessary to choose the hyperparameters and the exogenous variables that are used.

The hyperparameters are chosen with the optimization software, Optuna. Instead of trying parameters randomly or trying all possible combinations, which takes a lot of time, it uses an objective function and evaluates this function, each evaluation is called trial. Based on previous trials, the next hyperparameters are chosen [1].

The most important hyperparameters have been tried. The range of values has been selected based on recommendations from [25] and they have been adjusted when their best value was near one of the extremes of the interval.

The variables selected were the same as those in the other two models. Three additional variables were added: *day_of_week*, *month* and *quarter*. The impor-

tance of the variables based on gain is shown in Figures 5.3 and 5.5. The least important values were removed one by one until removing additional variables no longer improved the model performance. Figures 5.4 and 5.6 show the importance of the selected variables.




Figure 5.3: Importance of the variables for the day ahead market




Figure 5.4: Importance of the variables for the day ahead market after eliminating.




Figure 5.5: Importance of the variables for the intraday 1 market




Figure 5.6: Importance of the variables for the intraday 1 market after eliminating.

The Jarque-Bera test was applied to analyse if the residuals followed a normal distribution. The results shown in Table 5.14 indicate that they are not normal. The ADF and the KPSS tests were performed to evaluate if they are stationary, in Table 5.15 it is observed that they are stationary. The Breusch-Pagan test determines whether a series has heteroscedasticity or not, Table 5.16 contains the results, it shows that residuals do have heteroscedasticity. This is also seen in

Figures 8.13 and 8.14. The Ljung-Box test was applied to determine if there is autocorrelation in the residuals, Table 5.17 shows that there is autocorrelation. This is also seen in Figures 8.9, 8.11, 8.10 and 8.12. Figures 8.15 and 8.16 do not show a clear temporal trend in the residuals.

Market	p-value	Normal distribution
DA	0.000	No
IDA1	0.000	No

Table 5.14: Jarque-Bera test of the residuals of the lightGBM model.

Market	ADF p-value	KPSS p-value	Stationary
DA	0.000	0.600	Yes
IDA1	0.000	0.1	Yes

Table 5.15: ADF and KPSS tests of the residuals of the lightGBM model.

Market	p-value	Heteroscedasticity
DA	0.000	Yes
IDA1	0.000	Yes

Table 5.16: Breusch-Pagan test of the residuals of the lightGBM model.

Market	lag	p-value	Autocorrelation
DA	1	0.0	Yes
DA	2	0.0	Yes
DA	24	0.0	Yes
DA	96	0.0	Yes
IDA1	1	0.0	Yes
IDA1	2	0.0	Yes
IDA1	24	0.0	Yes
IDA1	96	0.0	Yes

Table 5.17: Ljung-Box of the residuals of the DA and IDA1 lightGBM models.

Table 5.18 contains the results of the LightGBM model. The elimination of some of the variables improves the results. There is a slight overfitting, but it could not be further reduced without compromising the model's performance in the test set.

Model	RMSE	MAE	MAPE	sMAPE	R^2
DA test, initial	24.66	18.31	inf	48.18	0.670
DA test, final	22.69	19.12	inf	47.04	0.721
DA train, final	18.56	14.41	inf	42.93	0.858
IDA1 test, initial	27.21	20.21	inf	54.34	0.587
IDA1 test, final	26.87	19.36	inf	54.51	0.617
IDA1 train, final	24.23	17.49	inf	51.74	0.765

Table 5.18: Results LightGBM

Figures 5.7 and 5.8 show the predicted and actual prices of the day-ahead and the intraday 1 market prices. For some of the observations they follow the real prices, but for others they do not. Especially, when there are shocks, the model does not follow as well the prices as for the rest of the observations.




Figure 5.7: LightGBM forecast compared to the real day-ahead price.




Figure 5.8: LightGBM forecast compared to the real intraday 1 price.

Chapter 6

Arbitrage strategy

The objective of an arbitrage strategy is to take advantage of the price differences between different markets to generate a profit. In this project, it is because of the price difference between the day-ahead market and the intraday 1 market, that a profit can be made. In the arbitrage strategy developed, there will exist a risk and a capital requirement for the strategy. This is due to prediction errors and the different market closure times.

Several strategies have been implemented. The simpler ones have been implemented first to evaluate their performance and to be able to compare the more complex ones with them. The strategies were implemented in Python. The training period goes from 15 June 2024 to 14 June 2025, the test from 15 June 2025 to 14 December 2025.

To evaluate the strategies, the metrics that have been used are: Profit and Loss, volatility, value at risk, Sharpe ratio, Sortino ratio, maximum Drawdown, Profit Factor and Win Rate.

Profit and Loss (PnL) is the difference between the profits of the strategy and the losses. In this project, it will be the difference between the selling and the buying prices. The selling price and the buying prices will depend on the prices of the day-ahead market and the intraday market. It will be possible to buy and sell in both markets.

Volatility can be calculated in different ways, in this project it is calculated using standard deviation Equation 6.3. This metric measures the fluctuation of prices over a period of time, it is an indicator of the risk of a trading strategy [13].

The value at risk (VaR) is the maximum loss for a given confidence level [12]. For example, if $VaR_{99} = 100$ €, there is a 1% probability of losing more than 100 €.

The Sharpe ratio is a metric that allows to compare the investment return to the risk of the investment. The higher the Sharpe ratio the better, it means there is more return compared to the risk taken [15].

In [34], the equations used to define the Sharpe ratio are:

$$D_t \equiv R_{Ft} - R_{Bt} \quad [34] \quad (6.1)$$

$$\bar{D} \equiv \frac{1}{T} \sum_{t=1}^T D_t \quad [34] \quad (6.2)$$

$$\sigma_D = \sqrt{\frac{\sum_{t=1}^T (D_t - \bar{D})^2}{T - 1}} \quad [34] \quad (6.3)$$

$$S_h \equiv \frac{\bar{D}}{\sigma_D} \quad [34] \quad (6.4)$$

In Equation 6.1, R_{Ft} refers to the return of the fund in a period t and R_{Bt} refers to the return of the benchmark in the same period [34]. In this project R_{Bt} is set to zero.

The Sortino ratio is a modification of the Sharpe ratio. Instead of taking into account all the risk, it only takes into account the downside risk [16].

$$\text{Sortino ratio} = \frac{\bar{R} - R_f}{\sigma_d} \quad (6.5)$$

In Equation 6.5, R_f is risk-free rate and $\bar{R}$ is the mean return. σ_d is the standard deviation of $\bar{R} - R_f$ only taking into account negative deviations [16].

Maximum Drawdown is the biggest fall of the accumulated profit and it is calculated as the difference between the highest peak until the time of the Maximum Drawdown and the accumulated revenue at the drawdown [14]. The Win Rate is the ratio of profitable days or quarter-hours to the total number of trading days or quarter-hours. The Profit Factor is the ratio of profits to losses.

These metrics have allowed, not only to choose the best strategy or strategies, but also to make small modifications in order to improve each strategy when it was possible. It is important to keep in mind that the strategies that have been chosen as the best ones may vary depending on the company's objectives and the maximum risk that the company establishes.

6.1 The Naive strategy

The selected naive strategy consists of always buying in the intraday 1 market and always selling in the day-ahead market. The accumulated PnL over the test period is 15,415.38€/MWh.

Figure 6.1 shows the daily and the accumulated PnL. From October, the strategy begins to generate loses, this can be explained by the changes in the day-ahead market, where prices started to have a 15-minute frequency. Tables 6.3 and 6.4 contain the results. These results will be compared to those of the other strategies.




Figure 6.1: Evolution of the accumulated PnL and the daily PnL for the naive strategy.

Metric	Value
PnL (€/MWh)	0.87
Volatility (€/MWh)	11.05
VaR95 (€/MWh)	-15.84
VaR99 (€/MWh)	-32.27
Sharpe ratio	0.080
Sortino ratio	0.122
Profit Factor	1.295
Win Rate	0.550

Table 6.1: Metrics of the naive strategy (15 min).

Metric	Value
PnL (€/MWh)	85.17
Volatility (€/MWh)	335.31
VaR95 (€/MWh)	-424.08
VaR99 (€/MWh)	-644.35
Sharpe ratio	0.254
Sortino ratio	0.462
Max. Drawdown (€/MWh)	2367.73
Profit Factor	1.966
Win Rate	0.629

Table 6.2: Metrics of the naive strategy (24 h).

6.2 The Moving Average strategy

A moving average method consists of using the mean of the past values of the series to forecast the next one [10].

This strategy is very simple and it does not use the market price forecast that has been developed before. The average price for each quarter-hour has been calculated using the six previous months.

For each quarter-hour, the mean price over the previous six months is calculated for both markets. The mean spread is calculated as the difference between the day-ahead price and the intraday 1 price for each quarter-hour and divided by the number of days that have been used for the calculation.

For this strategy, what is important is the sign of the mean spread. To decide if in a quarter-hour energy is bought from one market or the other, only the sign of the mean spread is used. If the sign of the mean spread at a specific quarter-hour is positive, energy will be sold in the day-ahead market and bought in the intraday market. If the sign is negative, energy is bought from the day-ahead and sold in the intraday.

The quantity that is traded in this strategy will always be the same, it will be independent of the quarter-hour.




Figure 6.2: Evolution of the accumulated PnL and the daily PnL for the moving average strategy.

Figure 6.2 shows the accumulated PnL in €/MWh and the daily PnL generated by the moving average strategy. This figure shows that the accumulated PnL is positive, but from October, there are losses.

The main metrics of this strategy are in Table 6.3 with the quarter-hour metrics and Table 6.4 with the daily metrics. The total PnL is 43,204.45€/MWh. The mean PnL is positive and the win rate is greater than 0.5, meaning that there are more days and more quarter hours in which there are gains. The big difference between the daily Sharpe Ratio and the daily Sortino Ratio comes from the large positive returns that the Sortino Ratio does not take into account and the Sharpe Ratio considers, increasing the standard deviation.

Performance metric	Value
PnL (€/MWh)	2.47
Volatility (€/MWh)	11.05
VaR95	-13.28
VaR99	-25.50
Sharpe ratio	0.224
Sortino ratio	0.425
Profit factor	2.085
Win Rate	0.621

Table 6.3: Metrics of the moving average strategy (15 min).

Performance metric	Value
PnL (€/MWh)	237.39
Volatility (€/MWh)	388.24
VaR95	-341.29
VaR99	-594.13
Sharpe ratio	0.611
Sortino ratio	1.565
Max. Drawdown	9,818.42
Profit factor	4.486
Win Rate	0.693

Table 6.4: Metrics of the moving average strategy (24 h).

The moving average strategy has been slightly modified in order to try to improve its performance.

6.2.1 The Moving Average with threshold

Establishing a threshold avoids trading when the mean spread of a quarter-hour is too small. This should decrease volatility. The rest of the strategy is exactly like the moving average strategy that has been presented before. Only a new condition on the absolute value of the spread is added.

To choose the threshold, several thresholds going from 0.1€/MWh to 2€/MWh have been tested. Tables 6.5 and 6.6 contain the main metrics.

Threshold	0	0.1	0.5	1	2
PnL (€/MWh)	2.47	2.46	2.36	2.17	1.66
Volatility (€/MWh)	11.05	10.98	10.65	10.20	9.24
VaR95	-13.28	-13.03	-12.00	-11.02	-8.32
VaR99	-25.50	-25.46	-25.03	-24.39	-22.60
Sharpe ratio	0.224	0.224	0.221	0.213	0.180
Sortino ratio	0.425	0.428	0.429	0.425	0.377
Profit Factor	2.085	2.111	2.200	2.326	2.518
Win Rate	0.621	0.625	0.632	0.641	0.644
Number of trades	14,529	14,529	12,476	10,038	5,706

Table 6.5: Metrics of the moving average strategy with threshold (15 min).

Threshold	0	0.1	0.5	1	2
PnL accumulated (€/MWh)	43,204.45	43,068.03	41,195.33	37,928.85	29,014.13
PnL (€/MWh)	237.39	236.64	266.35	208.40	159.42
Volatility (€/MWh)	388.24	387.43	378.10	359.70	291.13
VaR95	-341.29	-349.26	-356.66	-349.58	-307.98
VaR99	-594.13	-594.13	-553.69	-529.10	-438.04
Sharpe ratio	0.611	0.611	0.599	0.579	0.548
Sortino ratio	1.565	1.556	1.494	1.417	1.262
Max. Drawdown	9,818.42	9,858.52	10,203.16	10,369.06	10,096.97
Profit Factor	4.486	4.472	4.257	3.994	3.672
Win Rate	0.693	0.693	0.693	0.688	0.659

Table 6.6: Metrics of the moving average strategy with threshold (24 h).

6.3 The Sortino Ratio strategy

The previous strategy used the threshold to choose whether to trade or not for a specific quarter-hour. In this strategy the Sortino ratio is used to decide whether to trade. The Sortino ratio is calculated with a window of 6 months for each quarter-hour.

The results are in Tables 6.7 and 6.8. As the threshold increases, the Sortino ratio increases, the lower the volatility, the VaR and the max.drawdown, this gives a less risky strategy.

Threshold	0	0.1	0.5	1	2
PnL (€/MWh)	2.47	2.36	1.53	0.97	0.50
Volatility (€/MWh)	11.05	10.45	7.85	6.29	4.58
VaR95	-13.28	-11.94	-5.027	0.00	0.00
VaR99	-25.50	-24.180	-17.428	-10.875	0.00
Sharpe ratio	0.224	0.226	0.195	0.156	0.109
Sortino ratio	0.425	0.439	0.454	0.437	0.394
Profit Factor	2.085	2.108	3.072	4.267	6.868
Win Rate	0.621	0.622	0.676	0.716	0.774
Number of trades	14,529	12,721	5,157	1,943	582

Table 6.7: Metrics of the Sortino strategy (15 min).

Threshold	0	0.1	0.5	1	2
PnL accumulated (€/MWh)	43,204.45	41,067.03	26,723.38	16,909.73	8,760.71
PnL (€/MWh)	237.39	226.88	146.83	92.91	48.14
Volatility (€/MWh)	388.24	379.78	230.65	128.78	71.74
VaR95	-341.29	-336.31	-254.12	-127.43	-33.75
VaR99	-594.13	-558.88	-333.71	-188.40	-113.36
Sharpe ratio	0.611	0.602	0.637	0.721	0.671
Sortino ratio	1.565	1.519	1.644	1.927	2.206
Max. Drawdown	9,818.42	10,210.71	4,271.46	2,333.85	876.47
Profit Factor	4.486	4.268	4.916	5.968	8.602
Win Rate	0.693	0.688	0.733	0.769	0.743

Table 6.8: Metrics of the Sortino strategy (24 h).

6.4 Forecast-based strategy

The LightGBM forecast has been used to determine whether electricity should be bought in the day-ahead market or in the intraday 1 market. The forecast of both the day-ahead and the intraday 1 markets are generated. Depending on the sign of the predicted spread, energy is bought or sold in each market. Figure 6.3 shows the accumulated and the daily PnL. Unlike the previous strategies, there is no decrease in performance from October onwards. Results are shown in Tables 6.9 and 6.10.




Figure 6.3: Evolution of the accumulated PnL and the daily PnL for the forecast-based strategy.

Threshold	0	1	1.5	2
PnL (€/MWh)	2.90	2.88	2.85	2.86
Volatility (€/MWh)	11.18	10.89	10.72	10.51
VaR95	-12.48	-11.63	-11.16	-10.70
VaR99	-26.16	-25.01	-24.99	-24.05
Sharpe ratio	0.260	0.265	0.266	0.272
Sortino ratio	0.500	0.522	0.529	0.555
Profit factor	2.213	2.344	2.401	2.518
Win rate	0.614	0.623	0.628	0.635
Number of trades	10,058	9,233	8,862	8,448

Table 6.9: Metrics of the forecast-based strategy (15 min).

Threshold	0	1	1.5	2
PnL accumulated (€/MWh)	31,405.55	31,192.21	30,854.17	30,941.50
PnL (€/MWh)	201.32	199.95	197.78	198.34
Volatility (€/MWh)	242.31	240.39	238.25	233.46
VaR95	-134.83	-150.51	-153.25	-134.61
VaR99	-444.81	-464.85	-469.30	-471.55
Sharpe ratio	0.831	0.832	0.830	0.850
Sortino ratio	2.424	2.342	2.277	2.345
Max. Drawdown	571.27	576.81	573.68	573.43
Profit factor	9.348	9.254	9.125	9.874
Win rate	0.840	0.840	0.833	0.853

Table 6.10: Metrics of the forecast-based strategy (24 h).

In the residuals of the LightGBM model, it was seen that there was heteroscedasticity, the error variance is not constant. This suggests that there are quarter-hours in which the model gives more accurate forecasts than others. Figures 6.4 and 6.5 show the mean PnL and the percentage of winning trades for each quarter-hour.

The quarter-hours from 47 to 61 have negative mean PnL, this suggests that it is better not to trade. Tables 6.11 and 6.12 show the results when trading does not take place in these quarter-hours.




Figure 6.4: Average PnL per quarter-hour of the forecast-based strategy.




Figure 6.5: Win rate of the forecast-based strategy.

Threshold	All quarters	Selected quarters
PnL (€/MWh)	2.90	2.90
Volatility (€/MWh)	11.18	10.84
VaR95	-12.48	-11.53
VaR99	-26.16	-25.41
Sharpe ratio	0.260	0.268
Sortino ratio	0.500	0.528
Profit factor	2.213	2.420
Win rate	0.614	0.638
Number of trades	10,058	8,671

Table 6.11: Metrics of the forecast-based strategy after eliminating quarter-hours (15 min).

Threshold	All quarters	Selected quarters
PnL accumulated (€/MWh)	31,405.55	31,391.70
PnL (€/MWh)	201.32	201.23
Volatility (€/MWh)	242.31	227.05
VaR95	-134.83	-142.35
VaR99	-444.81	-410.82
Sharpe ratio	0.831	0.886
Sortino ratio	2.424	2.578
Max. Drawdown	571.27	594.23
Profit factor	9.348	10.681
Win rate	0.840	0.872

Table 6.12: Metrics of the forecast-based strategy after eliminating quarter-hours (24 h).

The forecast-based strategy does not achieve the highest PnL, but it responds better to market changes. It has the highest Sharpe ratio, and therefore the best balance between risk and return. Eliminating the quarter hours that generated losses may result in overfitting the strategy to the specific test period. Testing on another dataset would be necessary to confirm these improvements. For this reason, the forecast-based strategy operating in all quarter-hours is selected as the preferred strategy.

Chapter 7

Conclusion

This project has studied three forecasting models for electricity market price prediction: Linear Regression, SARIMAX and LightGBM. Linear Regression did not give satisfactory results because the residuals had autocorrelation and heteroscedasticity, indicating that the model assumptions were not satisfied and resulting in statistics that were not reliable. The autocorrelation suggested that a SARIMAX model would improve the results as it captures temporal dependencies.

SARIMAX did not improve the performance. The model gave poorer results than those of the Linear Regression.

The LightGBM model had the best results compared to the other two models. Its metrics showed higher forecasting accuracy. Also, the autocorrelation of the residuals and heteroscedasticity do not invalidate the model, since LightGBM is a tree-based algorithm.

The results show that the strategies are overall satisfactory. The models do capture the general market behaviour and follow the prices when they increase. However, they do have problems capturing important spikes. The strategy benefits from the forecast as it reduces the impact of market reforms.

7.1 Future work

One of the limitations of this work is the length of the dataset. Extending the dataset would give more robust evaluations for the trading strategy and would make it possible to confirm the current results.

For future work, evaluating the strategy over a longer period is one of the most important points. Another improvement would be to use more accurate forecast

variables. Instead of taking one cloud cover location and one wind forecast location, using weather forecasts corresponding to the locations of individual generation plants would help to improve the predictions.

Another important extension would be to optimise the volume bought. In this work it was assumed to be fixed and a value was not chosen. To do so, it is necessary to study the market prices to establish an upper limit of energy that can be bought without changing the market prices.

Bibliography

- [1] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019.
- [2] Fahad Radhi Alharbi and Denes Csala. A seasonal autoregressive integrated moving average with exogenous factors (SARIMAX) forecasting model-based time series approach. *Inventions*, 7(4):94, 2022.
- [3] ENGIE. La historia del grupo engie.
- [4] ENGIE. Proyectos.
- [5] ENGIE. 2025 management report and annual consolidated financial statements, 2026.
- [6] esios REE. Mapa instalaciones producción, 2025.
- [7] Red Eléctrica España. Potencia instalada, 2026.
- [8] James D. Hamilton. *Time Series Analysis*. Princeton University Press, 1994.
- [9] Ben Harack. Electricity grid: Key terms and definitions, 2010.
- [10] Rob J. Hyndman and George Athanasopoulos. *Forecasting: Principles and Practice*. OTexts, 2 edition, 2018.
- [11] Rob J. Hyndman and George Athanasopoulos. *Forecasting: Principles and Practice*. OTexts, 3 edition, 2021.
- [12] Corporate Finance Institute. Value at risk (var), 2019.
- [13] Corporate Finance Institute. Volatility, 2019.
- [14] Corporate Finance Institute. Maximum drawdown, 2020.
- [15] Corporate Finance Institute. Sharpe ratio, 2020.

- [16] Corporate Finance Institute. Sortino ratio, 2020.
- [17] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An Introduction to Statistical Learning with applications in Python*. Springer, 2023.
- [18] Robert E. Kass and Adrian E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795, 1995.
- [19] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree, 2017.
- [20] Peter C.B.; Kwiatkowski, Denis; Phillips and Peter Schmidt. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Cowles Foundation Discussion Papers*, 1992.
- [21] Yoon Y.T. Lee GJ. Lee, J. Renewable energy sources: From non-dispatchable to dispatchable, and their application for power system carbon neutrality considering system reliability. *J. Electr. Eng. Technol.*, 19, 2023.
- [22] Celia López de la Nieta Polonio. Análisis de los precios en el mercado diario eléctrico español. Trabajo de Fin de Grado, Universidad Politécnica de Madrid, 2022.
- [23] Palanisamy Manigandam, Md. Monzurul Ali, et al. Forecasting natural gas production and consumption in united states-evidence from SARIMA and SARIMAX models. *Energies*, 14(19):6021, 2021.
- [24] Open Meteo. Historical forecast api, 2026.
- [25] Microsoft. Lightgbm documentation: Parameters, 2026.
- [26] Rizwan Mushtaq. Augmented dickey fuller test. *SSRN*, 2011.
- [27] OMIE. Curvas agregadas de oferta y demanda.
- [28] OMIE. Funcionamiento del mercado diario.
- [29] OMIE. Omie, funciones.
- [30] OMIE. Reglas de funcionamiento de los mercados diario e intradiario de electricidad, 2024.
- [31] OMIE. Detalle del funcionamiento del mercado intradiario, n.d.

- [32] Ministerio para la Transición Ecológica y el Reto Demográfico (MITECO). Estructura del sector.
- [33] Alexandre Pavlovski. Dispatchable renewables: Selecting a right pathway. *Global Journal of Researches in Engineering: General Engineering*, 24(1), 2024.
- [34] W. F Sharpe. The sharpe ratio. *The Journal of Portfolio Management*, 1994.
- [35] Nicolás Vega Muñoz. Previsión del precio spot en el mercado eléctrico español con redes neuronales. Trabajo de Fin de Grado, Universidad Politécnica de Madrid, 2024.
- [36] Jeffrey M. Wooldridge. *Introductory Econometrics: A Modern Approach*. Cengage Learning, 2015.
- [37] Energía y Sociedad. Formación de precios en el mercado mayorista diario de electricidad.
- [38] Energía y Sociedad. Mecanismos de ajuste de demanda y producción.

Chapter 8

Appendices

8.1 Data-set

8.1.1 Weather information

Weighting coefficients

Date	Andalusia	Aragon	Asturias	Cantabria	Castile-La Mancha	Castile and Leon	Catalonia	Community of Madrid
06-2024	0.221	0.088	0.001	0.001	0.212	0.080	0.041	0.016
07-2024	0.222	0.087	0.001	0.001	0.214	0.078	0.041	0.016
08-2024	0.223	0.086	0.001	0.001	0.211	0.078	0.041	0.016
09-2024	0.228	0.085	0.001	0.001	0.207	0.085	0.041	0.017
10-2024	0.242	0.083	0.001	0.001	0.203	0.084	0.040	0.017
11-2024	0.239	0.082	0.001	0.001	0.203	0.084	0.040	0.017
12-2024	0.239	0.085	0.001	0.001	0.203	0.084	0.040	0.017
01-2025	0.241	0.084	0.001	0.001	0.202	0.083	0.040	0.018
02-2025	0.241	0.084	0.001	0.001	0.204	0.082	0.040	0.018
03-2025	0.235	0.083	0.001	0.001	0.205	0.092	0.040	0.018
04-2025	0.233	0.082	0.001	0.001	0.207	0.094	0.040	0.019
05-2025	0.236	0.081	0.001	0.001	0.204	0.094	0.039	0.019
06-2025	0.243	0.079	0.001	0.001	0.201	0.098	0.039	0.019
07-2025	0.249	0.079	0.001	0.001	0.198	0.104	0.038	0.018
08-2025	0.251	0.079	0.001	0.001	0.198	0.104	0.038	0.018
09-2025	0.248	0.079	0.001	0.001	0.195	0.108	0.038	0.021
10-2025	0.250	0.077	0.001	0.001	0.195	0.111	0.037	0.021
11-2025	0.255	0.082	0.001	0.001	0.191	0.112	0.037	0.020
12-2025	0.252	0.081	0.001	0.001	0.188	0.117	0.037	0.020
01-2026	0.256	0.080	0.001	0.001	0.187	0.116	0.037	0.020
02-2026	0.253	0.079	0.001	0.001	0.187	0.118	0.041	0.020
03-2026	0.252	0.079	0.001	0.001	0.187	0.118	0.041	0.020

Table 8.1: Monthly photovoltaic weighting coefficients, part I.

Date	Valencian Community	Extremadura	Galicia	La Rioja	Murcia	Navarre	Basque Country
06-2024	0.036	0.215	0.006	0.004	0.062	0.011	0.007
07-2024	0.036	0.218	0.006	0.004	0.060	0.010	0.007
08-2024	0.036	0.219	0.006	0.004	0.060	0.010	0.007
09-2024	0.036	0.214	0.006	0.004	0.059	0.010	0.007
10-2024	0.036	0.208	0.006	0.004	0.057	0.010	0.007
11-2024	0.036	0.213	0.006	0.004	0.057	0.010	0.007
12-2024	0.037	0.210	0.006	0.004	0.056	0.010	0.007
01-2025	0.037	0.209	0.006	0.004	0.056	0.010	0.008
02-2025	0.038	0.207	0.006	0.004	0.056	0.010	0.008
03-2025	0.038	0.202	0.006	0.004	0.055	0.012	0.008
04-2025	0.039	0.199	0.006	0.004	0.055	0.012	0.008
05-2025	0.041	0.200	0.006	0.004	0.055	0.012	0.008
06-2025	0.040	0.196	0.006	0.004	0.054	0.012	0.008
07-2025	0.039	0.190	0.005	0.006	0.052	0.011	0.007
08-2025	0.039	0.189	0.005	0.005	0.052	0.011	0.007
09-2025	0.039	0.186	0.005	0.005	0.055	0.011	0.007
10-2025	0.038	0.185	0.005	0.005	0.054	0.012	0.007
11-2025	0.038	0.181	0.005	0.005	0.053	0.012	0.007
12-2025	0.038	0.178	0.005	0.005	0.056	0.012	0.007
01-2026	0.038	0.174	0.005	0.010	0.055	0.013	0.007
02-2026	0.037	0.171	0.005	0.010	0.055	0.013	0.007
03-2026	0.037	0.168	0.005	0.009	0.061	0.013	0.007

Table 8.2: Monthly photovoltaic weighting coefficients, part II.

Date	Andalusia	Aragon	Asturias	Cantabria	Castile-La Mancha	Castile and Leon	Catalonia	Community of Madrid
06-2024	0.118	0.179	0.023	0.001	0.160	0.224	0.045	0.000
07-2024	0.118	0.179	0.023	0.001	0.161	0.223	0.046	0.000
08-2024	0.118	0.179	0.023	0.001	0.161	0.224	0.046	0.000
09-2024	0.118	0.179	0.023	0.001	0.161	0.224	0.045	0.000
10-2024	0.118	0.179	0.023	0.001	0.161	0.225	0.045	0.000
11-2024	0.118	0.179	0.023	0.001	0.160	0.227	0.045	0.000
12-2024	0.116	0.181	0.022	0.001	0.158	0.231	0.045	0.000
01-2025	0.116	0.182	0.022	0.001	0.158	0.231	0.045	0.000
02-2025	0.119	0.183	0.022	0.001	0.158	0.231	0.045	0.000
03-2025	0.116	0.184	0.022	0.001	0.157	0.232	0.044	0.000
04-2025	0.116	0.186	0.022	0.001	0.156	0.233	0.044	0.000
05-2025	0.116	0.184	0.022	0.001	0.156	0.234	0.044	0.000
06-2025	0.116	0.183	0.022	0.001	0.155	0.234	0.044	0.000
07-2025	0.116	0.183	0.022	0.001	0.155	0.236	0.044	0.000
08-2025	0.116	0.181	0.022	0.001	0.153	0.238	0.044	0.000
09-2025	0.114	0.181	0.022	0.001	0.154	0.239	0.045	0.000
10-2025	0.114	0.181	0.022	0.001	0.154	0.239	0.045	0.000
11-2025	0.114	0.183	0.022	0.001	0.153	0.239	0.043	0.000
12-2025	0.114	0.186	0.022	0.001	0.153	0.238	0.044	0.000
01-2026	0.114	0.186	0.022	0.001	0.153	0.238	0.044	0.000
02-2026	0.115	0.186	0.022	0.001	0.153	0.237	0.043	0.000
03-2026	0.114	0.188	0.021	0.001	0.152	0.237	0.043	0.000

Table 8.3: Monthly wind weighting coefficients, part I.

Date	Valencian Community	Extremadura	Galicia	La Rioja	Murcia	Navarre	Basque Country
06-2024	0.040	0.003	0.128	0.015	0.009	0.051	0.005
07-2024	0.040	0.003	0.127	0.015	0.009	0.050	0.005
08-2024	0.040	0.003	0.127	0.015	0.009	0.050	0.005
09-2024	0.040	0.003	0.127	0.015	0.009	0.050	0.005
10-2024	0.040	0.003	0.127	0.015	0.009	0.050	0.005
11-2024	0.040	0.003	0.127	0.014	0.009	0.050	0.005
12-2024	0.039	0.003	0.125	0.014	0.008	0.051	0.005
01-2025	0.039	0.003	0.125	0.014	0.008	0.051	0.005
02-2025	0.039	0.003	0.125	0.014	0.008	0.051	0.005
03-2025	0.039	0.003	0.124	0.014	0.008	0.051	0.005
04-2025	0.039	0.003	0.124	0.014	0.008	0.050	0.005
05-2025	0.039	0.003	0.124	0.014	0.008	0.050	0.005
06-2025	0.039	0.003	0.126	0.014	0.008	0.050	0.005
07-2025	0.039	0.003	0.125	0.014	0.008	0.051	0.005
08-2025	0.038	0.003	0.124	0.016	0.008	0.050	0.005
09-2025	0.038	0.003	0.124	0.016	0.008	0.050	0.005
10-2025	0.038	0.003	0.124	0.016	0.008	0.050	0.005
11-2025	0.038	0.003	0.124	0.016	0.008	0.050	0.005
12-2025	0.038	0.003	0.124	0.016	0.008	0.050	0.005
01-2026	0.038	0.003	0.124	0.016	0.008	0.051	0.005
02-2026	0.038	0.003	0.123	0.016	0.008	0.051	0.005
03-2026	0.038	0.003	0.123	0.016	0.008	0.051	0.005

Table 8.4: Monthly wind weighting coefficients, part II.

Weather forecast

	Cloud cover percentage		Wind	
Autonomous Community	Latitude	Longitude	Latitude	Longitude
Andalusia	36.6875	-6.1875	36.50	-6.000
Aragon	41.5	-0.625	41.75	-0.875
Asturias	43.5	-5.67	43.3125	-6.3125
Cantabria	43.43	-3.89	43.17	-3.73
Castile-La Mancha	39.0625	-3.3125	38.875	-1.8125
Castile and León	41.3125	-4.9375	42.0625	-4.0625
Catalonia	41.5625	0.8125	41.375	0.875
Community of Madrid	40.625	-3.1875	-	-
Valencian Community	38.5625	-0.5	40.5	-0.1875
Extremadura	39.75	-6.875	40.0	-6.1875
Galicia	42.875	-7.6875	43.5625	-7.625
La Rioja	42.22	-2.09	42.2	-2.0
Murcia	37.9375	-1.375	38.5	-1.125
Navarre	42.0625	-1.625	42.65	-1.520
Basque Country	42.75	-2.89	42.92	-2.480

Table 8.5: Latitude and longitude of the cloud cover and wind data points

8.2 Electricity Price Forecasting

8.2.1 Linear Regression

The Figures in this section show the residuals of the linear regression models of the day-ahead and the intraday markets.




Figure 8.1: Autocorrelation function of the residuals of the linear regression model for the day-ahead market.




Figure 8.2: Autocorrelation function of the residuals of the linear regression model for the intraday 1 market.




Figure 8.3: Partial autocorrelation function of the residuals of the linear regression model for the day-ahead market.




Figure 8.4: Partial autocorrelation function of the residuals of the linear regression model for the intraday 1 market.




Figure 8.5: Residuals over time for the linear regression model of the day-ahead market.




Figure 8.6: Residuals over time for the linear regression model of the intraday 1 market.




Figure 8.7: Residuals versus fitted values for the linear regression model for the day-ahead market.




Figure 8.8: Residuals versus fitted values for the linear regression model for the intraday 1 market.

8.2.2 LightGBM

The Figures in this section show the residuals of the LightGBM models for the day-ahead and the intraday markets.




Figure 8.9: Autocorrelation function of the residuals of the LightGBM for the day-ahead market




Figure 8.10: Autocorrelation function of the residuals of the LightGBM for the intraday 1 market




Figure 8.11: Partial autocorrelation function of the residuals of the LightGBM model for the day-ahead market.




Figure 8.12: Partial autocorrelation function of the residuals of the LightGBM model for the intraday1 market.



Figure 8.13: Residuals versus fitted values for the LightGBM model for the day-ahead market.




Figure 8.14: Residuals versus fitted values for the LightGBM model for the intra-day 1 market.




Figure 8.15: Residuals over time for the LightGBM model of the day-ahead market.




Figure 8.16: Residuals over time for the LightGBM model of the intraday 1 market.