CHALMERS
UNIVERSITY OF TECHNOLOGY

Wind Power Intraday Trading Strategies Informed by Imbalance Price Predictions

Master's thesis in Industrial Ecology

SIMON NORIN

DEPARTMENT OF ENVIRONMENTAL AND ENERGY SCIENCES

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Master's thesis 2026

Wind Power Intraday Trading Strategies Informed by Imbalance Price Predictions

SIMON NORIN

Department of Environmental and Energy Sciences
Chalmers University of Technology
Gothenburg, Sweden 2026

Wind Power Intraday Trading Strategies Informed by Imbalance Price Predictions
SIMON NORIN

Supervisor: Ludwig Klevtun, SR Energy
Supervisor: Sonia Yeh, Department of Environmental and Energy Sciences
Examiner: Sonia Yeh, Department of Environmental and Energy Sciences

Master's Thesis 2026
Department of Environmental and Energy Sciences
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2026

Wind Power Intraday Trading Strategies Informed by Imbalance Price Predictions
SIMON NORIN
Department of Environmental and Energy Sciences
Chalmers University of Technology

Abstract

The wind power sector depends on imperfect wind forecasts when trading energy volumes ahead of delivery, leading to large differences between contracted and actual production. These differences, known as imbalances, can either be adjusted through intraday market trades before the delivery hour — once updated wind forecasts are available — or settled afterward in the balancing market. The outcome of the balance settlement is determined by the imbalance price, which became highly volatile and extreme in Sweden during 2025 following several updates to the market designs, thereby driving up imbalance costs for wind power producers. This study examines how a risk-averse intraday trading approach, here named the *Baseline* strategy, can be modified into a more risk-taking strategy that leverages predictions of the imbalance price, referred to as the *Proposed* strategy. In the *Baseline* strategy, the intraday market is solely used for lowering the imbalance volume, while the *Proposed* strategy refrains from intraday trading during instances when a constructed imbalance price prediction model, indicates a more profitable imbalance price. Both strategies were evaluated on a 2-hour and 4-hour horizon before delivery, using intraday price series available for the horizons. The findings indicates that the *Proposed* strategy can deliver a 51% cost reduction of imbalance settlement costs at the 2-hour horizon relative to a no–intraday-trading case, but only a 25% reduction at the 4-hour horizon. Over the same horizons, the *Baseline* strategy achieves cost reductions of 35% and 38%, respectively. The study concludes that the risk-taking strategy renders trading opportunities, provided that the imbalance price could be accurately predicted. At the 2-hour horizon, the imbalance price prediction is accurate enough to capture advantageous price spread opportunities, whereas for the 4-hour horizon the simpler *Baseline* strategy results in a larger reduction in imbalance costs.

Keywords: wind power, intraday strategy, imbalance cost, imbalance price, mFRR, forecasting, prediction

Acknowledgements

I would like to thank my supervisor and examiner, Sonia Yeh, for her guidance throughout this project, for answering my questions, and for providing valuable feedback. Also, big thanks to the wind power company SR Energy, where I carried out this thesis. Special thanks to Ludwig Klevtun for his support and for never growing tired of my countless questions. Finally, I want to thank all the friends I have met at Chalmers over the years for making this period such a great time.

Simon Norin, Gothenburg, June 2026

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

ANN	Artificial Neural Network
BRP	Balance Responsible Party
DA	Day-ahead
GBDT	Gradient Boosted Decision Trees
ID	Intraday
MAE	Mean Absolute Error
mFRR	Manual Frequency Restoration Reserve
MLP	Multilayer Perceptron
MWh	Megawatt hours
MW	Megawatt
RMSE	Root Mean Square Error
SVR	Support Vector Regression
SvK	Svenska Kraftnät
VWAP3H	Volume Weighted Average Price during the last three hours

Contents

List of Figures

List of Tables

1

Introduction

The main goal of the Swedish electricity market is not only to supply electrical energy to end users, but also to contribute to the regulation and stabilization of grid frequency [1]. This includes a continuous balancing of electricity generation and consumption through a set of ancillary services designed to mitigate deviations, which are managed by the transmission system operator, Svenska Kraftnät (SvK). These services operate over both short and longer time periods enabling increases or decreases in power on both the demand and supply sides of the grid when a deviation occurs. Within the ancillary services, SvK provides two categories of power reserves to address frequency deviations: containment reserves, which prevent further deviation from the nominal frequency, and restoration reserves, which act to return the frequency to its nominal value of 50 Hz [2]. The latter category includes the manual frequency restoration reserve (mFRR), a service with substantial capacity for managing such deviations, referred to as imbalances.

The mFRR market is operated based on submitted bids, offering both upward regulation (increasing generation or decreasing consumption) and downward regulation (decreasing generation or increasing consumption) [3]. In Sweden, SvK is responsible for the procurement of these bids as well as their activation when required. Since 4 March 2025, the activation of bids has been automated in the power system, replacing the previous practice of manual activation [4]. This automated regime was introduced to improve balancing accuracy and to enable the transition of the market time resolution from one hour to 15 minutes, a change that would otherwise mean that manual activation became operationally too complex.

With the higher degree of automation, certain bids are bypassed to achieve a more precise correction of the imbalance volume, which have resulted in higher clearing prices in the mFRR market. Also, a larger share of the transmission capacity between bidding zones and between countries has recently become utilized, which forces activation closer to the physical location of the imbalance [5]. This, in turn, can lead to even higher prices when low-cost regulations capacity, for example from Norway, cannot be used to manage imbalances in southern Sweden.

The payment to the parties responsible for activated regulation bids is financed by power consumers and producers that are in imbalances, i.e., have deviations between their contracted power volumes and their actual realized production or consumption. This imbalance settlement is costly for the wind power industry, where the forecasted and actual output energy volume can differ substantially depending on accuracy in weather predictions [5].

During 2025, the Swedish wind power sector had substantial costs in the imbalance settlement market, much higher than previous years [6]. A high uncertainty in power output for wind power producers is normal, due to the difficulty of accurately forecasting wind speeds. What has changed during 2025 is that the price of these imbalance volumes has increased significantly following the market regime shifts described above. These price changes is what has incurred the higher costs for the wind industry.

The elevated and highly volatile imbalance costs in 2025 have increased business risk and may consequently weaken investment incentives for the development of new wind power generation capacity [6]. In the longer perspective, this could hinder Sweden's ability to expand fossil-free electricity production at the pace required to meet future greenhouse gas emission targets. To mitigate wind power imbalance costs, this study investigates a novel intraday trading strategy aimed at lowering them further than the branch standard strategy.

This thesis proposes that instead of solely trading wind power volumes on the intraday market - a market used to adjust energy bids and reduce imbalance settlements [7] - based on the predicted imbalance volume, more cost efficient decisions may be made by predicting the imbalance price and use that information. At every intraday trading opportunity, the intraday price could be compared to the predicted imbalance price in order to sell overproduction volume to highest price (intraday price or imbalance settlement price), and buy back volume to the lowest. This in order to reduce the high cost impacts from the imbalance market.

1.1 Purpose

The primary objective of this study is to reduce imbalance costs for the wind power sector in Sweden, with a specific focus on bidding zone SE4 during the year 2025. Furthermore, the study investigates the performance of models predicting the imbalance prices and wind power imbalance volumes in 2025, and examines how these forecasts can be utilized to take intraday market positions and thereby mitigate imbalance costs. Two main bidding strategies are assessed: the *Baseline Strategy*, trading on the intraday market today with volume forecasts; and the *Proposed Strategy*, using the price spread between intraday and imbalance price to determine whether a volume should be traded in the intraday market or let the same energy volume be settled in the balance market.

The study aims to address the following research questions:

- To what extent can incorporating imbalance price forecasts into intraday bidding (*Proposed* strategy) reduce imbalance settlement costs for wind power producers in SE4, compared to the volume-only trading strategy (*Baseline* strategy)?

- How does the forecast horizon affect the performance of the *Baseline* and *Proposed* strategy?

- How have the 2025 structural changes to the Swedish mFRR market affected the forecastability of imbalance prices, and what are the consequences for the viability of price-informed intraday trading strategies?

1.2 Delimitations

The project will be limited to the Swedish bidding zone SE4 and will examine potential price and volume effects in the wind power sector during only 2025. The analysis will use data with hourly aggregated timestamps for model training, including the period following the introduction of 15-minute market resolution in the electricity system in 2025. Lastly, the extracted and analyzed wind data will be data aggregated from all the wind parks in SE4 and may cause imbalance volumes and settlements to be canceled out. Therefore the results could be marginally changed if the strategies was adopted for a specific wind park instead.

2

Background

2.1 Electricity market in Sweden

The electricity in Sweden is produced mainly from hydro, nuclear and wind power, with hydro power located in the northern regions, nuclear in the southern regions and wind power more evenly spread across the country [8]. Over the full year, Sweden is a net exporter of electric energy and have only small need for import during certain days. Like all electricity markets, the produced energy must match the consumption at every given moment to avoid frequency deviations, called balancing, and is handled by the Swedish transmission system operator Svenska Kraftnät (SvK) [9].

Except SvK, there are several other roles - controlled by or reporting to SvK - in the Swedish electricity market important for the operations [10]. Firstly, to supply the consumers with electricity there must be producers who are directly producing energy, like wind power companies. Between producers and consumers exists electricity suppliers, who sell the electricity to the consumers and can act as both middleman without own production, or be the producer themselves. Above the parties connected to supplying and consuming energy, there also exist balance responsible parties (BRPs), to where all power producers and consumers are connected. They are responsible of assuring that their portfolio of production and consumption is in balance with the market, so that someone always consumes their production or produce their consumption. Often the electricity suppliers and BRPs are parts of electricity retailers with responsibility of both supply and balance. Most of the trading connected to energy supply and balance occurring before the delivery hour (day-ahead, see 2.1.1 and intraday, see 2.1.2), is handled through the power exchanger NordPool.

Furthermore, Sweden is divided in four different bidding areas called SE1, SE2, SE3 and SE4, see Figure 2.1. In every one of the bidding areas, the price of consuming energy (named day-ahead price or SPOT price) is set independently by matching the local supply and demand curves together with the available transmission capacities from other neighboring areas [11]. Only when the transmission capacity between two connected bidding areas is filled up, the energy prices in the two areas can differ. Ever since 2011 Sweden has had these zones, which were implemented to solve the issue regarding production in north and consumption in south [12]. Before, SvK had to curtail high production in the north (like hydro power) while temporarily increasing costly energy production in the south, since one price in the whole country could not reflect the transmission limits. The balance market, see section 2.1.3, is also handled within these four areas.

Figure 2.1: The bidding zones in Sweden and neighbouring countries [13].

In the following part of the electricity market chapter, the three main parts of the Swedish electricity market is presented: the day-ahead market, which determines energy schedules and SPOT prices one day ahead of the delivery hour; the intraday market, where trades are executed to adjust the positions between the closure of the day-ahead market and the delivery hour; and the balancing market, which is used to ensure balancing of supply and demand during the delivery hour.

2.1.1 Day-ahead market

The largest marketplace for electricity trading in Sweden is the day-ahead market which is operated by NordPool. In this market, predicted electricity consumption and generation are matched one day ahead of actual delivery in order to determine the day-ahead electricity price (SPOT price) [14]. Market participants, including both producers and consumers in forms of BRPs, submit bids specifying the quantities of energy they are willing to sell or purchase for each time interval of the following day, the corresponding prices, and the price area. A closed auction is then conducted at 12:00 (noon) for the entire next day, during which all buy and sell bids, within a price area, are cleared using a marginal pricing mechanism. The resulting SPOT price is thus determined - for every 15-min interval - by the marginal price required to satisfy demand. This price corresponds to the minimum price at which all the producers are willing to supply the total demand at that same price. A lower price would lead to too low production and too much consumption, and vice versa.

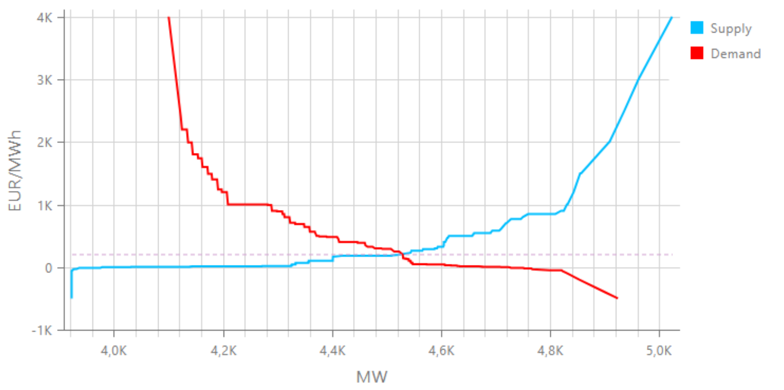

Figure 2.2: Bidding curves for supply (production) and demand (consumption) in the day-ahead market for a 15 min interval [15].

These bid curves differ for each 15-minute interval of the following day, as they depend on several factors (such as wind conditions, temperature, and time of day) and they together determine a market price. An example of these bidding curves is presented in Figure 2.2, where the intersection point represents the SPOT price for a given 15-minute interval. After the auction has been completed, the scheduled production, consumption, and transmission flows are published for each 15-minute interval in all the price areas of the upcoming day.

The primary reason for using a day-ahead electricity market is to ensure that electrical energy is generated at the lowest possible cost while keeping a transparent process [14]. This market design is particularly important for producers with time-coupling constraints, such as nuclear power plants, which often experiences long ramp-up times or high start-up costs. These factors limit their ability to adjust production within short time frames that would be required without of a day-ahead market [16]. For intermittent generation technologies, such as wind power that rely on meteorological forecasts, a longer time between bidding and actual production is instead disadvantageous. The long horizon reduces forecast accuracy and thereby increases the deviations that must be managed in the intraday or balancing markets to high costs, which is the main objective to handle in this study.

2.1.2 Intraday market

The intraday market represents the next market place for trading electricity, operating temporally closer to the physical delivery hour than the day-ahead market [7]. Its primary function is to enable the BRPs to adjust their energy positions with updated information, such as new weather forecasts or unexpected production outages. The objective with the intraday market is to increase the efficiency of the electricity production by moving electricity production to lower-cost resources and by avoiding the need for costly activations in the balancing market [17]. For wind power producers, the latter is especially important since that help reduce financial risks and costs in the balance market [7].

The intraday market has been organized on a quarter-hourly basis since 2025, accepting bids for each 15-minute interval during the current and following day. The main part of intraday trading is conducted through a continuous trading mechanism, which opens at 15:00, i.e. three hours after the closure of the day-ahead market. Market participants submit orders for both purchase and sale, and once a matching counter-order is identified, the transaction is executed immediately. This mechanism follows a pay-as-bid pricing rule, and a large number of trades are executed each minute [7]. In the methodology section, any reference to the intraday market will refer to the continuous intraday trading market.

Since 2024, the intraday market also consists of three intraday auctions (IDAs) with marginal pricing - IDA 1, IDA 2 and IDA 3 - similar to the day-ahead auction mechanism [18]. These auctions are held after the closure of the day-ahead auction, at 15:00 and 22:00 for delivery on the next coming day, and at 10:00 for same-day delivery within the time window 12:00–24:00. The introduction of these auctions is meant to complement continuous intraday trading, giving an updated and more accurate price signal to market participants [18]. The only period during which continuous intraday trading is curtailed occurs around the IDAs: trading cross price areas is suspended 20 minutes before and after each auction. The full timeline and details are illustrated in Figure 2.3. The IDAs will not be evaluated or used in this study due to the illiquidity of them.

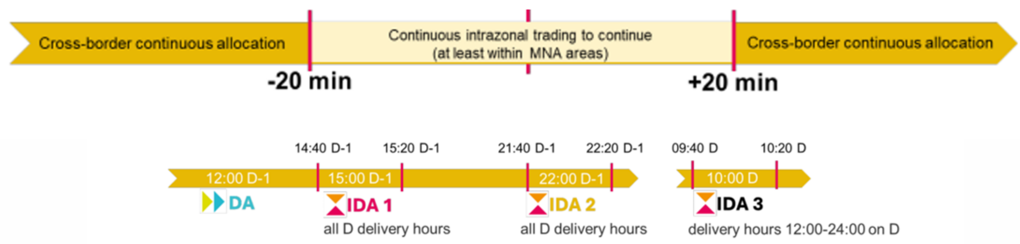

Figure 2.3: The intraday market as a timeline. Bottom line represents the three IDA auctions (IDA 1, IDA 2 and IDA 3), and the top row the continuous auction and its cross border closing during the IDAs during 40 minutes [19].

In order for BRPs to adjust their positions through buying or selling electricity, additional participants operate in the intraday market to provide the required flexibility [17]. These participants are referred to as resource owners - since they have flexible energy production and consumption - and play a crucial role in enabling BRPs to trade into balance [20]. Collectively, BRPs requiring adjustments and resource owners offering flexible capacity in the intraday market help the power system to be in production and consumption balance.

2.1.3 Balancing market

The balance market is the last one of the main market types for electricity in Sweden, and is the one that handles frequency deviations during the actual delivery hour, e.g. caused by wind power producers that are exposed to inaccurate wind speed forecasts [21]. Instead of Nord Pool that handled the two previous markets, it is SvK that is responsible of procuring bids and activating them when necessary. The bids are separated for either down or up regulation and connected to specific ancillary services operating on different timescales. Down regulation refers to lowering frequency during overproduction, through lowered production or increased consumption volumes. Contrary, up regulation refers to increasing grid frequency through increased production or lowered consumption [22].

The different ancillary services on the balancing market are offered by producers and consumers and have different requirements on endurance and activation time [21], see Table 2.1. The fastes, and often first, service being activated is the Fast Frequency Reserve (FFR) only handling up regulation, thus during a negative frequency deviation [2]. The procured bids of FFR should guarantee that the power provider can be in full operation within approximately a second after a deviation occurs (depending on the level of deviation), and be able to operate for 30 seconds.

Table 2.1: The different types of ancillary services, their activation time and endurance [2].

	Type	Activation time	Endurance
FFR	Remedial	1 second	30 seconds
FCR-D upward	Containment	10 seconds	20 min
FCR-D downward	Containment	10 seconds	20 min
FCR-N	Containment	1 min - 3 min	1 hour
aFRR	Restoration	5 min	1 hour
mFRR	Restoration	12.5 min	15 min - 30 min

To control the frequency and assure that no further deviation occurs, both for positive and negative deviations, the ancillary services Frequency Containment Reserves (FCR) exists [2]. FCR is divided into three main areas where bid can be placed: *FCR-D upward* activated during a negative deviation, when the frequency is located between 49.50 Hz and 49.9 Hz; *FCR-D down* activated during a positive deviation, when the frequency is located between 50.10 Hz and 50.50 Hz; and the FCR-N that is activated both upwards and downwards during normal (N) frequency between 49.90 Hz and 50.10 Hz.

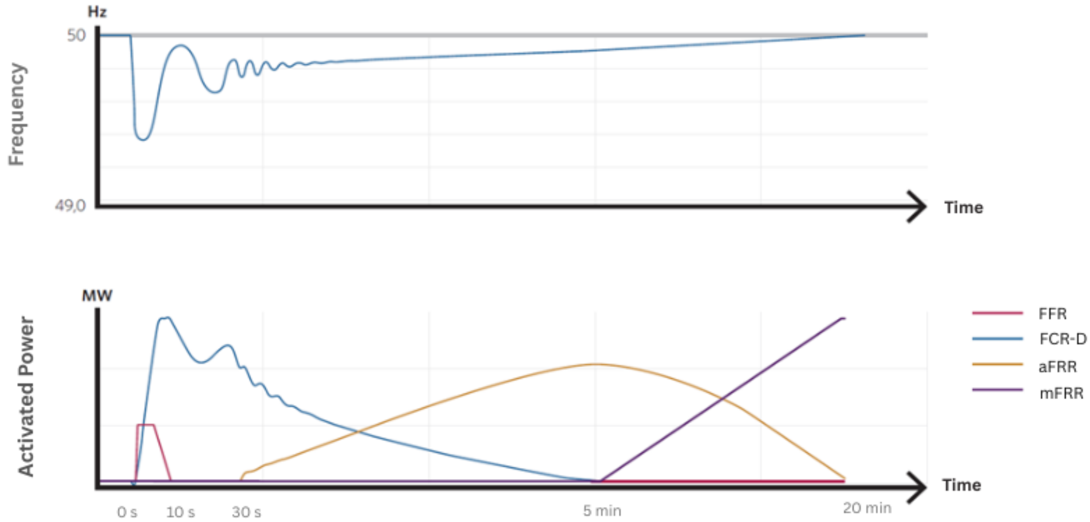

Figure 2.4: Activation of ancillary services during a larger downward deviation of the frequency. Modified from [23].

Lastly, when the other faster ancillary services have been activated to avoid frequency deviations, the Frequency Restoration Reserves (FRR) are activated to restore the frequency back to its nominal value of 50 Hz [2]. It often operate slower but with large volumes, and is divided into two parts: the Automatic Frequency Restoration Reserve (aFRR) that is activated automatically; and the Manual Frequency Restoration Reserve (mFRR) that is activated on command from SvK. An example of activation of ancillary services is shown in Figure 2.4, where different services are activated within different time intervals, and how that affects the frequency.

mFRR markets and pricing

It is the mFRR activations that sets the imbalance price. The mFRR market is divided in two parts: the Capacity Market and the Energy Activation Market (EAM) [24]. They are connected in that matter that an up or down regulation bid on the Capacity Market, requires the same bid volume or higher in the Energy Activation Market.

The Energy Activation Market is the actual market for the activations of the ancillary service mFRR, where bids are procured for both up and down regulation. Here, all the bids from the Capacity Market is transferred, but also bids from actors outside that market can contribute [26]. Only when there is a need for regulating the frequency through mFRR, known 7.5 minutes before the start of each market time quarter, the bids are activated according to the prices of the bids in Euro per MWh [27]. For up-regulation, the cheapest bids are used first, while for down-regulation, the most expensive bids are used first, see Figure 2.5. For each electricity area and time unit, there will be both a mFRR up price and a mFRR down price, set by the margin price for the activation. The dominating direction of regulation for the mFRR and the price of that direction is what determines the imbalance price in Sweden [28].

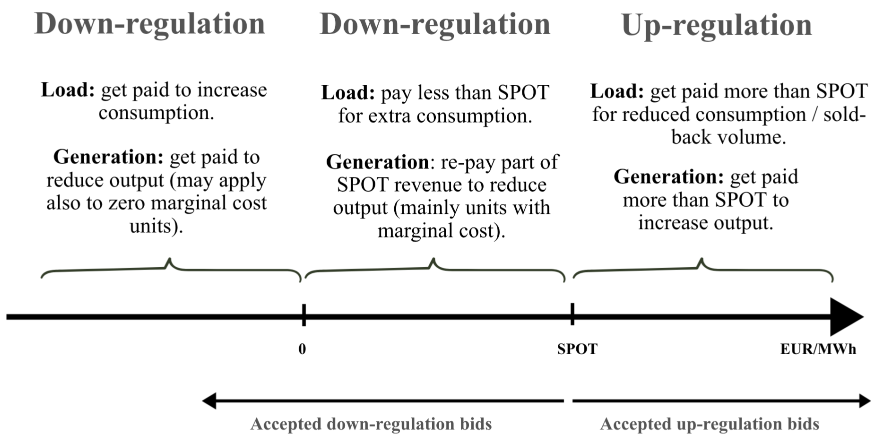

Figure 2.5: The reasoning behind the bids in the mFRR market for both producers (generation) and consumers (load), shown for up and down-regulation.

In the Swedish market, up-regulation bids are systematically higher than the day-ahead price [29]. This arises from the underlying economic incentives shown in Figure 2.5: producers require a price premium to supply additional energy above their day-ahead commitments, while consumers are only willing to reduce their consumption if they receive compensation at least equal to the cost they had in the day-ahead market.

Contrary, down-regulation bids are lower than day-ahead prices (i.e. both positive and negative prices) [29]. Logically for producers, they will only reduce their power output - and buy back volume - if the associated cost (or, in extreme cases, the payment they receive) for this is less than the revenue in the day-ahead market. Similarly, consumers will only increase their consumption if the marginal cost for that is lower than the price they would have paid for the same quantity in the day-ahead market.

Balance settlement

The above-mentioned mFRR bids and activations require active participation in the market and gives financial benefits for helping the grid frequency stabilization. However, both those parties that were activated for mFRR and those that did not participate, may have an imbalance between promised and actual production or consumption that must be financially settled [1].

The imbalance price decided by the mFRR market, is just like the day-ahead spot price: a price for energy. For example for a wind power producer, this follows the same logic: a positive imbalance price implies a profit for bringing extra energy (overproduction) to the balance market and a cost for underproduction. This while a negative imbalance price implies a cost for overproduction and a profit for underproduction [26]. This is regardless if there is up-regulation or down-regulation.

These costs, or incomes, are handled through the imbalance settlement, taking place after the delivery hour when the imbalance price and actual production or consumption is set. Firstly, the imbalance volume for every BRP is calculated with several factors which consist of: actual production and consumption; the trades on the day-ahead and intraday market; and an adjustment of possible activation of ancillary services [25]. The equation is presented below:

$$\begin{gathered}\text{Imbalance Volume [MWh]}\\ = \text{Production} - \text{Consumption} \pm \text{Trade} \pm \text{Imbalance Adjustments}\end{gathered}$$

For the imbalance volume for the wind power BRPs considered in this study, consumption is not taken into account, nor are any imbalance adjustments, since their mFRR activations are ignored.

The settlement of energy imbalance cost or payment is done through the Balance Settlement at eSett, where the imbalance price is multiplied with the imbalance differences. That results in either a negative value, implying a cost for having imbalance, or a positive value, implying a profit for having imbalance [25]:

$$\begin{gathered}\text{Balance Settlement [EUR]}\\ = \text{Imbalance Volume [MWh]} \cdot \text{Imbalance Price [EUR/MWh]}\end{gathered}$$

2.2 Regime changes 2025

During 2025, several changes to the electricity market has been implemented in order to increase efficiency, that directly and indirectly affect the mFRR market [30]. The two main changes has been the implementation of the automated mFRR EAM (Energy Activation Market) and Flowbased day-ahead calculation, both leading to changes in imbalance price causing high volatility and price spikes, see Figure 2.6.

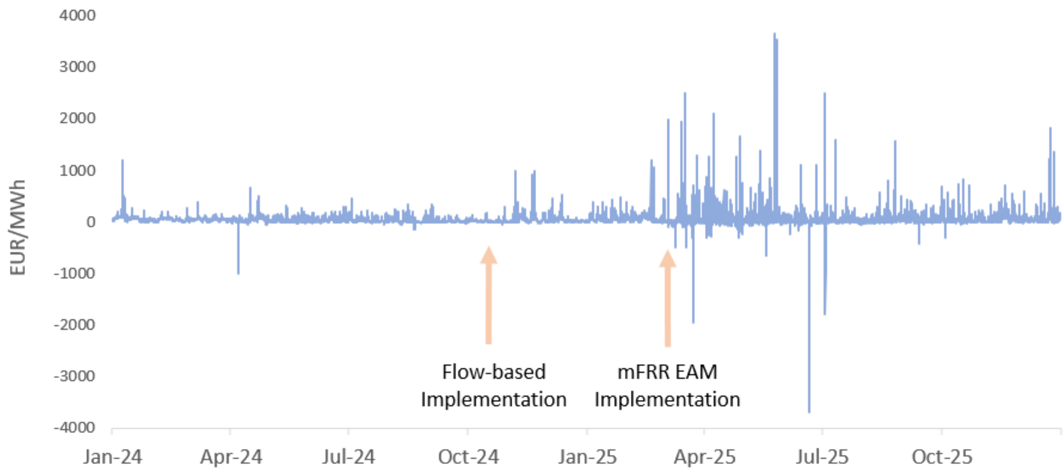

Figure 2.6: The volatile imbalance price in SE4 during 2025 compared to 2024 after the two regime changes. Data from [33].

The automated mFRR EAM market was introduced in March 2025 and meant several modifications [4]. In particular, manual activation of bids by telephone was replaced by an automatic activation, which selects bids when forecasts, close to the delivery hour, indicate a need for regulation. Bids continue to be activated according to the price order, but with a more precise matching of the required regulation volume, which may result in some bids being bypassed to achieve the exact needed volume activation [31]. This, has increased the marginal activation price and thereby led to higher imbalance prices compared with the previous regime of manual activation.

Secondly, during late 2024 the transmission capacity calculation on the day-ahead market was changed to the flow-based system, that allowed for more power capacity between bidding zones [34]. By increasing power flows between the bidding zones, the purpose was to even out spot price differences and lower the average spot price in all bidding zones. Since more power transmission capacity became available on the day-ahead market, the transmission buffer that previously was set aside for the intraday market and could help BRPs handle possible imbalances, was reduced [30]. With lower resources available for intraday market, higher share of imbalances was settled in the balance settlement, further driving imbalance prices [32].

The total imbalance *cost* for wind power, defined as the deviation between realized production and day-ahead market bids under the assumption of zero intraday trading, amounted to approximately 3,700,000 EUR in bidding zone SE4 in 2025 (at hourly temporal resolution). This is illustrated in Figure 2.7, which presents the time series of accumulated imbalance income. For comparison, the corresponding wind power imbalance *income* in SE4 during 2024 was approximately 750,000 EUR, illustrating how the wind industry is affected by the regime changes. The 3,700,000 EUR will later be used as a comparison value of how well the intraday strategies are performing.

Figure 2.7: Cumulative imbalance income for wind power producers in SE4 during 2024 and 2025. Data from [61].

3
Theory

3.1 Predictive models

In this section, the predictive models used in the study will be presented and described: Gradient Boosted Desicion Trees (GBDT), Support Vector Regression (SVR) and Multilayer Perceptron (MLP) as well as an Averaged model combining them. Above these four predictions, a fifth Naïve prediction model is applied in the method section, where the forecast value is simply set as the latest known value.

3.1.1 Gradient boosted decision trees (GBDT)

Decision trees are a powerful method for answering binary questions and number prediction in order to iteratively exclude outcomes and identify the most probable one [35]. Given an instance to be classified with a set of input variables, the model applies a sequence of binary splits and by walking the decision path in the tree, it produces an output that is expected to approximate the correct value. For large datasets and complex prediction tasks, aggregating multiple decision trees to solve the same problem typically yields higher predictive accuracy. In Gradient Boosted Decision Trees (GBDT), each successive tree is trained by the loss function with respect to the predictions of the previous trees. An example tree is visible in Figure 3.1.

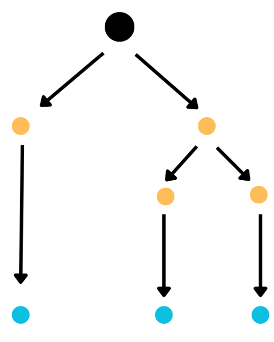

Figure 3.1: A simple decision tree with distinct outputs (blue) determined by the responses associated with the input variables (yellow).

GBDT models are particularly good at dealing with missing data and outliers, which makes them well suited for applications such as electricity market forecasting [37]. Features that appear near the root of the tree hold a greater influence on the final prediction, so the model indirectly performs a form of feature selection. However, GBDT models are also in risk of overfitting if the trees are too deep. The ensemble may then fit the training data too closely and becoming sensitive to noise in the test data.

For each time interval i, an input vector $\mathbf{x}_i \in \mathbb{R}^n$ is used to predict the target value y_i. In gradient boosting, the prediction is produced by an additive model $F(\mathbf{x})$, giving the sum of many small decision trees [36]. To train the model, a loss function L is defined to measure the difference between the observed value y_i and the model prediction $F(\mathbf{x}_i)$. In this work, the squared error loss is used:

$$L(y_i, F(\mathbf{x}_i)) = \frac{1}{2}\left(y_i - F(\mathbf{x}_i)\right)^2 .$$

Gradient boosting builds the model iteratively. At iteration m, the current model is named F_{m-1}. The next tree is fitted to reduce the loss by following the negative gradient of the loss with respect to the model output [38]:

$$r_{im} = -\left.\frac{\partial L(y_i, F(\mathbf{x}_i))}{\partial F(\mathbf{x}_i)}\right|_{F=F_{m-1}} .$$

For the squared error loss, this simplifies to:

$$r_{im} = y_i - F_{m-1}(\mathbf{x}_i).$$

The next regression tree f_m is then trained to approximate these residuals by least squares:

$$f_m = \arg\min_{f \in F} \sum_{i=1}^{N} \left(r_{im} - f(\mathbf{x}_i)\right)^2 .$$

Finally, the large model is updated by adding the newly trained tree scaled by a learning rate $\nu \in (0, 1]$, which controls the contribution of each tree to the final model:

$$F_m(\mathbf{x}) = F_{m-1}(\mathbf{x}) + \nu f_m(\mathbf{x}).$$

After m iterations, the final boosted model is given by $F_m(\mathbf{x})$.

3.1.2 Support vector regression (SVR)

Support Vector Regression (SVR) is a well used method for time series prediction, particularly for data characterized by high volatility and non-linearities [43]. SVR is a regression-based predictive model, but it uses a different strategy than regular regression models [39]. Rather than only minimizing the average prediction error between forecasted and observed values and using this to adjust the model, SVR introduces an ϵ-insensitive loss function. A tolerance tube of width ϵ is defined around the current prediction, within which deviations are not penalized. Only data points lying outside this ϵ-insensitive region are associated with a loss and are represented by support vectors. These support vectors are then used to optimize the regression function, as illustrated in Figure 3.2. The objective is to determine a function $f(x)$ that approximates the training data such that all observations lie within an error margin ϵ [40].

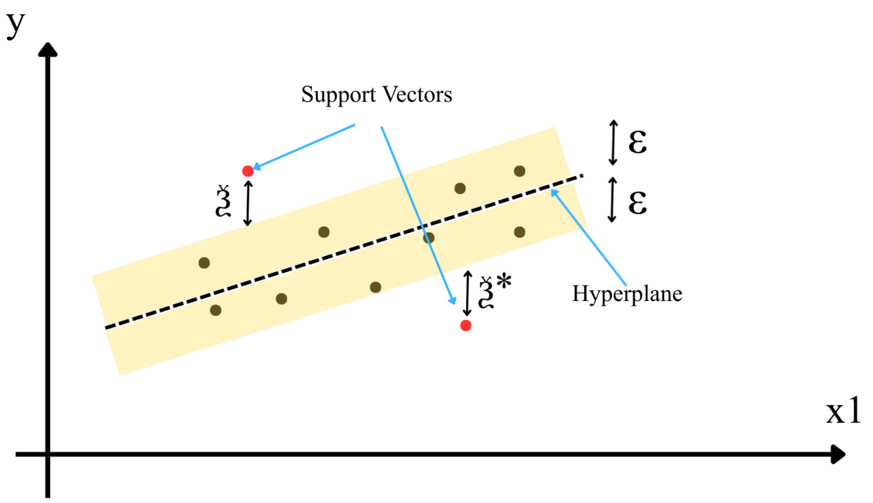

Figure 3.2: A one dimensional Support Vector Regression, with tolerance band ϵ and the data points ξ and ξ^* building the support vectors.

In its linear form, the predictor is:

$$f(\mathbf{x}) = \mathbf{w}^\top \mathbf{x} + b,$$

The goal is to both minimize the support vectors and flattening out the hyperplane [40]. Therefore the SVR tries to solve the following:

$$\min_{\mathbf{w},b,\boldsymbol{\xi},\boldsymbol{\xi}^*} \quad \frac{1}{2}\|\mathbf{w}\|^2 + C\sum_{i=1}^{N}(\xi_i + \xi_i^*),$$

Here, $C > 0$ controls the trade-off between model flatness and the magnitude of prediction errors outside the tube. To model nonlinear relationships, with several input variables, SVR can be extended using a kernel. The prediction function can be written as

$$f(\mathbf{x}) = \sum_{i=1}^{N}(\alpha_i - \alpha_i^*)\, K(\mathbf{x}_i, \mathbf{x}) + b,$$

where α_i, α_i^* are Lagrange multipliers.

In this work, the Radial Basis Function (RBF) kernel is used:

$$K(\mathbf{x}, \mathbf{z}) = \exp\left(-\gamma\|\mathbf{x} - \mathbf{z}\|^2\right),$$

where $\gamma > 0$ controls the kernel width [41]. Larger γ makes the model more local, while smaller γ gives a more global fit. Because the RBF kernel depends on distances, the input data should be scaled to comparable magnitudes [42]. Therefore, a standardization step is used before fitting the SVR.

3.1.3 Multilayer perceptron (MLP)

Another type of prediction model employed in this study is the Multilayer Perceptron (MLP), which belongs to the class of artificial neural networks (ANNs) [44]. An MLP consists of a set of interconnected nodes (neurons), organized in layers. The input variables are combined with associated weights as they propagate through one or several hidden layers, producing the model output, see Figure 3.3. By comparing the predicted output with the observed target value, an error term is obtained. This error is then implemented into the training procedure to adjust the model parameters.

MLPs are well suited for forecasting time series with nonlinear dynamics and can learn complex patterns from the data during training, without the need to specify them on beforehand [45]. For instance, if the difference between two input variables is more informative than the individual variable values themselves, an MLP can learn this interaction more effectively than many other statistical models.

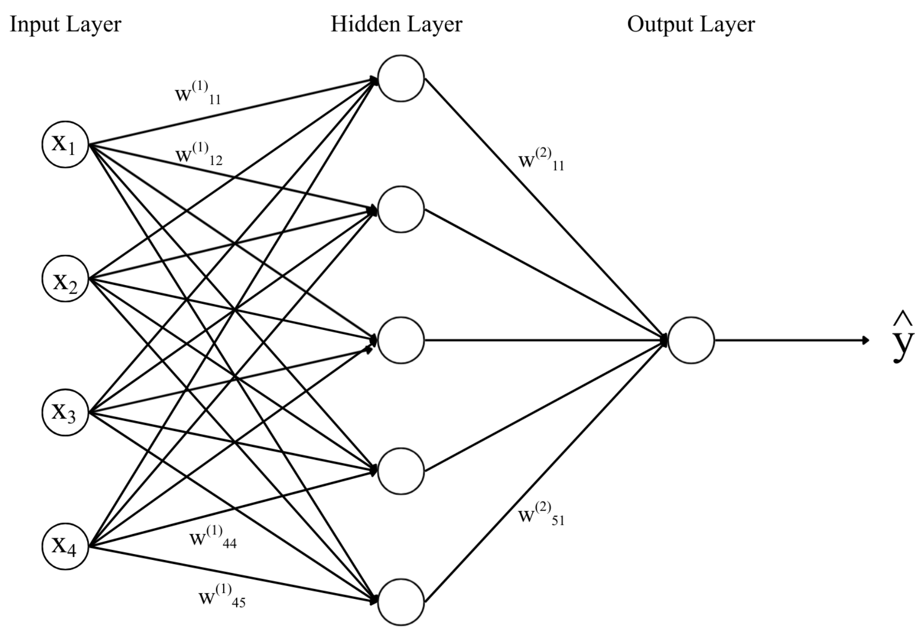

Figure 3.3: Schematic illustration of an MLP with four input variables (x), one hidden layer, and one output variable (y).

The input layer of the MLP is constructed with as many nodes as there are input variables, while the hidden layers are specified separately. A common design choice is to use one or several hidden layers with a decreasing number of neurons towards the output layer [44]. In each neuron, the outputs from the previous layer are multiplied by the weights (indicated by the arrows in Figure 3.3), added to a bias term, and then transformed by a nonlinear activation function. For a network with a single hidden layer, as depicted in Figure 3.3, this can be expressed as

$$\mathbf{h}^{(1)} = \phi\Big(\mathbf{W}^{(1)}\mathbf{x} + \mathbf{b}^{(1)}\Big),$$
$$\hat{y} = \mathbf{W}^{(2)}\mathbf{h}^{(1)} + b^{(2)},$$

where $\mathbf{x}$ is the input vector, $\mathbf{h}^{(1)}$ the hidden-layer activations, $\mathbf{W}^{(1)}$ and $\mathbf{W}^{(2)}$ the weight matrices, and $\mathbf{b}^{(1)}$ and $b^{(2)}$ the bias terms. The initial weights and biases are first set randomly and are then updated based on the error value between the model output and the true target. For the MLP considered here, the loss function is given by the squared error

$$L = \frac{1}{2}(\hat{y} - y)^2,$$

where y denotes the true target and $\hat{y}$ the model prediction. Using the backpropagation algorithm, the gradient of the loss is moving from the output layer backward through the network, and the weights and biases are updated accordingly. For the output layer, this can be written as

$$\mathbf{W}^{(2)} \leftarrow \mathbf{W}^{(2)} - \eta \frac{\partial L}{\partial \mathbf{W}^{(2)}}, \qquad b^{(2)} \leftarrow b^{(2)} - \eta \frac{\partial L}{\partial b^{(2)}},$$

where η is the learning rate. The same procedure is then applied to all hidden layers in the MLP until all weights and biases have been updated, after which the prediction model is evaluated again on the training data [44]. This training process is repeated for a specified number of epochs or until the model converges.

3.1.4 Averaged model

Individual models may have predictive inaccuracies at specific predictions, which can degrade their overall performance. Therefore, combining multiple models has become a standard approach to avoid the influence of individual model errors [46]. Previous studies on forecasting have demonstrated that forecast combinations frequently outperform the individual models, including those based on more complex methods such as machine learning [46].

A variety of approaches exist for combining predictive models. These approaches differ primarily in how they handle weighting. It is possible to assign higher weights to better-performing models, which can improve the average predictive performance. Nonetheless, several studies indicates that assigning equal weights to all models and using the arithmetic mean is a simple and effective strategy, often yielding performance comparable to more complex weighting combinations [46]. The simple average forecast is computed as

$$pred_{average,i} = \frac{pred_{1,i} + pred_{2,i} + \cdots + pred_{n-1,i} + pred_{n,i}}{n}$$

where n denotes the number of models and i indexes the prediction time point. In this study, the above mentioned more simple approach will be used for the Averaged model predictions.

3.2 Model evaluation

To assess the predictive performance of the models, two primary error metrics are employed in this study: the root mean square error (RMSE) [47] and the mean absolute error (MAE) [48]. For both, lower values indicate smaller prediction errors and therefore better model performance.

The RMSE is defined as

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

where y_i denotes the observed (true) value and $\hat{y}_i$ the corresponding model prediction at time index i. RMSE computes the square root of the mean of the squared prediction errors, thereby assigning a higher penalty to larger deviations. Consequently, RMSE is informative in applications where large errors are proportionally more harmful compared to lower errors [47]. However, this property also makes RMSE highly sensitive to outliers, where one specific value can influence the overall model performance heavily.

The MAE is defined as

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|,$$

and measures the mean of the absolute prediction errors [48]. In contrast to RMSE, MAE penalizes all deviations linearly, providing the typical predictive error. For both RMSE and MAE, the resulting values are expressed in the same units as the target variable, which give a direct and understandable model performance assessment [47].

3.3 State of the art research

This chapter presents other research within the imbalance price forecasting and different bidding strategies using that price forecast.

3.3.1 Imbalance price forecasting

Several previous studies have addressed the forecasting of imbalance electricity prices over various time horizons, employing a range of models and methodological approaches. A common practice is to apply some form of artificial neural network (ANN) model and compare its performance against one or more alternative prediction models [49, 50, 51, 52]. In all of the analyzed studies on imbalance price forecasting, the neural networks is the most accurate one compared to other statistical time-series models, see Table 3.1.

Table 3.1: Summary of the imbalance price forecasting research.

Study	Country	Year	Predictive Models	Best Model
[49]	Netherlands	2020	SVR & FFDNN	FFDNN
[50]	Netherlands	2019	ARMAX & ANN	ANN
[51]	Sweden (SE2, SE3)	2024	Logistic Reg. & ANN	ANN
[52]	Great Britain	2016	FCNN & RNN et. al	FCNN

The primary objective in many of these studies is to identify the most accurate prediction model, either by modeling up- and down-regulation prices separately or by representing the imbalance price as a single time series. Across all reviewed studies, it is consistently shown that the high volatility of imbalance prices limits forecasting accuracy. In cases where a naïve baseline model has been included for comparison, the relative performance gains of more advanced models are often small or, in some instances, negligible [50] [51]. This volatile and weakly predictable price behavior is observed across different geographical regions and different years within Europe.

One study compared a Feed-Forward Deep Neural Network (FFDNN), with a statistical approach based on Support Vector Machines (SVM) [49], and found that the ANN method achieved a slightly stronger predictive performance. The analysis was carried out using a Dutch imbalance price series for the year 2020. It formulated the task as a four-class problem with two classes for up-regulation and two for down-regulation: up greater than 55 EUR/MWh, up less than or equal to 55 EUR/MWh, down less than 0 EUR/MWh, and down equal to 0 EUR/MWh. In the model testing phase, the reported prediction accuracy for the FFDNN was 26–29%, whereas the SVM achieved 24–28 %, relative to an implicit baseline of 25% corresponding to an equal distribution. The study further notes that major global events influencing supply and demand (such as the Covid-19 pandemic) can affect model performance, and was shown by the lower predictive accuracy on the 2020 test data compared to the earlier test period.

Another study on imbalance price forecasting, conducted for the Great Britain market, evaluates models proposed in earlier research by comparing nonlinear statistical methods with several ANN-based approaches [52] and the results indicate only modest improvements in forecasting. The Fully Connected Neural Network (FCNN) achieved an RMSE of 14.5, while the previously proposed Regime-Switching model and Bayesian Dynamic Linear Model obtained RMSE with slightly worse performance values of 14.7 and 15.9, respectively. However, the ANN models clearly outperformed the statistical benchmarks in generating predictive density functions for the imbalance price, thereby improving the identification of price spikes. This advantage is explained by the capability of ANN models to more effectively use information in longer historical time series during training. It allows them to incorporate rare events and therefore increase the forecasting under such extreme events.

Furthermore, forecasting models typically use multiple input variables to predict the imbalance price, including SPOT prices, transmission capacities, generation, load, and meteorological data [51]. Several studies discuss the temporal shifting of input variables to prevent information leakage in both training and test datasets, i.e., to ensure that the model does not access information that would not have been available at the time of prediction [52]. This procedure commonly concerns variables that are published with a delay, such as actual load, generation, transmission flows, and weather observations. In contrast, predictive variables determined in advance, such as day-ahead prices and scheduled production, are considered safe to use without introducing leakage.

The literature also differs in the number of explanatory variables: some studies use a reduced set of inputs [50, 49], whereas others consider a larger set of predictors [51]. One study [49] reports that increasing the number of input variables can improve the accuracy of price forecasts. Another instead evaluates candidate input variables based on their mutual information with the target [51], i.e., to which extent each input variable gives additional information. These findings indicate that input variables differ in their predictive relevance and that predictive performance gains from adding variables depend on their information content.

3.3.2 Intraday bidding strategies

This section is divided into two parts: first, the rationale underlying different intraday trading strategies and the key aspects to be considered; and second, the outcomes of these strategies as reported in previous studies.

Rationality and aspects of intraday strategies

Intraday trading offers high potential for risk mitigation by enabling the reduction of imbalance volumes at a known price. In the absence of intraday trading, large imbalance volumes must be settled at an uncertain and volatile imbalance price [53]. However, for a wind power producer to engage in intraday trading informed by imbalance price prediction, several factors that influence the economic value of the intraday market must be taken into account. The most critical of these are the risks associated with incorrect forecasts of the spread between imbalance and intraday prices, as well as uncertainty in the prediction of imbalance volumes. In addition, there are regulatory requirements as a Balance Responsible Party (BRP) and potential illiquidity in the intraday market which may hinder the execution of intraday trades [54] [55].

When making use of the intraday market, the primary objective is that the avoided imbalance settlement cost exceeds the costs from the intraday trading [53]. Since both the imbalance volume of wind power and the imbalance price are unknown, intraday decision-making becomes a trade-off between the known costs and incomes of trading intraday, and the avoided incomes from the balance market. Two polar cases of intraday trading strategies - one risk-averse and one risk-taking - are illustrated in Figure 3.2.

These two strategies are discussed in the literature [56, 58] and represent two fundamentally different perspectives on the use of the intraday market. In the *risk-averse strategy*, wind power producers sell their forecast imbalance volume (both overproduction and underproduction) independently of the predicted imbalance price. This strategy can serve as a benchmark scenario against which alternative approaches can be evaluated [58]. The underlying rationale is to reduce imbalance volumes by intraday trading and thereby decrease exposure to imbalance cost risk, at the expense of missing situations where the imbalance market could yield lower costs or higher revenues. Only in two cases, the intraday trading is intentionally resulting in a net profit (green color in the left Table 3.2), compared to refraining from trading and leaving the volume be settled to the imbalance price. In the other two cases, the intraday trading is resulting in a net economic loss (red color in the left Table 3.2), compared to the balance settlement option.

Table 3.2: Two types of strategies on the intraday market for wind power producers. Green color represents an intraday trade of total economic loss and green color represents an intraday trade of total economic gain, compared to no intraday trading and leave volume to be settled to the imbalance price.

	Under production	**Over production**
ID price > Imb. price	Buy ID	Sell ID
ID price < Imb. price	Buy ID	Sell ID

Risk-averse strategy
(BRP Compliant)

	Under production	**Over production**
ID price > Imb. price	Sell ID	Sell ID
ID price < Imb. price	Buy ID	Buy ID

Risk-taking strategy
(BRP Non-Compliant)

In contrast, under the *risk-taking strategy*, intraday trading decisions are not driven by the forecasted imbalance volume. Instead, positions in the intraday market are taken by exploiting the price spread between the intraday market price and the forecasted imbalance price [56, 58]. This creates an arbitrage opportunity on the price differential: when the intraday price exceeds the predicted imbalance price, a certain volume is sold on the intraday market and repurchased on the balancing market at a lower price. When the predicted imbalance price instead exceeds the intraday price, volume is purchased intraday and sold at a higher price on the balancing market. As a consequence, a wind power producer may accept larger imbalance volumes as a result of intraday trading and is thereby exposed to a high risk that the realized imbalance price differs from the forecast. In this strategy, all the four trading combinations are intentionally resulting in a net profit compared to no trade (green colors in the right Table 3.2).

In addition to the economic incentives to participate in intraday trading to reduce high imbalance costs, BRPs (for example wind power producers) are subject to regulatory requirements to support overall system balance. According to the framework established by SvK, it is prohibited to maintain large imbalance volumes [54]. By implementing the risk-averse strategy, the wind power producer not only manages its own imbalance costs but also contributes to system balance and remains compliant with its BRP obligations. In contrast, the risk-taking strategy conflicts with the agreement signed by Swedish BRPs, as efforts to influence system imbalance volumes outside of designated ancillary service markets are prohibited [59]. In studies that consider the latter strategy — which is more used in theoretical analyses than in practical operations, although exceptions exist [56] — explicit trading limits are often introduced. That since, unlike in the risk-averse strategy where only the predicted imbalance volume is traded, there is no upper limit on the on possible traded volumes.

Lastly, the intraday market in Sweden is characterized by substantially lower traded volumes compared to the day-ahead market and by a smaller number of active participants [53]. Traded volumes are frequently so limited that the market is considered illiquid, meaning that desired volumes cannot always be bought or sold without large price movements [55]. Recent trends indicate increasing liquidity over the past few years, and regulatory measures such as the introduction of intraday auctions aim to increase the availability of buy and sell opportunities, thereby complementing continuous trading. Nevertheless, the illiquidity in the intraday market continues to hinder cost-efficient imbalance adjustment and must therefore be accounted for both in practical system operation and in analytical studies investigating bidding strategies.

Studies on intraday strategies

Several studies show that the intraday market both can render incomes when traded correctly and help reduce imbalance costs with an accurate imbalance price forecast. This with a risk-taking strategy based on the above-mentioned BRP non-compliant strategy [56, 57, 58], see Table 3.3.

A 2012 study of a wind power park in the Dutch electricity market shows the added value of more intelligently using the intraday and balancing markets by analyzing a range of bidding strategies [58]. The investigated strategies are grouped into two main categories: a risk-averse strategy in which available capacity is sold in the day-ahead market and then adjusted in the intraday market based on updated forecasts; and a risk-taking strategy in which an optimization model, informed by forecasted imbalance prices, selects between trading in the intraday or balancing market according to profitability. The results indicate that the optimization approach using imbalance price forecasts can increase profits (including loss reductions) by up to 17% relative to the business-as-usual strategy. In the model, they have also incorporated the fact that not all bids may be accepted in the intraday market, due to

Table 3.3: Summary of the intraday strategy research.

Study	Country	Year	Intraday Strategy	Results
[56]	Germany	2022	Risk-taking strategy	42 000 EUR per MW traded
[57]	Great Britain	2022	Risk-taking strategy (only sell part)	Prediction closer to delivery increases profit
[58]	Netherlands	2012	Risk-taking strategy	17% cost reduction

illiquidity in the market.

A 2022 study of the German electricity market analyzes a model in which the imbalance price is forecast on a day-ahead horizon for both upward and downward regulation and combined via a weighted sum of these two components [56]. The model then evaluates the profitability of purchasing 1 MW in the intraday market and realizing additional revenue in the imbalance settlement, as well as the opposite of selling 1 MW intraday. Over a two-year period (2017–2019), the authors investigate two different approaches to predicting the imbalance price: a logistic regression model and a model based on lagged system imbalance, demonstrating that the logistic regression model outperforms the lag-based approach. Under the assumption of trading only 1 MW at each intraday timestamp, the logistic regression strategy yields a total profit of 42,000 EUR, 47,000 EUR more than the simple approach. The study further notes that the underlying behavior on which the method is based—namely taking an intentional imbalance position —is not permitted in Germany. Nonetheless, some market participants are mentioned to engage in such practices.

In addition, the study conducted for the Great Britain electricity market [57] examines how imbalance price forecasts can be utilized either to trade energy at a fixed price intraday or to allow deviations to be settled through the imbalance mechanism. A key distinction from the previously mentioned studies, is that the latter only considers the sale of surplus energy (specifically, 1 MWh per half-hour interval) and focuses only on how improvements in imbalance price forecasting affect changes in profit, rather than on the absolute profit level. The comparison of different imbalance forecasting approaches demonstrates that shortening the forecasting horizon from day-ahead to intraday increases profits by approximately 2–5%.

Finally, the impact of price variation resulting from intraday market bidding and from adjustments to imbalance volumes in the balancing market is mentioned and, to some extent, analyzed in the literature [56, 60]. The intraday and balancing prices used in these studies are determined by prior market conditions and consequently, any systematic change in bidding strategies would alter traded volumes and thus influence the resulting price levels. For instance, a study of the Belgian power market for the period 2021–2025 [60] quantifies the relationship between imbalance volume and imbalance prices, estimating the marginal price effect per additional MW of imbalance. The results indicate that the up-regulation price increases by 0.41 EUR for each additional MW of system imbalance. This relationship is therefore incorporated into the bidding model for simulating a more realistic market behavior. In the German market study from 2022 [56], this potential effect is also recognized. However, it is argued that small transaction sizes on the order of 1 MW have a negligible impact on price formation and can thus be ignored.

4 Methods

This chapter provides the description and justification of the methods used in this study. The structure of the chapter is outlined in Figure 4.1. It begins with the procedures for data collection and preprocessing to mitigate the risk of information leakage. Subsequently, two predictive models were developed for the year 2025: one model for forecasting imbalance prices and another for predicting wind power energy imbalances. The model outputs were then applied to two intraday bidding strategies. Finally, the performance of the models was assessed in several aspects, among them model aging and sensitivity to changed model parameters.

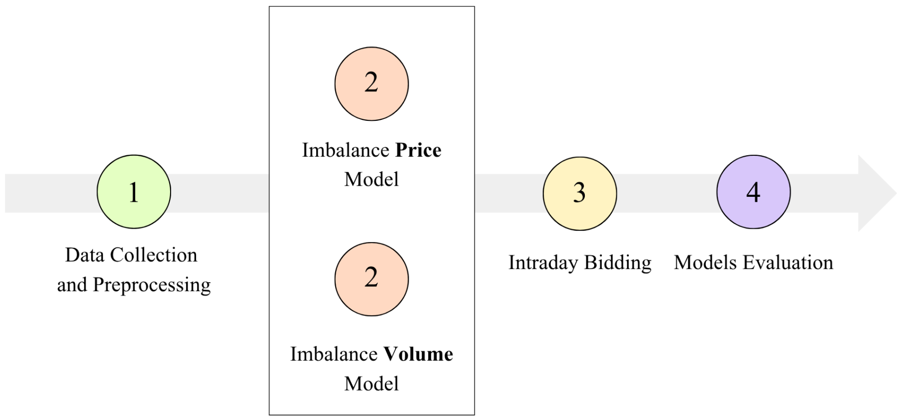

Figure 4.1: Schematic overview of the methodological framework employed in the study.

Data collection and preliminary data processing were performed in *Microsoft Excel*, whereas forecast modelling and intraday bidding were conducted in *Python 3.11*. The analysis relied on the *Python* libraries *NumPy*, *Pandas*, *os*, and *scikit-learn*, from which the classes and functions *StandardScaler*, *HistGradientBoostingRegressor*, *SVR*, *MLPRegressor*, *mean_squared_error*, and *mean_absolute_error* were used.

4.1 Data collection and preprocessing

The input data for the modeling were obtained primarily from two sources: the ENTSO-E Transparency Platform and the Nord Pool data platform [61][33]. From ENTSO-E, power system data were extracted, including total electricity generation and load in the SE4 bidding zone, as well as transmission imports from and exports to all neighboring bidding areas.

Table 4.1: Overview of categories, data names, units, and time shifts used in the dataset, where n denotes the forecast horizon hours.

Category	Data Name	Unit	Time Shift (hours)
Day-Ahead Price	SE4 SPOT Price	EUR/MWh	0
	DK2 SPOT Price	EUR/MWh	0
	GER SPOT Price	EUR/MWh	0
	LT SPOT Price	EUR/MWh	0
	PL SPOT Price	EUR/MWh	0
	SE3 SPOT Price	EUR/MWh	0
Day-Ahead Volume	SE4 DA Load	MWh	0
	SE4 DA Generation	MWh	0
	SE4 DA Import	MWh	0
	SE4 DA Export	MWh	0
	SE4 DA Wind Generation	MWh	0
Actual Volume	SE4 Actual Load	MWh	1+n
	SE4 Actual Generation	MWh	1+n
	SE4 Actual Import	MWh	1+n
	SE4 Actual Export	MWh	1+n
	SE4 Actual Wind Generation	MWh	1+n
Target Variables	Imbalance Price (SE4)	EUR/MWh	-
	Wind Actual - Wind DA (SE4)	MW	-
Intraday Bidding Variables	Intraday Price Closing (SE4)	EUR/MWh	-
	Intraday Price VWAP3H (SE4)	EUR/MWh	-

These datasets consisted of both realized production values and day-ahead forecasted values. From Nord Pool, market-related data were collected, including day-ahead SPOT prices in SE4 and the neighboring bidding areas: DK2 (Denmark), GER (Germany), LT (Lithuania), PL (Poland) and SE3 (Sweden). Lastly, intraday prices (both closing prices and the volume-weighted average price over the last three hours, VWAP3H), as well as imbalance prices in SE4 were collected.

The data were downloaded in Coordinated Universal Time (UTC) with an hourly temporal resolution for the full period from 2021-01-01 to 2025-12-31 for all timestamps, thereby yielding time series aggregated at the hourly level also after the introduction of the 15-minute market resolution. The data period was decided based on the requirement to have full time series for all the input variables, while the resolution of one hour was set to increase model performance by training and testing on data with similar resolution, even if that meant detailed quarterly information was lost.

The resolution setting was specified directly in the data extraction interfaces of both ENTSO-E and Nord Pool, and therefore needed no manual adjustment. Table 4.1 summarizes the collected time series, grouped into the following categories: day-ahead prices, day-ahead volumes, realized (actual) volumes, target variables (one for the imbalance-price model and one for the volume-imbalance model), and intraday prices employed in the bidding strategy presented in Section 4.3. For each variable, the table also reports the temporal shift applied to mitigate target leakage.

4.2 Forecasting

This section describes the forecasting methodology. That includes two main parts: how leakage was handled; and how the models were selected and their motivation.

4.2.1 Leakage handling

The risk of data leakage arised because the analysis was based exclusively on historical data and did not incorporate real-time testing. Data leakage was therefore addressed in two principal forms: train–test contamination, i.e., training the model on test data; and target leakage, i.e., training on variables that would not be available at the actual time of prediction [69]. Both types of leakage can induce overfitting, leading to inflated performance during backtesting while the model actually has lower predictive accuracy under real-world operating conditions.

To mitigate target leakage, the dataset was preprocessed by applying temporal shifts to selected variables so that their timestamps reflected realistic information availability. The raw input for the forecasting models consisted of hourly observations in which all variables were initially aligned on the same timestamp. Directly using all variables from a given timestamp to predict the target at that same timestamp would introduce target leakage, as it would assume access to information that was not available at prediction time. To address this, specific time series (those subject to slower publication or data upload) were shifted forward in time. This procedure was consistent with the methodological choices reported in previous studies on imbalance price and volume forecasting, such as [50] and [51].

In particular, the realized (actual) energy variables obtained from the ENTSO-E platform [61] (SE4 load, generation, import, and export) are published with an approximate one-hour delay (e.g., at the end of hour 12:00, only data up to hour 11:00 are available). Accordingly, the realized energy values were shifted one hour forward in time so that, at each prediction time, only those values that would have been available in practice were presented to the model. The actual energy series were further shifted by the specified forecast horizon (in hours). In contrast, day-ahead energy forecasts and SPOT prices were available ahead of time and therefore did not require temporal shifting.

Lagged target variables were also shifted to maintain consistency with operational information availability, according to Table 4.1. For imbalance price lags, which are known up to the current hour, the corresponding series were shifted so that, for a prediction at time t, the model could access lagged values from t back to $t-5$ (five lagged observations). The wind volume imbalance lags were shifted similarly, but with an additional one-hour shift to account for their longer information delay. The complete shifting scheme and the models input variables are illustrated in Figure 4.2, where the forecast horizon is set to 2 hours.

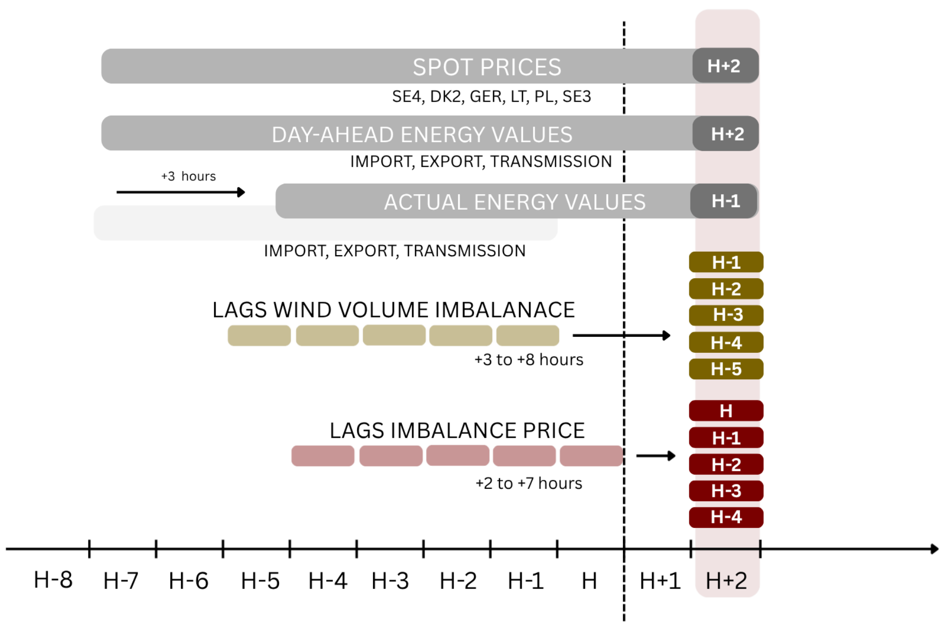

Figure 4.2: An illustrative image of how the input data was treated in the models, when predicting 2 hours ahead in time.

Lastly, to avoid train-test contamination mentioned earlier, the train data (used to train the model) and the test data was clearly separated. In the case where 2025 was the evaluated year, the training data used was between year 2021 and 2024.

4.2.2 Models

This study employed two distinct forecasting frameworks: one dedicated to predicting the imbalance price (EUR/MWh) in SE4, and another dedicated to predicting the wind imbalance volume (MWh), defined as the deviation between realized wind power generation and the day-ahead scheduled volume in SE4. Due to limited availability of historical updated weather forecast data in this study, the wind imbalance volume was modeled using the same forecasting methodology as for the imbalance price, with an identical set of models, input variables, and model configurations. The only difference concerned the lagged target variables which were the same as the target.

The complete prediction framework is illustrated in Figure 4.3, where the wind imbalance volume and the imbalance price are predicted separately. This design choice was made to enable the implementation of bidding strategies presented in Section 3.3.2, in which both the wind imbalance volume and the imbalance price had to be available as inputs. Although the wind imbalance volume may exhibit some correlation with the imbalance price in SE4, this study assumed that other factors—such as generation or network outages, consumption imbalances, and transmission constraints—together have a stronger influence on the imbalance price.

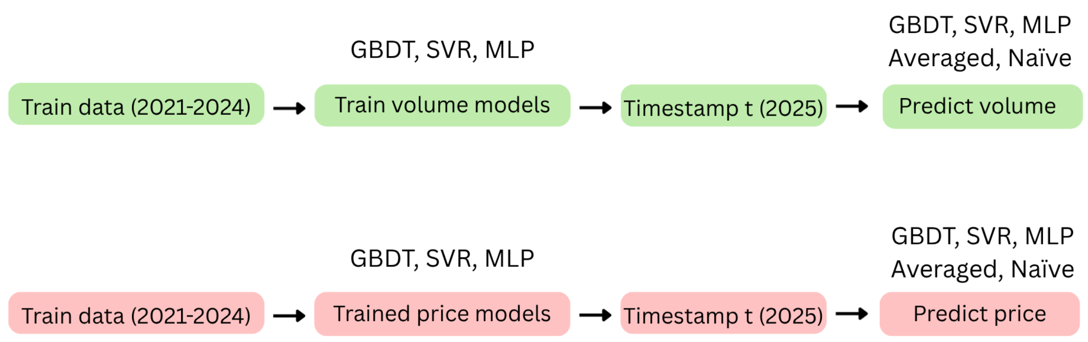

Figure 4.3: The procedure for the wind imbalance volume model (green) and imbalance price model (red).

Furthermore, the specific models used in this study was the gradient boosting decision tree (GBDT), support vector regression (SVR), and multilayer perceptron (MLP). For the year 2025, the three models were applied to predict both imbalance volume and imbalance price and in addition, an ensemble forecast was constructed by computing the arithmetic mean of the three model outputs. A Naïve prediction was also implemented, in which the predicted imbalance price and wind imbalance volume were assumed to be equal to the most recently observed value at the time of prediction.

Below, the five prediction approaches are described, together with the rationale for their use and details of their implementation. As the primary objective of the study was to predict the imbalance price, all model design choices and justifications presented below are formulated with the goal of maximizing imbalance price forecasting accuracy. Nevertheless they were used for the imbalance volume predictions.

Naïve model

Within the imbalance price forecasting, several previous studies have employed a 'naïve' model either as a benchmark against which more advanced forecasting techniques are compared [51], or to illustrate that simple methods can yield competitive performance even for complex time series [62]. In the existing literature on imbalance price prediction, the specific naïve model varies, but common variants include setting the forecast equal to the day-ahead SPOT price or to the most recently traded intraday price [51] [62].

In this study, which focuses on intraday trading in relation to the imbalance price, the Naïve price model was instead defined as the most recently observed imbalance price available at the time the forecast is produced. The underlying rationale of not using the intraday or SPOT price for predictions, was that the predicted imbalance price serves as the trading decision variable. Using the intraday price as the imbalance price prediction implied an expectation of no price differential, and using the SPOT price as the naïve prediction assumed no imbalance activations and therefore, no trading opportunities.

For the imbalance volume forecast, the Naïve model was defined in a similar manner, with the predicted volume set equal to the last observed imbalance volume at the prediction time.

GBDT model

The GBDT model was selected due to its capability to learn effectively from minimally preprocessed or raw data, as well as its ability to capture extreme events [35]. This characteristic is particularly advantageous for imbalance price forecasting, where certain price spikes often lack clear explanatory drivers. For the wind imbalance volume predictions, the GBDT model was likewise employed due to its capacity to represent abrupt changes in imbalance volumes, which may to a large extent not be directly reflected in the input variables.

The hyperparameters that were modified relative to the *scikit-learn* default configuration of the GBDT model and its associated preprocessing are listed in Table 4.2. No explicit preprocessing of the input variables was applied, as the model relied on splits based on the raw values and was therefore insensitive to scaling of the inputs [63]. The primary motivation for adjusting the hyperparameters was to better handle the volatile and weakly predictability of the imbalance price timeseries.

Table 4.2: Overview of the GBDT model. The hyperparameters changed - in relation to the pre-defined ones - are shown, together with the preprocessing.

Model	Preprocessing	Changed hyperparameters
GBDT [65]	None	`learning_rate` = 0.05 (0.1)
		`iterations` = 500 (100)
		`max_depth` = 6 (no limit)

One hyperparameter that was altered from the original *scikit-learn* implementation was the `learning_rate`, which handles the step size in the gradient descent procedure. Although values below 0.1 are generally recommended, previous studies indicated that further reductions could enhance predictive performance. Consequently, a `learning_rate` of 0.05 was adopted [64]. When the `learning_rate` was decreased, it was necessary to increase the number of trees to preserve the model's capacity to fit the data before early stopping, and therefore the number of `iterations` was increased from 100 to 500. Finally, the `max_depth`, which specifies the maximum number of splits in each decision tree, was reduced from the default "no limit" to 6. A very small `max_depth` restricts the model to capturing mainly linear or near-linear relationships, whereas very large depths increase the risk of overfitting. For the imbalance price time series, a `max_depth` of 6 was chosen as an appropriate compromise between underfitting and overfitting.

SVR model

The Support Vector Regression (SVR) model was incorporated due to its established capability to approximate nonlinear relationships, which is advantageous for forecasting imbalance prices [43]. Furthermore, SVR has previously been employed for imbalance price forecasting, where it showed competitive predictive performance [49]. An overview of the SVR configuration, including preprocessing steps and modified hyperparameters, is provided in Table 4.3.

Table 4.3: Overview of the SVR model. The hyperparameters changed - in relation to the pre-defined ones - are shown, together with the preprocessing.

Model	Preprocessing	Changed hyperparameters
SVR [66]	StandardScaler	`C` = 10.0 (1.0)

The performance of SVR is sensitive to the scaling of the input variables, as large-magnitude variables may dominate the distance computations and therefore affect the regression relative to variables with smaller absolute values. To mitigate this, a preprocessing step based on the StandardScaler from *scikit-learn* was applied, which standardized each input timeseries to have zero mean and unit variance [64].

Relative to the default configuration in the *scikit-learn* implementation, the only hyperparameter that was modified was the regularization parameter `C`. This parameter controls the trade-off between maximizing the flatness of the regression function and penalizing deviations larger than the ϵ loss. Increasing `C` reduces the tolerance for errors and places a higher penalty on deviations, thereby increasing the model's capability to represent extreme price spikes [67].

MLP model

The multilayer perceptron (MLP) model was selected to represent the neural network architectures since that are models that has frequently been reported in the literature with high predictive performance for imbalance price forecasting [50, 51]. The ability to capture extreme events, also increased the expectations of achieving high model performance during the volatile year of 2025. Since the MLP model is sensitive to the scale of the input features, it was decided to normalize the input data also with this model. An overview of the MLP configuration is provided in Table 4.4.

Table 4.4: Overview of the MLP model. The hyperparameters changed - in relation to the pre-defined ones - are shown, together with the preprocessing.

Model	Preprocessing	Changed hyperparameters
MLP [68]	StandardScaler	`hidden_layers` = [128, 64, 32] ([100,])
		`learning_rate` = 0.0005 (0.001)
		`early_stopping` = True (False)

Three hyperparameters were modified relative to the default configuration in *scikit-learn* [68]. The architecture of the `hidden_layers` was expanded to [128, 64, 32] neurons, in order to better handle the complexity of the imbalance price forecasting task, which involves a large number of input variables and complex nonlinear relationships. The specific number of layers and neurons was chosen as a compromise between model learning and computational cost, given that increased architectural complexity generally increases training time.

Furthermore, the `learning_rate` was reduced from the default value of 0.001 to 0.0005 to enable more slow weight updates, which can be advantageous when learning from noisy timeseries [68]. Finally, the `early_stopping` option was set to `True` to mitigate model overfitting.

Averaged model

An ensemble model was constructed by computing, for each forecasted timestamp, the arithmetic mean of the predictions generated by the three base models (GBDT, SVR, and MLP). The motivation was that averaging could smoothen out irregularities of individual models and thereby improve predictive performance [46]. No weights were applied when computing the average prediction at each timestamp. This choice was motivated both by the simplicity of the implementation and by previous studies indicating that weighted averaging not necessarily yield better performance [46]. The Averaged model was thus defined as:

$$pred_{\text{averaged},i} = \frac{pred_{\text{GBDT},i} + pred_{\text{SVR},i} + pred_{\text{MLP},i}}{3},$$

where i denotes a timestamp within the year 2025.

4.3 Intraday bidding model

The imbalance price and the wind imbalance volume predictions were incorporated into an intraday bidding model. The central concept was to compare the forecasted imbalance price with the current intraday price at the corresponding delivery hour, in order to determine how to allocate the forecasted imbalance volume. The procedure was based on the two strategies presented in the theory section 3.3.2: one fully based on the risk-averse strategy - named *Baseline* strategy -; and one more intelligent intraday strategy - named *Proposed* strategy - based on the risk-taking strategy. The two strategies are presented in Table 4.5.

Table 4.5: The two types of strategies evaluated in the thesis: one *Baseline* strategy and one *Proposed* strategy. Green color is a profitable intraday trade and red color is a costly intraday trade, compared to refraining from intraday trading. The blue color refers to no intraday trading.

	Under production	**Over production**
ID price > Imb. price	Buy ID	Sell ID
ID price < Imb. price	Buy ID	Sell ID

Baseline strategy
(Risk-averse strategy)

	Under production	**Over production**
ID price > Imb. price	No trade	Sell ID
ID price < Imb. price	Buy ID	No trade

Proposed strategy
(Modified from risk-taking strategy)

In the *Baseline* strategy—used as a benchmark for the *Proposed* strategy —the entire predicted imbalance wind volume was always traded in the intraday market, regardless of the relationship between the predicted imbalance price and the intraday market price. As the risk-averse strategy described in the theoretical section on intraday bidding, the primary objective of this baseline strategy was to systematically reduce any potential imbalance volume by buying or selling in the intraday market, without consideration to whether the transaction was economically advantageous or not.

This design led to two configurations that yielded a positive net profit, compared to refraining from intraday trading, highlighted in green in the left panel of Table 4.5: (i) selling an overproduction volume in the intraday market when the intraday price exceeds the imbalance price; and (ii) buying back an underproduction volume in the intraday market when the intraday price is lower than the imbalance price. The remaining two configurations generate a net cost from intraday trading, highlighted in red: (iii) buying back an underproduction volume when the intraday price is higher than the imbalance price; and (iv) selling an overproduction volume when the intraday price is lower than the imbalance price.

The *Proposed* strategy was derived from the risk-taking strategy presented in the theory section, but modified to comply with Swedish regulations for BRPs, see the right part of Table 4.5. Instead of always executing trades whenever there was a spread between the intraday price and the imbalance price (like the risk-taking strategy in the theory section), the strategy was reformulated to never increase the predicted wind imbalance volume, as this would violate the BRP agreement [54]. Consequently, two trading cases were kept but limited to only trade the imbalance volume, visual in green: (i) buying intraday volume in situations of predicted underproduction when the intraday price is lower than the predicted imbalance price; and (ii) selling intraday volume in situations of predicted overproduction when the intraday price is higher than the predicted imbalance price. Both cases simultaneously reduce the imbalance volume and yield a positive expected net profit, compared to refraining from intraday trading and let the volume be settled to the imbalance price. These two cases were the same as the two green colored situations in the *Baseline* strategy.

The two remaining configurations in the original risk-taking strategy — in which profit would be generated by trading further into imbalance — were excluded in the *Proposed* strategy. In these cases, no intraday trade was executed and the existing imbalance volume was left un-traded and settled to the imbalance price, visual in blue in the strategy table above. That meant still taking advantage of the price spread, but never increase the existing imbalance volume. SvK prohibits self-regulation, i.e., earning revenue by on purpose increasing or decreasing power production in order to use imbalance settlement [59]. However, there is no explicit regulation against refraining from intraday trading, and such inaction therefore does not imply a direct violation of the BRP agreement for a wind power producer.

The two strategies were applied with all prediction model variants: Naïve, GBDT, SVR, MLP, and the Averaged model. For the *Baseline* strategy, this yielded five result series (one for each imbalance volume prediction model). For the *Proposed* strategy, 25 result combinations were obtained, corresponding to all pairwise combinations of the five imbalance volume models and the five imbalance price models.

4.3.1 Selection of intraday price series

Since the intraday bidding model was only backtested for the year 2025, the available intraday historical price data were limited. In an ideal study, one would have access to average hourly intraday prices for all horizon hours up to the delivery hour. However, Nord Pool only provides, for each delivery hour, the closing intraday price and the intraday volume-weighted average price over the last three hours of trading (VWAP3H) [70]. Consequently, instead of analyzing the predicted imbalance price and predicted imbalance volume across multiple horizon hours while using the intraday price at each horizon, the bidding strategies in this study were constrained to use only these two available time series and their associated horizon hours.

In Scenario 1, the intraday price used in the bidding model was the closing intraday price, i.e. one of the time series published by Nord Pool. The intraday market closes one hour before the start of the delivery hour. For example, for hour 12 the last transactions for that hour occur at 10:59. Employing this closing price as the relevant intraday price in the bidding model implied a forecasting horizon of two hours for both the wind power imbalance volume prediction and the imbalance price prediction. This setup is illustrated in the upper timeline of Figure 4.4.

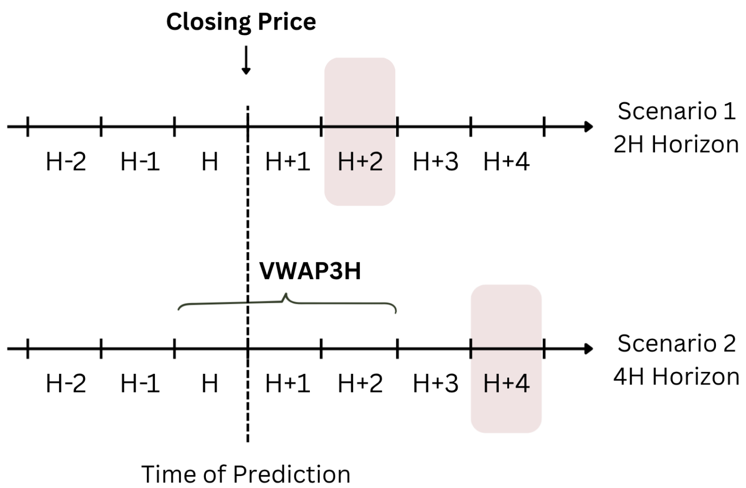

Figure 4.4: The two intraday price series used for the two forecast horizon scenarios. The closing price is used at a horizon of 2 hours and the VWAP3H price is used at a horizon of 4 hours.

In Scenario 2, the VWAP3H price was instead used and was assumed to represent the real intraday price during the final three hours of trading before market closure for the given delivery hour. The VWAP3H price is defined as the volume-weighted average intraday price over the last three hours of trading for a specific delivery hour, and thus represents the average transaction price for buying or selling energy in the intraday market over that period [70]. Because individual bids and matching prices underlying this average were not available, it was assumed that VWAP3H represents the intraday price at the midpoint of that three-hour window. Accordingly, the imbalance price and imbalance volume models had to provide forecasts with a 4-hour horizon in order for the prediction time to fall within the interval associated with this known intraday price. This configuration is illustrated in the lower timeline of Figure 4.4.

4.3.2 Cost calculations

The artificial bids and accepted energy volumes in the intraday bidding procedure were determined based on the forecasted wind imbalance volume, which could be either positive or negative. When the price signals indicated favorable conditions for intraday purchases or sales, the intraday model used only 50% of the predicted wind imbalance volume. This threshold was introduced as a conservative limit to more accurately reflect real operations, given that the intraday market is not fully liquid and not all submitted bids are necessarily accepted [55]. The results in this study therefore gives a conservative result for achievable cost savings.

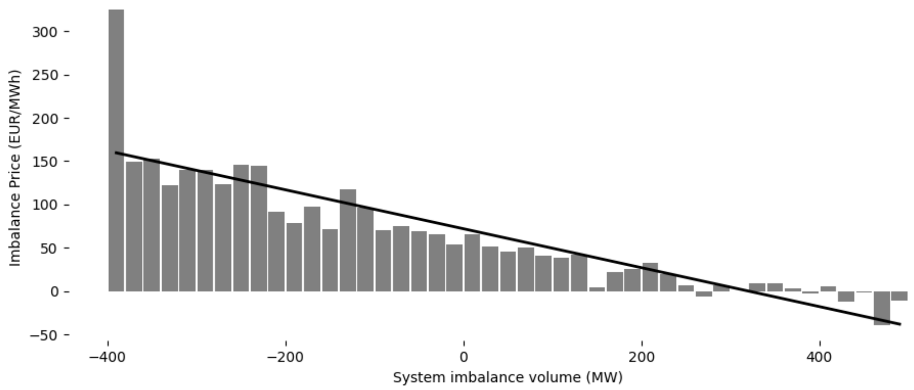

Figure 4.5: The general imbalance price change on the market for different system imbalance volume. The slope is 0.22 EUR/MWh^2.

To further simulate the reality with the intraday model, an additional adjustment was applied to the historical imbalance price whenever wind energy volumes were traded intraday, as such trades directly affect the total imbalance volume settled in the balancing market [60]. To quantify how the imbalance price in SE4 varies as a function of the imbalance volume, 2025 imbalance volume data (total grid imbalance volume) from ENTSO-E were analyzed with the corresponding imbalance prices [61]. The imbalance volumes were grouped into bins of 20 MWh with a minimum of 6 observations per bin, and all other values were excluded. A linear regression applied to the resulting binned data indicated a price decrease of 0.22 EUR/MWh for each additional MWh of imbalance power. This linear relationship was then used to compute the adjusted imbalance price when an intraday transaction of a given wind volume was executed. Figure 4.5 illustrates the binned data points and the corresponding estimated price slope.

The traded volume for each hour of the year 2025, indexed by i, was computed as:

$$\text{Traded Volume}_i\,[\text{MWh}] = x_i \cdot 0.5 \cdot \text{Predicted Diff}_i\,[\text{MWh}]$$

$$x_i = \begin{cases} 1, & \text{if an ID trade was executed,} \\ 0, & \text{if the deviation was settled in the balancing market} \end{cases}$$

Here, x_i was a Boolean decision variable indicating whether an intraday trade was conducted ($x_i = 1$) or whether the entire imbalance volume was settled in the balancing market ($x_i = 0$). For two cases in the *Proposed* strategy (the no trade cases) the x_i variable was set to 0, while for the *Baseline* strategy, the x_i variable was always set to 1.

The traded volume was then multiplied by the intraday price and aggregated over the whole year:

$$\text{Income ID}\,[\text{EUR}] = \sum_{i=1}^{n} \text{Traded Volume}_i\,[\text{MWh}] \;\cdot\; \text{ID Price}_i\,[\text{EUR/MWh}]$$

where ID Price_i corresponded either to the closing intraday price (Scenario 1) or to the VWAP3H price (Scenario 2). A positive value of Income ID represented a net intraday trading revenue, whereas a negative value indicated a net intraday trading cost.

The revised imbalance revenue was then calculated based on the updated imbalance volume (after possible intraday trading) and a modified imbalance price obtained by applying the $0.22\,\text{EUR/MWh}^2$ factor:

$$\begin{gathered} \text{New Imbalance Volume}_i\,[\text{MWh}] = \\ (\ \text{Actual Volume}_i\,[\text{MWh}] \;-\; \text{DA Volume}_i\,[\text{MWh}]\ ) \;-\; \text{Traded Volume}_i\,[\text{MWh}] \end{gathered}$$

$$\begin{gathered} \text{New Imbalance Price}_i\,[\text{EUR/MWh}] = \\ \text{Original Imbalance Price}_i\,[\text{EUR/MWh}] \\ +\ (\ \text{Traded Volume}_i\,[\text{MWh}] \;\cdot\; 0.22\,[\text{EUR/MWh}^2]\ ) \end{gathered}$$

The total revised imbalance revenue over the full year was then given by:

$$\text{New Imbalance Income}\,[\text{EUR}] = \\ \sum_{i=1}^{n} \text{New Imbalance Volume}_i\,[\text{MWh}] \;\cdot\; \text{New Imbalance Price}_i\,[\text{EUR/MWh}]$$

where a positive value corresponded to a net imbalance revenue, whereas a negative value indicated a net imbalance cost. Subsequently, the sum of the intraday trading income and the revised imbalance income was compared with the original imbalance income (i.e., the case without intraday trading) for the year 2025. The original imbalance income applied in this comparison was −3,700,000 EUR, as shown in Figure 2.7.

The entire procedure was repeated for each of the five imbalance volume models in the *Baseline* strategy, as well as for all 25 combinations of imbalance volume and imbalance price forecasts in the *Proposed* strategy. This analysis was then conducted for both intraday price time series, yielding two distinct scenarios: a 2-hour forecast horizon (Scenario 1) and a 4-hour forecast horizon (Scenario 2).

4.3.3 Volume calculations

The impact of intraday trading on wind power imbalance volumes was also evaluated, as it was necessary to demonstrate that the proposed models did not generate trading decisions that resulted in larger imbalance volumes for wind power producers than in a scenario without intraday trading. This requirement arised from the need to comply with the obligations specified in the BRP agreement [54].

The original imbalance volume, corresponding to a situation without intraday trading, was first calculated and subsequently used to compute the root mean square error (RMSE) of the imbalance volume:

$$\text{Original Imbalance Volume}_i\,[\text{MWh}] = \text{Actual Volume}_i\,[\text{MWh}] - \text{DA Volume}_i\,[\text{MWh}]$$

$$\text{RMSE original volume}\,[\text{MWh}] = \sqrt{\frac{1}{n}\sum_{i=1}^{n} \left(\text{Original Imbalance Volume}_i\,[\text{MWh}]\right)^2}\,.$$

The new imbalance volume after the intraday trades was then computed and used to derive the RMSE of the adjusted imbalance volume:

$$\text{New Imbalance Volume}_i\,[\text{MWh}] = \\ \text{Original Imbalance Volume}_i\,[\text{MWh}] - \text{Traded Volume}_i\,[\text{MWh}]$$

$$\text{RMSE new volume}\,[\text{MWh}] = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\text{New Imbalance Volume}_i\,[\text{MWh}])^2}\,.$$

The RMSE value of the new imbalance volume was then benchmarked against the original volume error and the percent decrease of the imbalance volume after the trade was calculated. The entire procedure was repeated for each of the five imbalance volume models in the *Baseline* strategy, as well as for all 25 combinations of imbalance volume and imbalance price forecasts in the *Proposed* strategy. This analysis was then conducted for both the 2-hour forecast horizon and a 4-hour forecast horizon.

4.4 Forecasting models evaluation

This section describes the methods for: evaluating the prediction models performance; how the price model performance has changed under the regime changes of the imbalance market; a sensitivity analysis of the hyperparameters; and model aging evaluation.

4.4.1 Model performance

The imbalance price and wind imbalance volume prediction models were all evaluated with RMSE and MAE, for how good their predictions were compared to the actual value. The two evaluation values were calculated accordingly:

$$\text{RMSE}\,[\text{MWh or EUR/MWh}] = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\text{Value}_i - \text{Predicted Value}_i)^2}$$

$$\text{MAE}\,[\text{MWh or EUR/MWh}] = \frac{1}{n}\sum_{i=1}^{n}|\text{Value}_i - \text{Predicted Value}_i|$$

where the Value_i was the actual value for a certain timestamp i under 2025, either the imblanace price or the wind imbalance volume. The Predicted value_i was instead the predicted imbalance price or the predicted wind imbalance volume. This was done for all the five price models and the same five volume models. The models were assessed on the two different horizon hours used in the intraday bidding strategies, namely 2-hour and 4-hour horizons.

4.4.2 Evaluation of price model under regime changes

The 2025 imbalance price model performance was additionally compared with its performance in 2024, i.e., prior to the major regime-changing implementations. For this extra evaluation, only the *Averaged* model was considered. When 2024 was used as the test year, the training data were set to end in 2023, thereby preventing any information leakage into the model.

To enable a meaningful comparison between the 2025 imbalance price predictions and the corresponding 2024 predictions—beyond a simple comparison of absolute error magnitudes—the model outputs were benchmarked against SPOT price–based predictions for both years. Since the imbalance price is set to the SPOT price in hours without balancing activation [29], the SPOT price provides a relevant reference baseline. This baseline has also been adopted in previous studies when constructing benchmark comparisons [51], as it is typically known in advance and therefore gives a reference that does not vary with forecast horizon. In the present study, an objective was to outperform this baseline.

The evaluation of the averaged imbalance price model was conducted over forecast horizons from 1 to 24 hours, i.e., predicting the imbalance price from one hour ahead up to 24 hours ahead. For each forecast horizon, the RMSE of the model was computed and compared with the RMSE obtained from the SPOT price benchmark. Subsequently, the RMSE of the averaged model was expressed as a percentage of the SPOT RMSE. This procedure was carried out separately for 2024 and 2025, and the resulting relative RMSE values were visualized in a line chart, see Figure 5.11 .

4.4.3 Sensitivity analysis

A sensitivity analysis of the imbalance price and wind imbalance volume models was conducted to assess the selected hyperparameters in the predictive models. For each model (GBDT, SVR, and MLP), the three most influential hyperparameters with respect to predictive accuracy were identified. For each of these hyperparameters, three distinct values were evaluated (with the exception of the early stopping criterion in the MLP model). It was the value used in the study, as well as one larger and one smaller value. The models were then trained and tested for all 27 possible hyperparameter combinations, for both price and volume prediction, using a 2-hour forecasting horizon over the year 2025.

For the GBDT model, the three hyperparameters varied in the sensitivity analysis are reported in Table 4.6. These hyperparameters were specifically selected because they had been modified relative to the default settings provided in [65], thereby offering a basis to evaluate how accurately the theory informed the model settings.

Table 4.6: Overview of the choosen hyperparameters for sensitivity analysis for the GBDT model. The values in bold are the ones used for the results in this study.

Model	Changed hyperparameters	Values
GBDT [65]	`learning_rate`	[0.03, **0.05**, 0.1]
	`max_iterations`	[300, **500**, 800]
	`max_depth`	[3, **6**, 9]

For the SVR model, the only hyperparameter that differed from the default configuration in [66] was the regularization parameter `C`, making it a candidate for the sensitivity analysis. The epsilon parameter was additionally considered due to its role in determining the magnitude of errors disregarded during training. The gamma parameter was included because it governs the complexity of the decision function, i.e., how good the model can fit local patterns in the data. All hyperparameter variations considered in the SVR sensitivity analysis are summarized in Table 4.7.

Table 4.7: Overview of the choosen hyperparameters for sensitivity analysis for the SVR model. The values in bold are the ones used for the results in this study.

Model	Changed hyperparameters	Values
SVR [66]	`C`	[1, **10**, 100]
	`epsilon`	[0.01, **0.1**, 1.0]
	`gamma`	[0.01, **scale**, 0.1]

For the MLP model, the hyperparameters examined in the sensitivity analysis were, analogously to the GBDT case, those that had been changed relative to the default settings. These hyperparameters are listed in Table 4.8.

Table 4.8: Overview of the choosen hyperparameters for sensitivity analysis for the MLP model. The values in bold are the ones used for the results in this study.

Model	Changed hyperparameters	Values
MLP [68]	`hidden_layers`	[(64,32), **(128,64,32)**, (256,128,64)]
	`learning_rate`	[0.0001, **0.0005**, 0.001]
	`early_stopping`	[False, **True**]

4.4.4 Model aging evaluation

The aging of the imbalance price and imbalance volume forecasting models was evaluated to quantify how their predictive accuracy changed from the beginning to the end of the year 2025. This was assessed by applying the price and volume models to the last six months of 2025, while varying the age of the models, i.e., the gap between the end of the training period and the start of the test month. Each of the six months was considered separately. The procedure began with July, and for each subsequent month, the corresponding training set was extended by one month relative to the previous test month.

Four model ages were analyzed: 0 months, 3 months, 6 months, and 12 months. For the 0-month age setting, the models were trained using data up to the end of the month before the test month. For the other age settings, the training period ended the specified number of months before the beginning of the test month. For each age, the root mean squared error (RMSE) was computed for each of the six test months and then averaged across these months. This analysis was executed for the Gradient Boosted Decision Trees (GBDT), Support Vector Regression (SVR), Multilayer Perceptron (MLP), and Averaged models, for both the imbalance price and imbalance volume forecasts at a 2-hour prediction horizon.

5
Results and Discussion

This chapter is structured into several subsections: the performance of the imbalance price and imbalance volume forecasting; the results from the 2-hour and 4-hour scenarios; a detailed subsection about the trading decisions; a robustness and sensitivity analysis; and a comparison with previous studies.

5.1 Predictive performance

The development and evaluation of price and volume forecasting models constitute a central component of this study. In this section, the price and volume prediction model's performance will be presented.

5.1.1 Performance of price models

Figure 5.1 illustrates an example of how the forecasting of the imbalance price was executed in practice, showing the first two weeks of August 2025. The orange line represents the actual imbalance price, and the blue line a prediction from the Averaged model, at a 2-hour forecasting horizon.

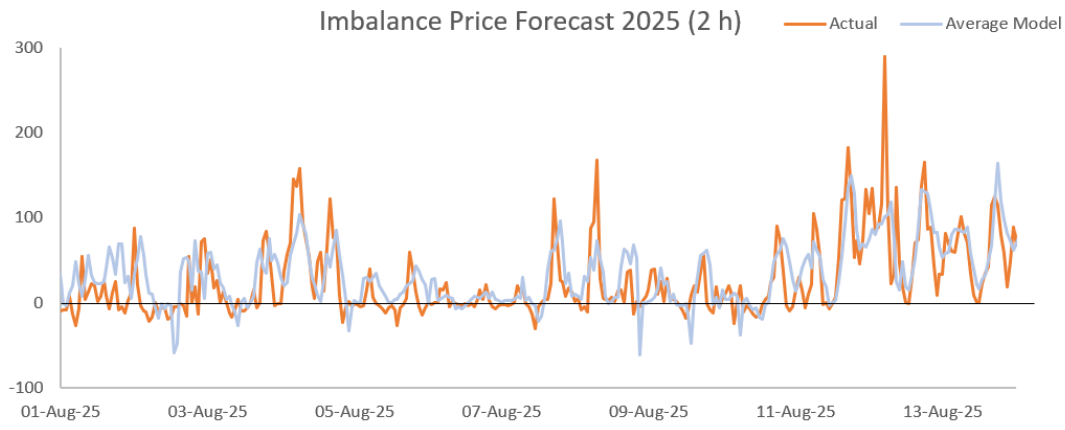

Figure 5.1: An example of price forecast modelling with the Averaged model compared to the actual imbalance price. Example is from first week in August 2025.

The figure demonstrates that the model generally tracks the realized imbalance price closely. Furthermore, it is able, to some extent, to anticipate both upward and downward price spikes. However, for most of the extreme events, the forecasts have a systematic delay corresponding to the 2-hour forecast horizon. That due to the strong influence of the lagged target variable inputs, that is the latest known imbalance price value at the time of prediction. This shifting in price spike predictions, introduces a risk that intraday trading is executed on wrong price spreads, and thus leading to non-profitable decisions.

In Table 5.1, the RMSE and MAE values for the imbalance price prediction are reported for the five individual models, with the lowest values highlighted in bold to indicate the best performing models. As expected, the RMSE values are consistently higher than the MAE values due to the squaring of the error term in RMSE, which penalizes larger prediction errors more strongly than the MAE. Moreover, both metrics produce the same ranking of models, suggesting that the models have a similar capability in handling both predictions during large deviations (peaks) and during lower imbalance price periods.

Table 5.1: The price model performance for 2h and 4h horizon time.

Price Models 2025	**RMSE 2h**	**MAE 2h**	**RMSE 4h**	**MAE 4h**
Naïve	187,9	62,0	203,6	75,6
GBDT	146,4	50,0	**145,3**	**48,0**
SVR	144,4	48,6	146,9	51,4
MLP	147,0	51,0	147,3	51,4
Averaged	**142,9**	**47,2**	145,8	49,3

For the 2-hour forecast horizon, the Averaged model achieves the best performance according to both RMSE and MAE, whereas the Naïve model gives the poorest performance overall. It is also shown that the remaining three models (GBDT, SVR, and MLP) perform more similarly to the Averaged model than to the Naïve model, which has clearly worse performance. This improvement can be explained by the input variables used by the thinking models, which increases predictive accuracy by adding market and volume information.

For the 4-hour forecast horizon, the GBDT model has the highest performance, with the Averaged model ranking second and achieving a comparably error. Relative to the 2-hour horizon, all error metrics increase, which can be explained by the greater difficulty of accurately forecasting imbalance prices at longer horizons.

Across both forecast horizons, the four thinking models (GBDT, SVR, MLP, and Averaged) exhibit similar error levels, differing only by a few percent in performance. This indicates that the models extract and exploit largely similar information from the input data and that additional model complexity, as in the MLP, does not give a substantial performance gain in price prediction forecasting.

5.1.2 Performance of volume models

Since the imbalance volume predictions are generated in this study solely due to the absence of updated weather forecasts (which would normally be used to determine the wind imbalance volume), the evaluation of the volume model's performance is less comprehensive than that of the imbalance price prediction model. Nonetheless, examining performance remains relevant, as the predicted imbalance volume influences intraday bidding strategies and imbalance cost reductions.

The RMSE and MAE of the imbalance volume predictions at forecast horizons of 2 and 4 hours for the year 2025 are reported in Table 5.2. For both horizons, the Averaged model gives the best performance in forecasting the day-ahead wind imbalance volume among all considered models, whereas the Naïve model consistently yields lowest performance. Because the imbalance volume time series is less volatile than what is observed in the imbalance price series, the MAE and RMSE values are proportionally more similar to each other than in the case of the price model.

Table 5.2: The wind volume model performance for 2h and 4h horizon time.

Volume Models 2025	**RMSE 2h**	**MAE 2h**	**RMSE 4h**	**MAE 4h**
Naïve	49,0	29,2	60,2	37,1
GBDT	36,9	22,7	38,4	24,3
SVR	40,4	24,7	45,1	28,3
MLP	40,7	27,1	45,7	31,8
Averaged	**35,0**	**22,5**	**36,5**	**23,8**

A further observation regarding the volume model's performance is that the GBDT model outperforms both the SVR and MLP models at each horizon, which is reflected in the both the low RMSE and MAE values. It is thus the model that contributes most to the strong performance of the Averaged model.

5.2 Intraday scenario 1 (2-hour horizon)

In this section, the 2-hour horizon scenario will be more detailed evaluated for different model combinations, both for the *Baseline* strategy and the *Proposed* strategy. The imbalance cost reduction numbers and the wind imbalance volumes changes will be presented.

5.2.1 Imbalance cost reduction

In the heatmap in Figure 5.2, the imbalance cost reduction for the 2-hour forecasting horizon (Scenario 1) is shown for different model combinations. All values are expressed in percent and indicate each model's cost reduction relative to the imbalance level in 2025 reported in Figure 2.7. The five combinations in the right part of the figure, enclosed by the blue dotted contour, represent the *Baseline* strategy with five different volume forecasts. The 25 combinations within the dotted black contour correspond to the *Proposed* strategy. Within each of these two strategies, the model combination achieving the highest imbalance cost reduction is highlighted (with a solid outline).

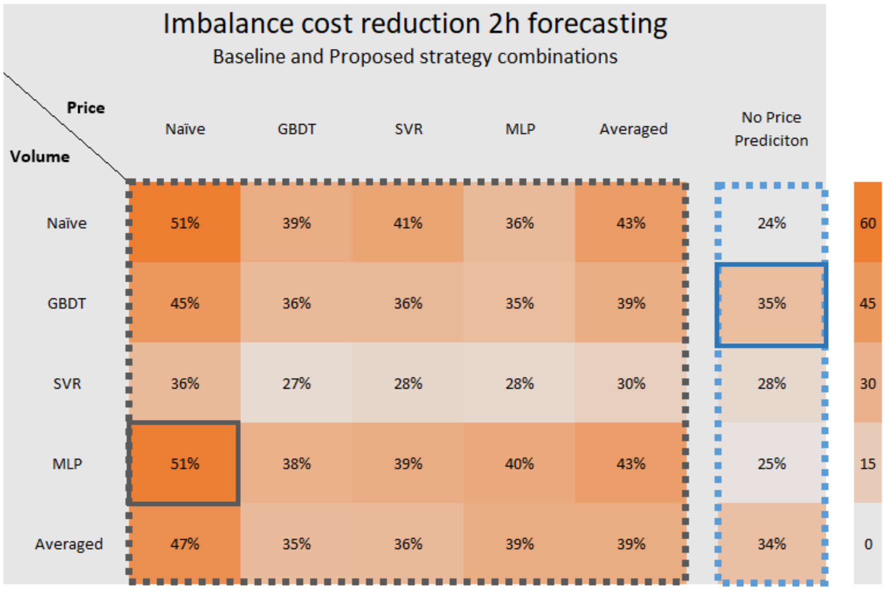

Figure 5.2: Reduction of day-ahead imbalance costs due to intraday trading under 2-hour forecasting in 2025, including both the decrease in imbalance costs and the net revenues from intraday trading. The blue dotted contour encloses the *Baseline* strategy combinations, and the black dotted contour encloses the *Proposed* strategy combinations.

From Figure 5.2, it is clear that certain strategy combinations gives higher imbalance cost reduction performance than others. Within the five *Baseline* strategy configurations, the best-performing case is obtained when the GBDT model is used for volume prediction, yielding a 35% reduction in imbalance costs relative to the case without any intraday trading. Most of the model combinations under the *Proposed* strategy achieve even greater cost reductions than this best *Baseline* configuration, with the maximum reduction of 51% observed when the Naïve price prediction model is combined with the MLP volume prediction model.

For the *Proposed* strategy overall, the SVR volume model consistently exhibits the poorest performance in terms of cost reduction, resulting in the lowest improvement among all volume models. In contrast, the Naïve price model delivers robust cost reduction performance in combination with all volume models. This suggests that the Naïve price model is particularly successful at capturing imbalance price spikes over the year, since it is already clear that its overall predictive accuracy is worse than for the other models.

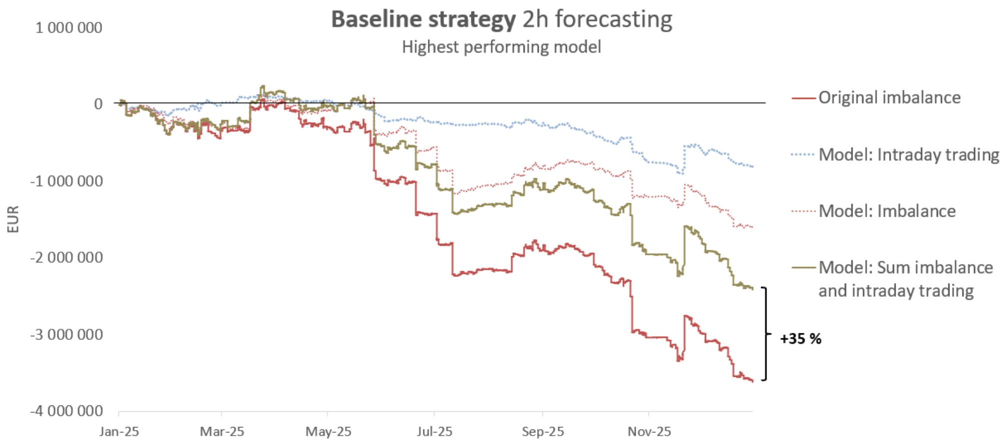

Figure 5.3: Highest cost reduction in the *Baseline* strategy at 2-hour horizon, represented by the GBDT volume model. The total imbalance cost reduction is 35%.

Figure 5.4: Highest cost reduction in the *Proposed* strategy at 2-hour horizon, represented by the MLP volume model and the Naïve price model. The total imbalance cost reduction is 51%.

Furthermore, the two best-performing configurations, highlighted with solid lines in Figure 5.2, are illustrated in the line charts in Figure 5.3 for the *Baseline* strategy and Figure 5.4 for the *Proposed* strategy. Each line chart contains four accumulated cost time series: *Original imbalance*, representing the imbalance settlement cost in the absence of intraday trading; *Model: Intraday trading*, representing the revenues and costs associated with intraday transactions executed by the model; *Model: Imbalance*, representing the imbalance settlement cost after intraday trading; and *Model: Sum imbalance and intraday trading*, defined as the sum of the latter two series. The original imbalance time series is identical in both figures, whereas the remaining cost series differ between the two strategies.

Several differences are observable when comparing the two line charts corresponding to the highest performing combination of each strategy. The models are based on distinct risk-management approaches: the *Baseline* strategy is more risk-averse (consistently reducing the imbalance volume), whereas the *Proposed* strategy is more risk-seeking. This distinction is visible in the dotted blue line, which represents the intraday trading revenues and costs.

In the *Baseline* strategy, the intraday trading result is negative (approximately –1,000,000 EUR), indicating that a substantial number of intraday transactions are executed at a net cost in order to reduce the imbalance volume. In contrast, under the *Proposed* strategy, intraday trading yields a positive net income over the year, showing that only trades with high expected profitability are executed intraday.

This difference in trading behavior is clearly reflected in the resulting imbalance costs. The *Baseline* strategy yields substantially lower imbalance costs - around -2,000,000 EUR - compared with the *Proposed* strategy, which incurs approximately -2,500,000 EUR in imbalance costs. Despite the fact that the *Proposed* strategy takes on more risk and achieves a smaller reduction in imbalance settlement costs, the overall cost reduction (defined as the sum of imbalance settlement and net intraday costs relative to the original imbalance cost) is greater for the *Proposed* strategy. These results indicate that accepting a controlled level of risk in the imbalance market can be economically beneficial.

5.2.2 Imbalance volume change

The model configurations in the heatmap above are also presented in Figure 5.5, where the magnitude of the reduction in the day-ahead imbalance volume across the various trading scenarios are shown instead.

The percentage values represent the reduction from the original imbalance volume (i.e., the deviation between day-ahead scheduled and actual production) to the imbalance volume after intraday trading, all quantified using RMSE. Higher percentage values indicate a greater reduction in imbalance volume, and a value of 100% would correspond to intraday trading achieving a fully balanced position. As before, the configurations within the black dotted lines correspond to the *Proposed* strategy, consisting of 25 model combinations, whereas the configurations within the blue dotted lines denote the 5 *Baseline* strategy combinations.

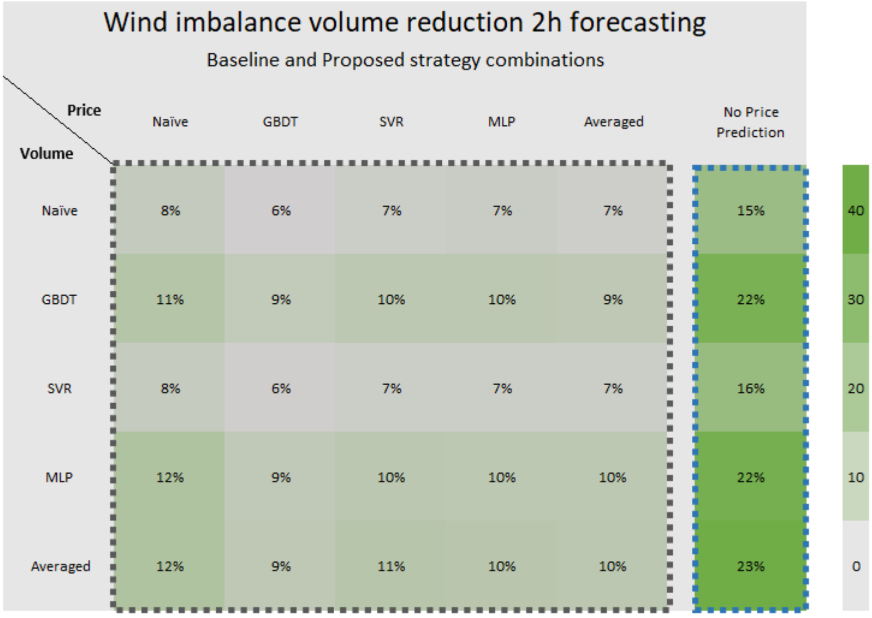

Figure 5.5: Reduction of day-ahead wind imbalance volume resulting from intraday trading under the 2-hour forecasting, for different combinations of imbalance price and volume models. The blue dotted contour encloses the *Baseline* strategy combinations, and the black dotted contour encloses the *Proposed* strategy combinations.

The results indicate that the largest reduction in wind power imbalance volume is achieved under the *Baseline* strategy, with reductions of approximately 15–25%. Under the *Proposed* strategy, the reduction in wind imbalance volume is more modest, on the order of 5–10%. This outcome is expected, as substantially fewer intraday trades are executed in this case, particularly in periods when the price spread showed for settling in the balancing market.

Figure 5.5 illustrates why the *Baseline* strategy is more advantageous from an electricity market perspective and more closely aligned with the BRP requirement to trade into balance to the greatest extent possible. Nonetheless, the *Proposed* strategy combinations do not lead to negative effects over the year as a whole since they do not increase the wind imbalance volume and therefore do not clearly violate the agreement. For example, the best-performing model in terms of cost reduction within the *Proposed* strategy—combining the MLP-based volume prediction and Naïve price prediction—achieves an average reduction in wind imbalance volume of 12% during 2025. This indicates that the intraday trading can be beneficial both financially and for the energy market.

5.3 Intraday scenario 2 (4-hour horizon)

In this section, the 4-hour horizon scenario will be more detailed evaluated for different model combinations, both for the *Baseline* strategy and the *Proposed* strategy.

The imbalance cost reduction numbers and the wind imbalance volumes changes will be presented.

5.3.1 Imbalance cost reduction

The heatmap in Figure 5.6 illustrates the imbalance cost reduction given with the 4-hour forecasting horizon (Scenario 2) in 2025 for the various combinations of price and volume forecasting models. All values are expressed in percentage terms and represent each model's relative cost reduction with respect to the 2025 imbalance level reported in Figure 2.7. In line with Scenario 1, the 25 combinations enclosed by the black dotted line correspond to the *Proposed* strategy, while the 5 combinations within the blue dotted line correspond to the *Baseline* strategy. The best performing combination within each strategy is highlighted with a solid outline.

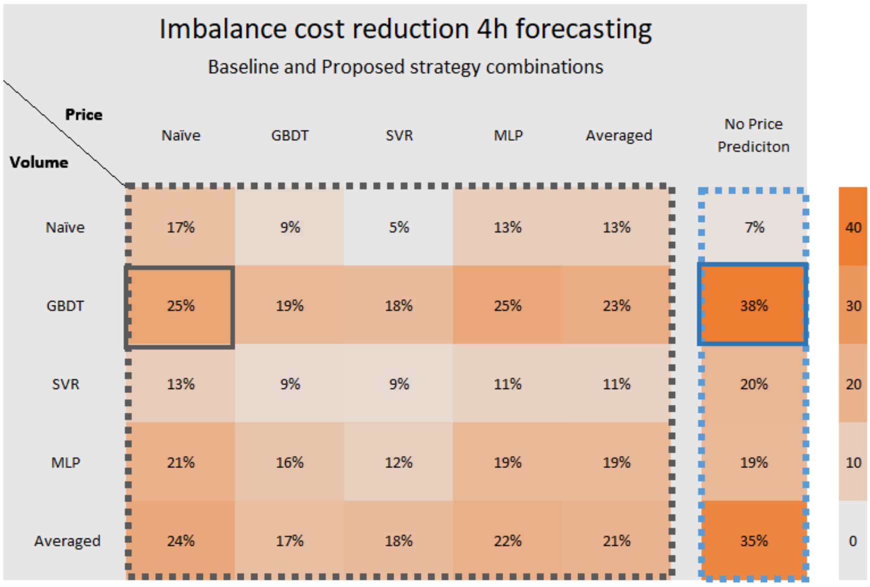

Figure 5.6: Reduction of day-ahead imbalance costs due to intraday trading under 4-hour forecasting in 2025, including both the decrease in imbalance costs and the net revenues from intraday trading. The blue dotted contour encloses the *Baseline* strategy combinations, and the black dotted contour encloses the *Proposed* strategy combinations.

For the *Baseline* strategy at the 4-hour horizon, the GBDT volume forecasting model again yields the best performance, achieving a cost reduction of 38% relative to the no-trading scenario. This value is similar to the highest cost reductions obtained under the *Baseline* strategy in the 2-hour forecasting scenario (Figure 5.2). However, the figure clearly shows that the remaining *Baseline* strategy configurations — specifically the Naïve, SVR, and MLP models — give lower cost reduction performance at the 4-hour horizon compared with the 2-hour horizon.

In contrast, the *Proposed* strategy, represented by the 25 combinations within the black dotted line in Figure 5.6, achieves lower cost reductions than the *Baseline* strategy at the 4-hour horizon, which is the opposite of what was observed for the 2-hour horizon. In this case, the best-performing model combination—GBDT for volume forecasting and the Naïve model for price forecasting—reduces imbalance costs by only 25%. Furthermore, it is evident that within the *Proposed* strategy, the SVR volume forecasting model is associated with small cost reductions, consistent with the observations from the 2-hour scenario.

The overall lower cost-reduction performance of the *Proposed* strategy at the 4-hour forecast horizon — relative to the 2-hour horizon - can be explained by a degradation in the quality of the imbalance price forecasts. This is supported by the observation that the *Baseline* Strategy continues to give comparatively results for some of the models, indicating that the reduced cost savings in the *Proposed* Strategy mailnly arises from trades based on inaccurate predicted imbalance prices. An additional reason may be that, at the 4-hour horizon, the price spread between the predicted imbalance price and the intraday price (here represented by the intraday VWAP3h) is neither as large nor as favorable as the spread used in the 2-hour forecasting case that was based on the intraday closing price.

The model combinations yielding the highest cost reductions within the *Baseline* and *Proposed* strategies are shown in Figure 5.7 and Figure 5.8, respectively. These correspond to the configurations highlighted with solid lines in Figure 5.6, achieving cost reductions of 38% and 25%.

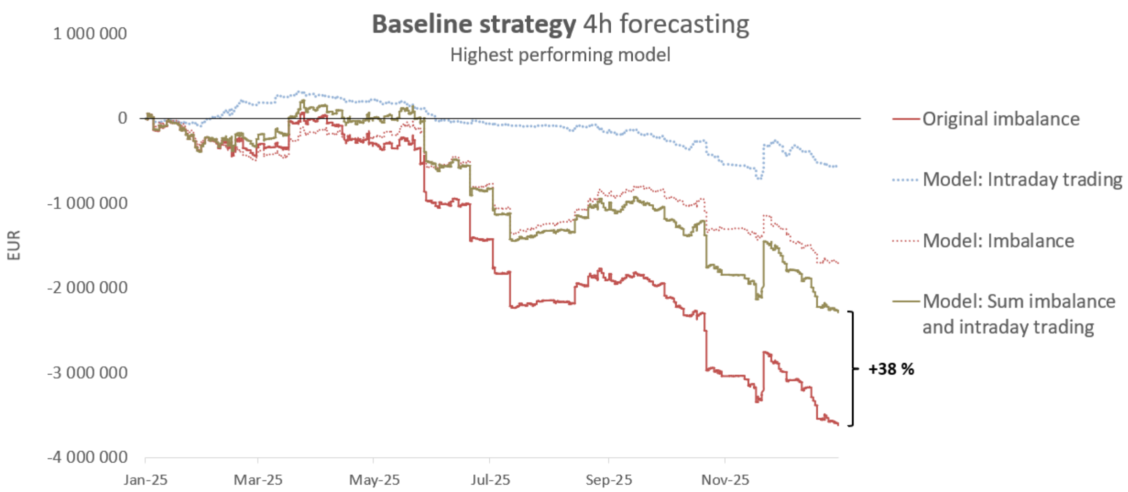

Figure 5.7: Highest cost reduction in the *Baseline* strategy at 4-hour horizon, represented by the GBDT volume model. The total imbalance cost reduction is 38%.

Figure 5.8: Highest cost reduction in the *Proposed* strategy at 4-hour horizon, represented by the GBDT volume model and the Naïve price model. The total imbalance cost reduction is 25%.

Analogous to the analysis in the previous section for the 2-hour forecast horizon, the line charts show the original imbalance settlement cost, the cost (or income) from intraday trading, the new imbalance settlement cost, and the sum of the latter two. In the risk-averse *Baseline* strategy (Figure 5.7), it is visible that over the full year, this strategy gives a net cost from intraday trading (on the order of -1,000,000 EUR). This results from a large number of trades being executed at relatively unfavorable prices in order to reduce the imbalance volume. Nevertheless, the new imbalance cost is substantially reduced, leading to an overall cost reduction of 38%.

By contrast, in the *Proposed* Strategy (Figure 5.8), intraday trading generates a net income, of approximately 1,000,000 EUR, since only the most profitable trades are executed. However, in this 4-hour horizon setting, the resulting new imbalance settlement cost remains at nearly the same level as the original imbalance settlement cost, so that the total cost reduction remains relatively low, 25%. This illustrates that the *Proposed*, more risk-seeking strategy is only advantageous when the trading decisions are based on more accurate imbalance price forecasts, which appears not to be the case in this 4-hour horizon scenario.

5.3.2 Imbalance volume change

The change in wind imbalance volume for the 4-hour forecasting horizon is presented in Figure 5.9. The figure illustrates the extent to which the original wind imbalance volume is reduced through intraday trading across the different strategies and their combinations, as described in Section 5.2.2. A value of 100% corresponds to a complete mitigation of the initial wind imbalance volume during intraday trading, such that no imbalance volume remains to be settled in the balancing market.

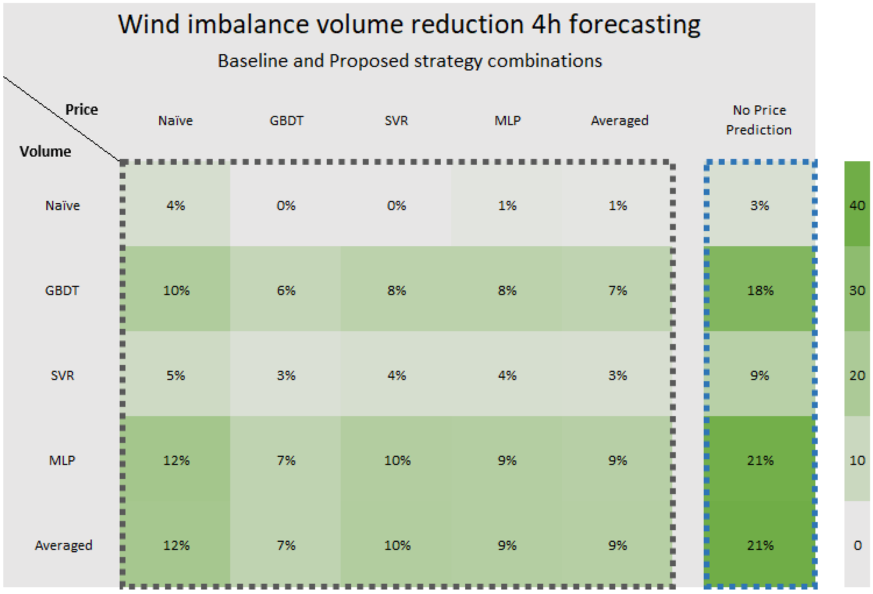

Figure 5.9: Reduction of day-ahead wind imbalance volume resulting from intraday trading under the 4-hour forecasting, for different combinations of imbalance price and volume models. The blue dotted contour encloses the *Baseline* strategy combinations, and the black dotted contour encloses the *Proposed* strategy combinations.

Analogously to the 2-hour forecasting scenario, the heatmap for the 4-hour horizon indicates a greater volume reduction under the *Baseline* strategy (the five configurations enclosed by the blue dotted line) than under the *Proposed* strategy (the 25 combinations enclosed by the black dotted line). This outcome is again explained by the larger number of trades executed within the *Baseline* strategy in order to reduce imbalance volumes.

For the *Baseline* strategy at the 4-hour horizon, the five volume reduction values are lower than those obtained for the 2-hour horizon, which can be explained by the small deterioration in volume forecast accuracy at longer horizons. By contrast, the results for the Naïve volume model change substantially, suggesting that at extended horizons, volume forecasts based solely on historical values provide limited predictive value.

Within the *Proposed* strategy combinations, many configurations give only small volume reductions, in some cases negligible or only a few percentage points. This represents a large difference from the 2-hour forecasting results and is explained by the smaller cost reductions observed at the 4-hour horizon. Nevertheless, for the 4-hour horizon, none of the *Proposed* strategy combinations results in a negative volume reduction, i.e., an increase in imbalance volume after intraday trading that would make a clear violation of the BRP agreement.

5.4 Intraday bidding selections

The best cost reduction results, outlined in Figure 5.2 for the 2-hour forecast horizon, and in Figure 5.6 for the 4-hour forecast horizon, are presented below in Figure 5.10. Each point in the subfigures represents a delivery hour in 2025. The y-axis represents the price differential between the intraday market price (either two or four hours before delivery depending one the graph) and the realized imbalance price, while the x-axis represents the realized wind power imbalance volume (overproduction or underproduction).

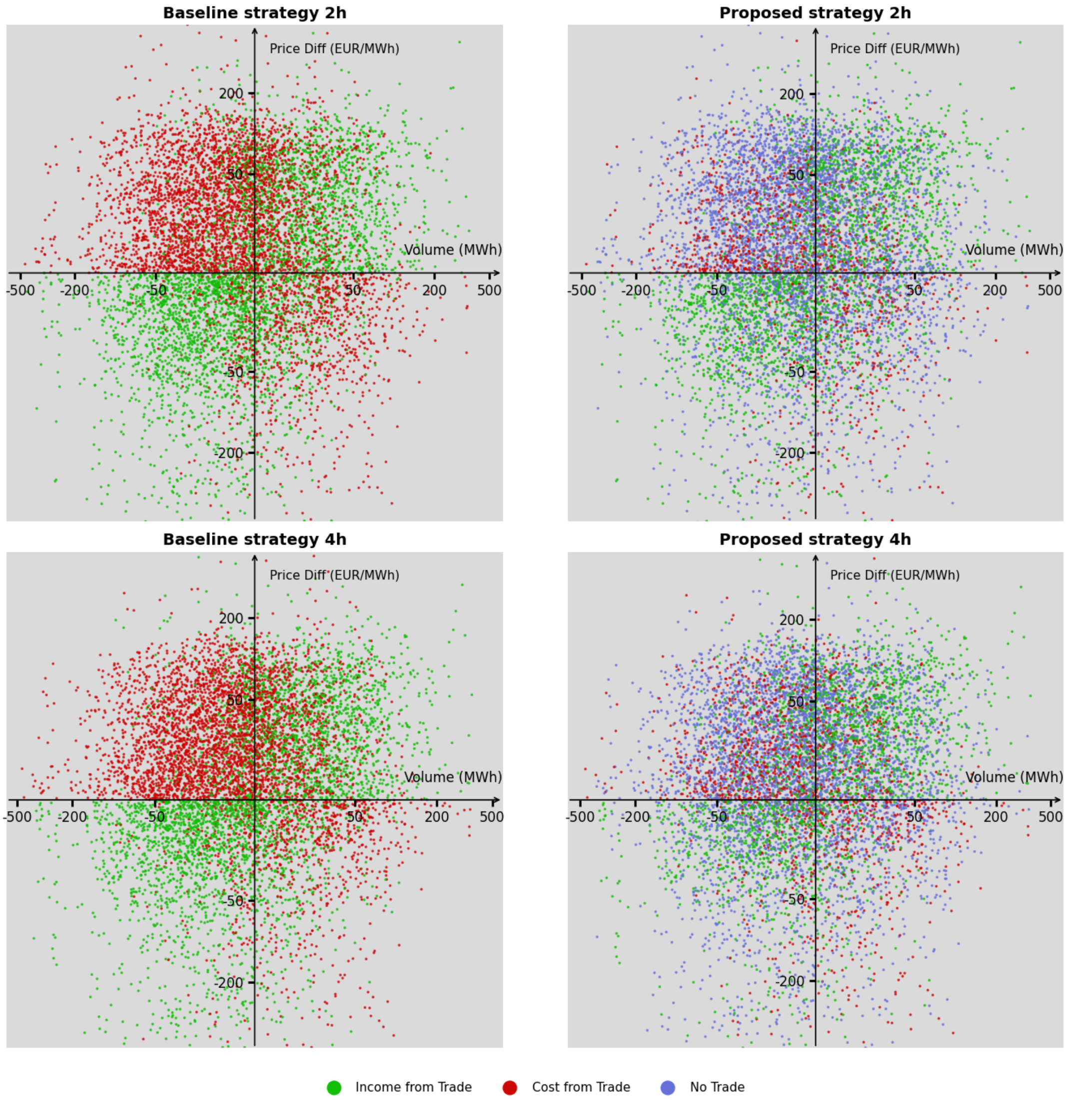

Figure 5.10: A scatter plot over the best model configurations for the *Baseline* and *Proposed* strategy for the 2 and 4-hour forecasting.

All the dots are located at the same places in the subfigures and the only aspect that differs is the color, representing how the models made decisions. The colors follow the same convention for the quadrants as in the Methods section, where the two strategies were introduced in Table 4.5. Red points indicate trades that are more costly than refraining from intraday trading, whereas green points indicate trades that are more profitable than not trading. Blue points appear only for the *Proposed* strategy and again represent hours without any intraday trading.

The two left-hand panels of Figure 5.10 illustrate the best-performing *Baseline* strategies under the 2-hour and 4-hour forecasting scenarios, highlighted earlier in the heatmaps as the 35% and 38% cost reduction cases. In the *Baseline* strategy, trading decisions are solely based on forecasts of the imbalance volume, that is, predictions of the x-axis values in the figure. A positive forecast (overproduction) triggers a sell order in the intraday market, whereas a negative forecast (underproduction) triggers a buy order. Since the model does not incorporate information on the expected price spread (y-axis), trades fall into two quadrants where they are loss-making (red) and into the other two quadrants where they are profitable (green). The mixing of both red and green points within the same quadrant is based on volume prediction model error, i.e., instances where overproduction is predicted when underproduction actually occurs, and vice versa. A red trade in the quadrants one and three, or a green trade in quadrants 2 and 4, indicates this prediction error and increases both the imbalance settlement volume and increases the cost for wind power.

For both of the *Baseline* strategies, the mixing of the colors is higher around the y-axis, suggesting that the model struggles with predicting the sign of the volume during lower magnitudes of imbalance volumes. However during larger magnitudes of imbalance volume, there is less color mixing, and thus in those cases more accurate prediction of the volume direction. It is also visible that there is more mixing on the right side of the y-axis compared to the left side, indicating that the volume model more often predicts the overproduction cases wrongly compared to the underproduction cases. Between the 2-hour and 4-hour *Baseline* strategies, there is no or small visual difference between the plots.

The right panels in Figure 5.10 display the most favorable outcomes for the *Proposed* strategy for both the 2-hour and 4-hour forecasting horizons: the two cases highlighted earlier in the heatmaps of 51% and 25% cost reductions. The primary objective of this strategy was to eliminate the red dots observed in the *Baseline* configuration in quadrants two and four, and to replace them with blue dots indicating "no trade". The results indicate that, when imbalance price predictions are incorporated, the model adjusts its decisions such that many of the previously non-profitable trades in quadrants two and four disappear.

However, this adjustment is also introducing an increase in non-trade decisions (blue dots) in quadrants one and three, where they replace both red trades (incorrectly predicted in the *Baseline* strategy) and green trades (correctly predicted in the *Baseline* strategy). This effect is particularly clear in quadrant one, where a large share of red trades is replaced by blue decisions. This arises because these trades were incorrectly predicted in terms of volume and, under the *Proposed* strategy, are consequently classified as belonging to quadrant two rather than quadrant one.

Comparing the *Proposed* strategy at the 2-hour and 4-hour horizons, some visible differences appear. The 2-hour forecasting model exhibits higher accuracy in quadrant two relative to the 4-hour forecasting model, as indicated by a larger proportion of blue (no-trade) dots and a lower frequency of red dots. Moreover, in quadrant three, the 4-hour model frequently misclassifies observations as "no trade" even though they should be executed as trades. The difference between the *Proposed* strategies at the different horizons is due to the imbalance price prediction, which complicates the correct estimation of the spread between the intraday market price and the imbalance price, at the longer horizon.

5.5 Robustness and stability analysis

In this section, the predictive models will be further analyzed in order to understand: how the regime changes on the imbalance market affects the price predictions; how the models performance could be increased by different initial settings; and how the aging of the models affect the results.

5.5.1 Analysis of the regime changes

Although all cost and volume analyses in this study were conducted for 2025, it was also of interest to examine how the predictive price models performed in 2024, i.e., prior to the major regime changes described in Section 2.2. This comparative analysis was carried out not only for 2- and 4-hour horizons, but for all prediction horizons from 1 to 24 hours, in order to more accurately capture differences between the years. The primary objective of this comparison was to align the model performance evaluation with previous studies, also made before 2025.

The performance of the price forecasting models across different prediction horizons is shown in Figure 5.11. In this figure, the SPOT prediction denotes the RMSE between the realized imbalance price and a benchmark forecast that uses the SPOT price as the predicted imbalance price, computed separately for 2024 and 2025. For each year, this SPOT-based RMSE was normalized to 100% and used as a reference for assessing relative model performance. For both 2024 and 2025, the RMSE of the Averaged imbalance price predictive model was computed for the 24 different prediction horizons and then expressed as a percentage of the corresponding SPOT-based RMSE for that year. Lower percentage values indicate higher model performance, with 0% corresponding to a perfect imbalance price forecast.

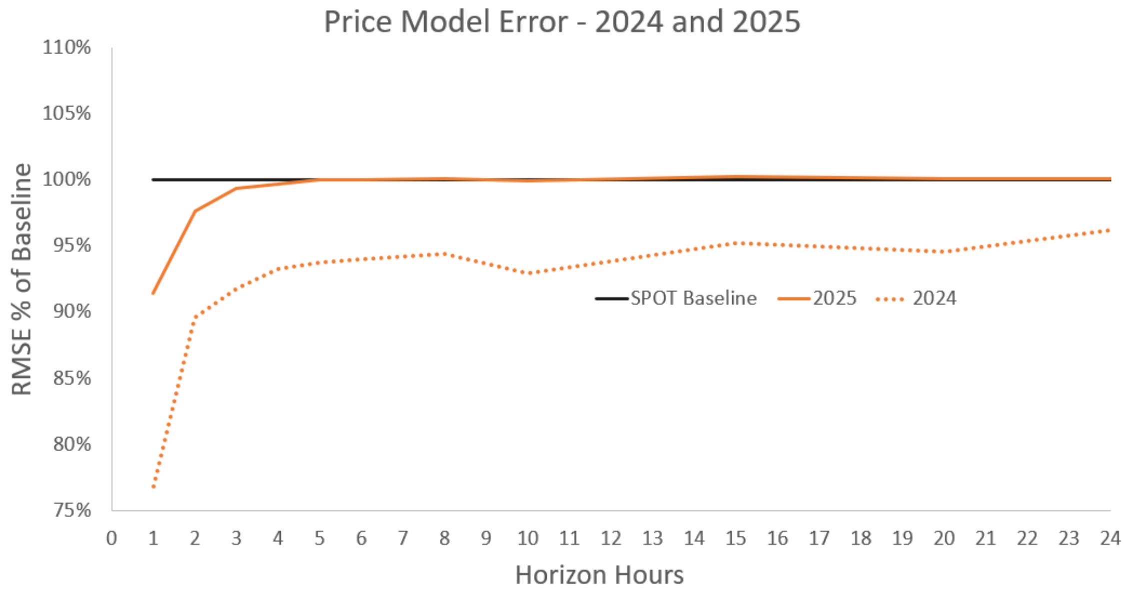

Figure 5.11: The imbalance price Averaged model performance during 2025 and 2024 compared with the SPOT imbalance price guess. The data is indexed according to the SPOT baseline for the years 2025 and 2024 separately.

Figure 5.11 clearly illustrates that the predictive performance of the price model deteriorated following the regime changes in the balancing market. The dotted line representing 2024 lies below the solid line representing 2025 across all 24 forecast horizons, indicating a higher predictive performance during 2024 compared to 2025.

For both 2024 and 2025, the models reaches a plateau beyond a prediction horizon of approximately 5 hours, suggesting that extending the horizon further does not highly affect forecast accuracy. For horizons shorter than 5 hours, predictive accuracy improves for both years, with the highest performance observed at the 1-hour horizon: approximately 90% of the SPOT benchmark for 2025 and 75% for 2024.

The primary explanation for the convergence of the imbalance price prediction accuracy over longer forecast horizons in both years lies in the fading informational content of lagged variables, e.g. past imbalance prices, as the horizon increases. For instance, the most recent observed imbalance price include high predictive information for the immediately following time step, but its contribution to forecast accuracy is lower when used to predict 10 hours ahead. Although both years exhibit a plateau in predictive performance, there is a significant difference in the level at which this plateau occurs: in 2025 the performance stabilizes at approximately 100%, whereas in 2024 it stabilizes in the range of 90–95%.

This indicates that, as the predictive relevance of lagged variables decreases with increasing horizon length, the non-lagged explanatory variables—specifically the day-ahead market prices and volume variables, which are not varying across forecast horizons—have some predictive power in 2024 but not in 2025. This difference between the years is most certainly explained by the higher price volatility observed in 2025, during which day-ahead production numbers are not any more controlling price spikes and imbalance activations.

5.5.2 Sensitivity analysis

A sensitivity analysis was also carried out in this study for the three different forecasting models - GBDT, SVR and MLP - and evaluated both for the price and volume predictions under the 2-hour forecast horizon. In the Table 5.3 below, the three hyperparameters iterated for every model are presented, as well as the best, worst and thesis combination with their respective RMSE value. All of the hyperparameters combinations are presented in the Method section 4.4.3.

In the GBDT and SVR model, 27 combinations of hyperparameters were tested, while in the MLP case only 18 combinations. The results indicate that the model combination used in the thesis all is between the best and worst combinations in the RMSE evaluations - often closer to the highest performing one - indicating that the theoretically informed hyperparameter setting had accuracy. Between the lowest performing combination (highest RMSE) and the highest performing combination (lowest RMSE), often the difference is not larger than the difference between the models themselves.

The only parameter setting that stands out, is the lowest performing combination of the MLP model. Instead of the number of hidden layers having the largest influence on the predictive outcome, the early stopping appeared to be of higher importance. Both for the MLP volume and price model, when no early stopping was applied, the model was overfitted and not performing well during the test year.

Due to time limits in this project, the best hyperparameter setup found in the sensitivity analysis was not tested in the intraday bidding strategies in order to see how this affected the cost reduction and imbalance volume changes. However, it is not certain that an increased predictive performance would yield a higher performing intraday trading scheme. Mostly since the previous evaluations show that models with an averagely poor performance, like the Naïve model, still can capture important price and volume events that overall give high cost reduction, although this value is not directly shown in the model's predictive performance.

Table 5.3: Sensitivity analysis for the models at the 2-hour forecasting horizon.

GBDT 2h					
Target	**Setup**	**Learning rate**	**Iterations**	**Max depth**	**RMSE**
Price	**Thesis**	0.05	500	6	**146.4**
	Best	0.03	300	3	**143.8**
	Worst	0.1	800	9	**152.3**
Volume	**Thesis**	0.05	500	6	**36.9**
	Best	0.03	300	6	**35.5**
	Worst	0.1	800	3	**39.8**

SVR 2h					
Target	**Setup**	**C**	**epsilon**	**gamma**	**RMSE**
Price	**Thesis**	10	0.1	scale	**144.4**
	Best	1	0.1	0.01	**144.1**
	Worst	1	1	0.1	**145.8**
Volume	**Thesis**	10	0.1	scale	**40.4**
	Best	10	1	0.01	**36.0**
	Worst	1	0.01	0.1	**42.0**

MLP 2h					
Target	**Setup**	**Hidden layers**	**alpha**	**Early stopping**	**RMSE**
Price	**Thesis**	(128,64,32)	0.001	True	**147.0**
	Best	(128,64,32)	0.01	True	**145.4**
	Worst	(64,32)	0.01	False	**166.8**
Volume	**Thesis**	(128,64,32)	0.001	True	**40.7**
	Best	(64,32)	0.0001	True	**38.3**
	Worst	(128,64,32)	0.001	False	**60.0**

5.5.3 Model aging

Model aging is a way to analyze the predictive deterioration of a model, caused by training data becoming less representative the longer the test period is, because of market or user behavior changes. The model aging test conducted in this study was evaluated with different model ages: 0, 3, 6 and 12 months, described deeper in section 4.4.4. The results from the test is visual in Table 5.4, showing the aging of the three main models (GBDT, SVR and MLP) both for volume prediction and price prediction at a horizon of 2-hours.

The aging test results show that the prediction models trained with more recent data are outperforming models trained on earlier data in predictive accuracy. The highest performance within every predictive model (marked in green in the table) is situated in the lower age span, while the lowest performance (marked in red in the table) is found in the higher age span. This indicates that the models are aging and that the models must be re-trained with more recent data regularly.

Table 5.4: The RMSE of the models with different aging averaged under a monthly resolution test period during the last six months of 2025, with model training lagged with 0, 3, 6 or 12 months.

Models	**Age: 0 months**	**Age: 3 months**	**Age: 6 months**	**Age: 12 months**
GBDT Price 2h	96.1	**94.8**	95.1	**97.2**
GBDT Volume 2h	35.1	**34.2**	34.7	**35.8**
SVR Price 2h	**95.3**	95.4	95.6	**95.6**
SVR Volume 2h	**34.5**	35.0	35.2	**36.8**
MLP Price 2h	**94.3**	95.9	95.4	**102.9**
MLP Volume 2h	**62.4**	70.4	**80.5**	77.1

To decide the frequency of updated training, the difference in performance - within one model and different ages - must be evaluated. For the GBDT and SVR models, the RMSE difference between the recently trained model and the aged model was small, only around 1%, indicating that re-training yearly is enough. That also defends the method used in this study, where the predictions models supporting the intraday trading, never had an age higher than 12 months. Also, for the GBDT model, the best performing one is not of 0 months age, but 3 months, indicating that other factors could influence predictive performance more than a few months younger model.

In the model aging test, the MLP model stood out. For the price predictive model, the RMSE performance between the lowest and highest model ages was 10%, while for the volume model, this difference was 23%. This indicated that re-training was important at a higher frequency than only at a yearly basis to keep the predictive accuracy. One possible explanation is that the MLP model only can capture the price spikes introduced in the regime changes if those are present also in the training data. This stands in contrast to the GBDT and SVR models, which are probably more likely to just include the lagged variables in the model without reflecting about frequency of previous price spiking.

5.6 Results compared to the previous studies

When comparing the results of this study with the previous work on imbalance price forecasting and intraday bidding strategies, both similarities and differences are found.

5.6.1 Imbalance price model comparison

For the imbalance price prediction, this study found the MLP model (based on a neural network) to not outperform the other tested models. In fact, the MLP had lower predictive performance - measured in RMSE and MAE - than the GBDT, SVR and the Averaged model over horizons of both 2 and 4 hours. In the previous studies, the neural network models almost always had the highest performance, contrary to what was presented here.

One possible explanation is that the regime changes - not tested in previous available studies - strongly affected the MLP predictive performance, while the other models incorporated information during training beneficial for predicting volatility, i.e. weighting the lagged imbalance prices higher. Another explanation could be a faster model aging of the MLP model - visible in the aging section above - causing the MLP model to mispredict imbalance prices in the end of the tested year to a higher degree than the other models.

However, a result presented in this study, corresponding with previous research, was the small or almost neglectable predictive performance difference between the models. Only a few percentages differed in the error values (excluding the Naïve model), which was also shown in previous studies and explained by the volatile nature of the time series, decreasing overall predictivity and reducing ability for certain models to stand out. Naturally, after the regime changes, this applied also to the models tested for imbalance price prediction in this study.

In section 5.5.1, the imbalance price prediction performance for the averaged model was presented for both the year 2024 and 2025, as well as for all horizon hours from 1 to 24 hours, instead of only the 2 and 4-hour horizon for the year 2025 used in the intraday trading. The results showed that the model performance prior to the regime changes (during 2024) was higher compared to the 2025 predictions. The 2024 model was outperforming the SPOT baseline prediction with a few percentages also at the day-ahead horizon. This was in line with what the Swedish study presented, where an ANN model clearly outperformed the SPOT guess at the 24-hour horizon during 2024 in SE2. The difference in price area evaluated in this study (SE4) and in the other study (SE2), introduces difficulties in further comparing the results.

5.6.2 Intraday trading comparison

A more risk-taking intraday strategy - as the *Proposed* strategy used in this study - gives cost benefits in the same way as those strategies presented in previous studies. Although previous studies used a strategy where trades are executed in order to increase the imbalance volume, somewhat similar results are found here (with a strategy that was not violating the BRP agreement).

For instance, one study found a 17% cost reduction when applying an imbalance price trading strategy, compared to the *Baseline* strategy. In this study, for the 2-hour forecasting, the cost reduction difference between the analyzed strategies are about 15 percentage units, i.e. in the same magnitude. However, the resulting numbers from the previous studies are at a day-ahead forecast horizon, while here, already at a 4-hour horizon the arbitrage opportunity vanishes. This further cements the effects that the regime changes have had on the imbalance price market.

This relates also to another conclusion from the studies: that the intraday trading profits increases closer to delivery hour. For instance, one study showed that forecasting imbalance price intraday instead of day-ahead, could increase intraday trading incomes by 2-5%. This cost improvement is visual also in the results presented here, but in higher absolute percentage numbers.

6
Conclusion

This thesis incorporated an intraday trading strategy for wind power producers in Sweden, that instead of only trading the imbalance volume, incorporated an imbalance price prediction. By using the price spread between the intraday price and the forecast imbalance price, trading opportunities appeared that helped reduce the high costs for imbalance settlement. By intentionally taking imbalance volume to the balance settlement through the *Proposed* strategy, and thereby increase risk, a higher cost reduction was achieved than for the branch standard risk-averse *Baseline* strategy.

During 2025 in price area SE4, the highest performing *Proposed* strategy could reduce imbalance settlement costs by 51% compared to no trading, while the corresponding *Baseline* strategy only managed 38%. This indicated that the imbalance price prediction model was somewhat accurate and could capture the price spread in order to profit on those opportunities. It was also concluded that the individual price model with lowest error metric - and thus best predictive accuracy - necessarily did not yield the highest cost reduction. For instance, the Averaged model had highest predictive accuracy at the 2-hour horizon, while the Naïve model - with lowest predictive performance at the same horizon - gave the best cost reduction results. That indicated that capturing price signals and directions, was of higher importance than the best guess over the year.

Furthermore, the results from the different prediction horizons evaluated - two hours and four hours ahead of delivery hour - indicated that the imbalance price prediction was heavily degraded at longer prediction horizons. When the 2-hour horizon scenario was evaluated, the best *Proposed* strategy rendered a larger cost reduction compared to all the *Baseline* strategies. The opposite occurred when the forecast horizon was increased in the 4-hour scenario: the best *Baseline* strategy outperformed all the *Proposed* strategies combinations. The risk taken in the *Proposed* strategy did therefore only render increased incomes at the short 2-hour forecast horizon. For the *Baseline* strategy, the imbalance cost reductions were kept at similar levels at both the time horizons, indicating that the volume model's predictive performance was not as degraded as the price predictive model at longer horizons.

The study also investigated how the ability to predict the imbalance price was affected by the regime changes before and during 2025. When the performance was compared to a simple SPOT price guess under both 2024 and 2025, it was visible how the predictive accuracy was degraded. For instance, during 2024 it was possible to capture imbalance price activations up to 24 hours before delivery, while for 2025, activations were only predictable maximum 5 hours ahead of delivery. Although the degraded predictive accuracy under 2025, intraday trading informed by imbalance price prediction was still profitable at 2-hour horizon, showing the opportunities of the strategy.

Lastly, it was discussed how the *Proposed* strategy could violate the BRP agreement for the wind power producers. The agreement does not explicitly prohibit such a trading strategy - particularly given that its objective is to avoid increasing imbalances in the wind power sector - but certain forecasting errors or mispredicted events may still occur that could increase the resulting imbalance volumes. However, that could be the case also within the *Baseline* strategy. Also, the imbalance volume results do not indicate an increase in settled imbalance volumes under any of the strategies, although the *Baseline* strategy provides the greatest mitigation of wind power imbalance volumes.

6.1 Future studies

One addition to this study - in order to increase the *Proposed* strategy's performance - would include a confidence interval to the imbalance price prediction. This could then be incorporated also to the intraday strategy. For instance, only if the prediction was of a certain accuracy, the intraday trading would be executed. That could help exclude trades that are predicted wrongly, or decrease the volume traded under specific prediction uncertainty.

Another possible study improvement would analyze how the intraday market responds to increased or decreased intraday trading volumes. In this thesis, no intraday price adjustments were included when calculating the results, since available intraday data was limited and did not include all bid sizes and respective prices. However, if the model could incorporate how the intraday market prices responds to both high and low volume trades, the results would more accurately reflect real deployment of the *Proposed* intraday strategy.

Bibliography

[1] Svenska Kraftnät. (2024, 17 June). *Balance Settlement*. Received 2026, 23 January. Available: https://www.svk.se/en/stakeholders-portal/electricity-market/balance-settlement/

[2] Svenska Kraftnät. (2025, 24 March). *Information on different ancillary services*. Received 2026, 23 January. Available: https://www.svk.se/en/stakeholders-portal/electricity-market/provision-of-ancillary-services/information-on-different-ancillary-services/

[3] Next Kraftwerke. (n.d.). *What is mFRR (manual Frequency Restoration Reserve / R3)?*. Received 2026, 23 January. Available: https://www.next-kraftwerke.com/knowledge/mfrr

[4] Svenska Kraftnät. (2025, 4 March). *Automatiserad nordisk energiaktiveringsmarknad för mFRR*. Received 2026, 23 January. Available: https://www.svk.se/utveckling-av-kraftsystemet/systemansvar–elmarknad/ny-nordisk-balanseringsmodell-nbm/automatiserad-nordisk-energiaktiveringsmarknad-for-mfrr/

[5] Svenska Kraftnät. (2025, 21 November). *Om obalanskostnader och Svenska kraftnäts ansvar*. Received 2026, 23 January. Available: https://www.svk.se/om-kraftsystemet/om-systemansvaret/balansering-av-kraftsystemet/om-obalanskostnader-och-svenska-kraftnats-ansvar/

[6] Varberg Energi. (2025, 7 April). *Extrema balanskostnader hotar sol och vind*. Received 2026, 9 June. Available: https://www.varbergenergi.se/nyheter/extrema-balanskostnader-hotar-sol-och-vind

[7] Svenska Kraftnät. (2025, 26 November). *Intradagsmarknaden – justering av dagen före-handel*. Received 2026, 10 February. Available: https://www.svk.se/om-kraftsystemet/om-elmarknaden/intradagsmarknaden–justering-av-dagen-fore-handel/

[8] SCB. (2026, 15 January). *Elförsörjning 2025*. Received 2026, 28 January. Available: https://www.scb.se/hitta-statistik/statistik-efter-amne/energi/tillforsel-och-anvandning-av-energi/manatlig-elstatistik-och-

byten-av-elleverantor/pong/tabell-och-diagram/elforsorjning/

[9] Svenska Kraftnät. (2025, 14 February). *Vårt uppdrag*. Received 2026, 28 January. Available: https://www.svk.se/om-oss/verksamhet/

[10] Svenska Kraftnät. (2025, 19 March). *Roller i kraftsystemet*. Received 2026, 28 January. Available: https://www.svk.se/om-kraftsystemet/oversikt-av-kraftsystemet/roller-i-kraftsystemet/

[11] Svenska Kraftnät. (2024, 26 April). *Operations and Electricity Markets*. Received 2026, 28 January. Available: https://www.svk.se/en/national-grid/operations-and-electricity-markets/

[12] Loiacono, L., & Rizzo, L., & Stagnaro, C. (2025). Impact of bidding zone reconfigurations on electricity prices: Evidence from Sweden. *Energy Economics*, 141, 1-13. Available: https://doi.org/10.1016/j.eneco.2024.108106

[13] Svenska Kraftnät. (2025, 16 January). *Elområden*. Received 2026, 5 February. Available: https://www.svk.se/om-kraftsystemet/om-elmarknaden/elomraden/

[14] Svenska Kraftnät. (2025, 26 November). *Dagen före-marknaden – fysisk handel med el*. Received 2026, 4 February. Available: https://www.svk.se/om-kraftsystemet/om-elmarknaden/dagen-fore-marknaden–fysisk-handel-med-el/

[15] Nord Pool. (2026, 4 February). *Aggregated Bidding Curves*. Received 2026, 4 February. Available: https://data.nordpoolgroup.com/auction/day-ahead/aggregated-bid-curves?deliveryDate=latest&bidCurveArea=SE

[16] Klæboe, G., & Braathen, J., & Eriksrud, A.L. et al. (2022). Day-ahead market bidding taking the balancing power market into account. *TOP*, 30, 683-703. Available: https://doi.org/10.1007/s11750-022-00645-1

[17] HU, X., & Jaraitė, J., & Kažukauskas, A. (2021). The effects of wind power on electricity markets: A case study of the Swedish intraday market. *ENERGY ECONOMICS*, 96, 105-159. Available: https://doi.org/10.1016/j.eneco.2021.105159

[18] entso-e. (n. d.). *Intraday Auctions on SIDC*. Received 2026, 10 February. Available: https://www.entsoe.eu/network_codes/cacm/implementation/ida/

[19] NEMO Committee. (2022, 6 October). *SIDC IDA pre-launch event presentation*. Received 2026, 10 February. Available: https://www.nemo-committee.eu/assets/files/sidc-ida-pre-launch-event-presentation.pdf

[20] Energiforsk. (2021). *INTRADAGMARKNADEN RAPPORT 2021:797.* Received 2026, 10 February. Available: https://energiforsk.se/media/30200/intradagmarknaden-energiforskrapport-2021-797.pdf

[21] Svenska Kraftnät. (2025, 1 September). *Balansmarknaden.* Received 2026, 13 February. Available: https://www.svk.se/om-kraftsystemet/om-elmarknaden/balansmarknaden/

[22] Svenska Kraftnät. (2025, 1 September). *Vad är nedreglering och varför behövs det?.* Received 2026, 13 February. Available: https://www.svk.se/om-kraftsystemet/om-elmarknaden/balansmarknaden/vad-ar-nedreglering-och-varfor-behovs-det/

[23] Svenska Kraftnät. (2025, 28 October). *Så kan du som privatperson bidra med stödtjänster.* Received 2026, 18 February. Available: https://www.svk.se/om-kraftsystemet/om-systemansvaret

[24] Svenska Kraftnät. (2025, 22 October). *Manuell frekvensåterställningsreserv (mFRR).* Received 2026, 18 February. Available: https://www.svk.se/aktorsportalen/bidra-med-reserver/om-olika-reserver/mfrr/

[25] eSett. (2025, 29 October). *Nordic Imbalance Settlement Handbook* . Available: https://www.esett.com/app/uploads/2025/10/NBS-Handbook-v5.2.pdf

[26] Svenska Kraftnät. (2023, 22 September). *Välkommen till Träffa Balansmarknad* . Received 2026, 18 February. Available: https://www.svk.se/4a7721/contentassets/a91afa18da584cd793d29a88626bdaeb/introduktion-till-mfrr–traffa-balansmarknad-mote-17-2023-09-22-.pdf

[27] Fingrid. (n. d.). *Balancing energy and balancing capacity markets (mFRR).* Received 2026, 18 February. Available: https://www.fingrid.fi/en/electricity-market/reserves_and_balancing/balancing-energy-and-balancing-capacity-markets/#activating-the-mfrr-energy-bid-

[28] Svenska Kraftnät. (2024, 15 August). *Införande av enpris.* Received 2026, 18 February. Available: https://www.svk.se/utveckling-av-kraftsystemet/systemansvar–elmarknad/ny-nordisk-balanseringsmodell-nbm/inforande-av-enpris/

[29] Statnett. (2023, 29 November). *Wind and Solar in the Nordic Reserve Markets.* Received 2026, 18 February. Available: https://www.statnett.no/globalassets/for-aktorer-i-kraftsystemet/.pdf

[30] Achour. N. (2025, 26 February). Konsekvenserna av mFRR EAM och flödesbaserad kapacitetsberäkning. *Sigholm*. Received 2026, 18 February. Available: https://www.sigholm.se/artiklar/mfrr-eam

[31] Svenska Kraftnät. (2025, 14 January). *Välkommen till Träffa Balansmarknad* . Received 2026, 18 February. Available: https://www.svk.se/48e9b0/siteassets/om-oss/mot-svenska-kraftnat/traffa-balansmarknaden/introduktion-till-mfrr-eam-2025-01-14_for-publicering.pdf

[32] Nordenergi. (2026, 10 June). *Impact of Recent Market Changes on the Nordic Balancing Market* . Received 2026, 24 February. Available: https://www.nordenergi.eu/wp-content/uploads/2025/06/mFRR-EAM-and-consequences-for-the-balancing-market_2025.pdf

[33] Nord Pool. *Balance market (mFRR)*. Received 2026, 20 January. Available: https://data.nordpoolgroup.com/power-system/manual-frequency-restoration-reserve?deliveryDate=latest&deliveryAreas=SE1,SE2,SE3,SE4&viewMode=All

[34] Svenska Kraftnät. (2024, 29 October). *Flödesbaserad kapacitetsberäkningsmetod*. Received 2026, 25 February. Available: https://www.svk.se/utveckling-av-kraftsystemet/systemansvar–elmarknad/flodesbaserad-kapacitetsberakningsmetod/

[35] Stermedia. (n.d.) *Gradient Boosting Decision Trees*. Received 2026, 9 March. Available: https://stermedia.ai/gradient-boosting-decision-trees/

[36] Medium. (2023, 21 December). *Gradient Boosted Decision Tree — Clearly Explained*. Received 2026, 9 March. Available: https://medium.com/@ruchi.awasthi63/gradient-boosted-decision-tree-clearly-explained-bd1d8c7d9923

[37] Persio, L.D., & Fraccarolo, N. (2023). Energy Consumption Forecasts by Gradient Boosting Regression Trees. *Mathematics*, 11(5), 1068. Available: https://doi.org/10.3390/math11051068

[38] Friedman, J.H., (2000). Greedy Function Approximation: A Gradient Boosting Machine. *The Annals of Statistics*, 29(5). Available: https://doi.org/10.1214/aos/1013203451

[39] Medium. (2025, 15 January). *Introduction to Support Vector Machines and Suppor Vector Regression*. Received 2026, 10 March. Available: https://medium.com/h7w/introduction-to-support-vector-machines-svm-and-support-vector-regression-svr-b7cd76204def

[40] MatLab. (n.d.). *Understanding Support Vector Machine Regression.* Received 2026, 10 March. Available: https://www.mathworks.com/help/stats/understanding-support-vector-machine-regression.html

[41] SciKit-Learn. (n.d.). *rbf_kernel.* Received 2026, 10 March. Available: https://scikit-learn.org/stable/modules/generated/sklearn.metrics.pairwise.rbf_kernel.html

[42] Awad, M., & Khanna, R. (2015) *Efficient Learning - Theories, Concepts, and Applications for Engineers and System Designers.* Apress Berkeley, CA. Available: https://doi.org/10.1007/978-1-4302-5990-9

[43] Science Direct. (2018). *Support Vector Regresion.* AI Summary. Received 2026, 10 March. Available: https://www.sciencedirect.com/topics/computer-science/support-vector-regression

[44] Baladram, S. (2024). *Multilayer Perceptron, Explained: A Visual Guide with Mini 2D Dataset.* Medium. Received 2026, 12 March. Available: https://medium.com/data-science/multilayer-perceptron-explained-a-visual-guide-with-mini-2d-dataset-0ae8100c5d1c

[45] SciKit-Learn. (n.d.). *1.17. Neural network models (supervised).* Received 2026, 12 March. Available: https://scikit-learn.org/stable/modules/neural-networks-supervised.html

[46] Wang, X., & Hyndman, R., & Li, F., & Kang, Y. (2023). Forecast combinations: An over 50-year review. *International Journal of Forecasting*, 39(4), 1518-1547. Available: https://doi.org/10.1016/j.ijforecast.2022.11.005

[47] Kosourova, E. (2025). *RMSE Explained: A Guide to Regression Prediction Accuracy.* Datacamp. Received 2026, 30 March. Available: https://www.datacamp.com/tutorial/rmse

[48] Waples, J. (2025). *Mean Absolute Error Explained: Measuring Model Accuracy.* Datacamp. Received 2026, 30 March. Available: https://www.datacamp.com/tutorial/mean-absolute-error

[49] Abdelmotteleb, I., & Esmat, A., & Tijm, S. , & Gibescu, M. (2023). Deep Learning-Based Imbalance Market Price Range Predictions in the Day-Ahead Horizon [Conference Paper]. *2023 IEEE Belgrade PowerTech* Available: https://doi.org/10.1109/PowerTech55446.2023.10202684

[50] Terpstra, N.R. (2020). *Day-ahead and imbalance price forecasting on the Dutch Electricity Market: a comparison between time series and artificial neural networks models.* Eindhoven University of Technology. Available:

https://research.tue.nl/en/studentTheses/day-ahead-and-imbalance-price-forecasting-on-the-dutch-electricit/

[51] Svedlindh, T., & Yngveson, K. (2025). *Price Formation and Forecasting Models in the Electricity Market*. Uppsala Universitet. Available: https://www.diva-portal.org/smash/get/diva2:1983102/FULLTEXT01.pdf

[52] Ganesh, V.N., & Bunn, D. (2024). Forecasting Imbalance Price Densities With Statistical Methods and Neural Networks. *IEEE TRANSACTIONS ON ENERGY MARKETS POLICY AND REGULATION*, 2(1), 30-39. https://doi.org/10.1109/TEMPR.2023.3293693

[53] Svenska Kraftnät. (2025, 26 November). *Intradagsmarknaden – justering av dagen före-handel*. Received 2026, 22 April. Available: https://www.svk.se/om-kraftsystemet/om-elmarknaden/intradagsmarknaden–justering-av-dagen-fore-handel/

[54] Svenska Kraftnät. (n.d.). *Avtal om Balansansvar för el*. Received 2026, 21 April. Available: https://www.svk.se/4aa60d/siteassets/aktorsportalen/bsp-och-brp/kommande-avtal-bsp-och-brp/kommande-avtal-brp/1-brp-huvuddel-_-5938-2.pdf

[55] Haglund, K., & Lööf, A. (2024). *An Analysis of the Swedish Intraday Market*. Lund Tekniska Universitet. Available: https://lup.lub.lu.se/luur/download?func=downloadFile&recordOId=9169608&fileOId=9169610

[56] Koch, C. (2022). Intraday imbalance optimization: incentives and impact of strategic intraday bidding behavior. *Energy Systems*, 2, 409-435. Available: https://doi.org/10.1007/s12667-021-00445-9

[57] Browell, J., & Gilbert, C. (2022). Predicting Electricity Imbalance Prices and Volumes: Capabilities and Opportunities. *Energies*, 15(10), 3645. Available: https://doi.org/10.3390/en15103645

[58] Chaves-Àvila, J.P., & Haakvort, R., & Ramos, A. (2012). Short-term strategies for Dutch wind power producers to reduce imbalance costs. *Energy Policy*, 52, 573-582. Available: https://doi.org/10.1016/j.enpol.2012.10.011

[59] Svenska Kraftnät. (2025, 23 September). *Webbinarium om självreglering för BRP:er*. Received 2026, 22 April. Available: https://www.svk.se/om-oss/mot-svenska-kraftnat/evenemang/2025/webbinarium-om-sjalvreglering-for-brper/

[60] Bruneel, R., & Schuurmans, M., & Patrinos, P. (2025). Optimal intraday power trading for single-price balancing markets: An adaptive risk-averse strategy using mixture models. *Applied Energy*, 389, 125754. Available:

https://doi.org/10.1016/j.apenergy.2025.125754

[61] entsoe. (n.d.). *Transparency Platform*. Available: https://transparency.entsoe.eu/

[62] Narajewski, M. (2022). Probabilistic Forecasting of German Electricity Imbalance Prices. *energies*, 15(14). Available: https://doi.org/10.3390/en15144976

[63] SciKit-Learn. (n.d.). *Compare the effect of different scalers on data with outliers*. Received 2026, 23 April. Available: https://scikit-learn.org/stable/auto_examples/preprocessing/plot_all_scaling.html

[64] SciKit-Learn. (n.d.). *Supervised learning*. Received 2026, 27 April. Available: https://scikit-learn.org/stable/supervised_learning.html

[65] SciKit-Learn. (n.d.). *HistGradientBoostingRegressor*. Received 2026, 27 April. Available: https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.HistGradientBoostingRegressor.html

[66] SciKit-Learn. (n.d.). *SVR*. Received 2026, 27 April. Available: https://scikit-learn.org/stable/modules/generated/sklearn.svm.SVR.html

[67] SciKit-Learn. (n.d.). *Scaling the regularization parameter for SVCs*. Received 2026, 27 April. Available: https://scikit-learn.org/stable/auto_examples/svm/plot_svm_scale_c.html#sphx-glr-auto-examples-svm-plot-svm-scale-c-py

[68] SciKit-Learn. (n.d.). *MLPRegressor*. Received 2026, 27 April. Available: https://scikit-learn.org/stable/modules/generated/sklearn.neural_network.MLPRegressor.html

[69] Mucci, T. (n.d.). *What is data leakage in machine learning?*. IBM. Received 2026, 1 April. Available: https://www.ibm.com/think/topics/data-leakage-machine-learning

[70] Nord Pool. *Hourly statistics*. Received 2026, 9 April. Available: https://data.nordpoolgroup.com/intraday/intraday-hourly-statistics?deliveryDate=latest&deliveryArea=SE4

DEPARTMENT OF ENVIRONMENTAL AND ENERGY SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY