MS-E2177 Seminar on Case Studies in Operations Research, Spring 2026

Reconstruction of Finnish reserve market merit order curves

Final Report

Team:

Veikko Hokkanen (Project Manager)
Mikael Kari
Joel Niemi
Timo Norrkniivilä
Tuomas Rajala

Client:

Hytrade

May 22, 2026

Contents

1 Background

1.1 Finnish electricity market

The operation of a power system is based on the constant equilibrium between generation and consumption of electricity. In practice, supply and demand must be equal at every moment. In Finland and the broader Nordic synchronous system – comprising Sweden, Norway, and Denmark – this balance is reflected in maintaining the system frequency at 50 Hz (Sihvonen, 2025). The interconnection of national systems into a shared synchronous area is designed to improve market performance while strengthening security of supply (Nord Pool, 2026).

The Nordic electricity market structure consists of multiple distinct trading platforms. Transactions take place in financial markets, the day-ahead market, intraday markets, reserve capacity markets, and balancing markets (Fingrid Oyj, 2026b). The segmentation of these marketplaces is illustrated in Figure 1.

Figure 1: Figure 1. Finnish electricity market segments. Modified from Fingrid Oyj (2026b).

On financial electricity markets, derivative instruments such as futures and options are traded in connection with electricity prices. Their primary purpose is to manage risk and reduce exposure to price fluctuations. These contracts may be arranged either bilaterally or through organized exchanges. One example of a bilateral contract is a power purchase agreement (PPA), where the electricity price and delivery volume are agreed upon in advance for a specified period. Derivatives enable electricity producers to hedge against declining electricity prices, while purchasers can protect themselves from rising prices, especially in situations where electricity is sold onward at a fixed rate (Fingrid Oyj, 2026b).

The day-ahead market is the main segment of the wholesale electricity market. Trading takes place daily for every 15-minute interval of the following day. Production and consumption bids must be submitted by 13:00 Eastern European Time (EET). Information regarding available transmission capacity is provided to market participants at 10:30 EET. After the auction closes, prices are published and determined according to the balance between supply and demand (Fingrid Oyj, 2026b). Supplier bids consisting of electricity quantities and prices are arranged according to the merit order principle, which defines the order in which generation units are dispatched based on their variable production costs. All accepted suppliers, including those offering electricity at lower prices, receive the market equilibrium price (Lundgren, 2012; Nord Pool, 2025).

The intraday market is used to balance electricity production and consumption closer to the actual delivery period. Various uncertainties, including weather conditions, equipment malfunctions, and changes in consumption patterns, may influence electricity demand. Likewise, production can be disrupted by transmission line outages or technical limitations at power plants. Trading in the intraday market is continuous and begins after the closure of the day-ahead market. In Finland, intraday trading remains open until the beginning of the delivery hour (Fingrid Oyj, 2026b). Around 80 % of Finland's electricity demand is traded through the day-ahead and intraday markets (Lieskoski et al., 2024).

Power system balancing and frequency control are maintained through reserve and balancing markets, which make it possible to correct real-time imbalances between generation and consumption. Transmission system operators (TSOs) are primarily responsible for maintaining system balance and ensuring that sufficient flexible capacity is available (Lieskoski et al., 2024). Reserve capacity may be offered by electricity generation units, consumption resources, storage systems, or hybrid solutions that satisfy the required technical criteria and market regulations (Fingrid Oyj, 2025b). Market bids must be submitted no later than 45 minutes before the start of the delivery period (Fingrid Oyj, 2026b).

In the Nordic power system, frequency stability relies primarily on three categories of reserves: Frequency Containment Reserve (FCR), Frequency Restoration Reserve (FRR), and Fast Frequency Reserve (FFR). FCR is separated into products for normal operation and for disturbance situations, the latter including both upward and downward components (FCR-D up and FCR-D down). FRR is subdivided into automatic (aFRR) and manual (mFRR) forms (Lieskoski et al., 2024). These reserve products differ in terms of response time, activation mechanism, and direction of control, allowing them to address system needs effectively under varying conditions (Fingrid Oyj, 2026a). This project focuses particularly on the mFRR market.

The role of mFRR is to return system frequency toward the nominal 50 Hz level. In addition, it can support congestion management to safeguard operational reliability (Fingrid Oyj, 2026c). Over the long term, the procured volume of mFRR capacity is determined by projected imbalance magnitudes, the availability of backup generation resources, and the expected need for voluntary balancing energy bids to address both common and exceptional imbalance situations (Fingrid Oyj, 2025a).

Activation of mFRR occurs through Fingrid's control signals, either manually or automatically (Fingrid Oyj, 2025b). The implementation of the Nordic automated mFRR Energy Activation Market (mFRR EAM) has significantly modified the operation of the mFRR market by introducing automated bid selection and activation as well as a 15-minute market time unit. The activation time requirement for mFRR is 12.5 minutes, which is slower than for the other reserve products discussed earlier, indicating that mFRR is not intended for the fastest balancing actions. Participation in the mFRR market requires a minimum bid size of 1 MW, and bids must be submitted in increments of 1 MW (Fingrid Oyj, 2026c).

To participate as an mFRR reserve provider, an agreement with Fingrid must be established. The reserve product must also comply with all technical specifications defined for the product. Furthermore, the provider is required to sign an agreement with the balance settlement company eSett, which is responsible for the settlement of mFRR energy (Fingrid Oyj, 2025c).

Trading in the mFRR market takes place through both capacity and energy markets. Capacity bids are submitted hourly and cleared in the day-ahead stage, whereas energy bids are submitted for each 15-minute market time unit. Capacity market bids for the following day must be submitted by 08:30, while energy market bids must be submitted 45 minutes before the operational period. Within the energy market, activations can be either scheduled or direct. Scheduled activation means that an activation signal is predetermined for the beginning of a specific market time unit, while direct activation allows activation at any point during the unit. In addition to the energy and capacity markets, balancing capacity may also be procured through balancing capacity agreements established via a separate tendering process. Participants are obligated to offer the entire contracted capacity for the full agreement period (Fingrid Oyj, 2025d).

Several different bid structures are available in the mFRR energy market. An alternative bid is submitted together with other bids, but only one bid from the group may be activated during a single market time unit. In a multipart bid structure, lower-priced bids belonging to the same multipart group must be activated before higher-priced bids. A conditionally linked bid may only be activated if the linked bids were activated during the two preceding market time units. The availability of a technically linked bid depends on whether the connected bid was directly activated during the previous market time unit (Fingrid Oyj, 2025d).

An mFRR capacity bid must contain information regarding capacity, price, hour, bidding zone, and transmission area. Similarly, an energy bid must specify power, price, location, reserve-related information indicating whether the balancing offer concerns reserve power, the activation type, bid divisibility, and the minimum volume that can be activated (Fingrid Oyj, 2025d).

In the mFRR capacity market, the clearing price is determined according to the highest accepted bid in both the upward and downward regulation directions, meaning that all accepted participants receive the same market price. In the energy market, pricing is likewise based on activated bids. For upward regulation, the price corresponds to the highest activated bid in areas without transmission bottlenecks, while for downward regulation the price is determined by the lowest accepted bid. As in the capacity market, all activated bids are compensated at the common market clearing price. In addition, providers are required to offer the same volume in the energy market as committed in the capacity market unless prevented by force majeure circumstances. Failure to comply results in penalties. This requirement applies regardless of whether the bids are scheduled, directly activated, or linked (Fingrid Oyj, 2025d).

Balancing energy bids that have been accepted as capacity bids in the mFRR energy market must remain available for activation during at least the first three 15-minute market time units of each hour. These bids must be directly activatable. If a capacity bid has also been accepted for the subsequent hour, then a directly activatable bid must additionally be available for the fourth market time unit. Resource aggregation is permitted. In principle, aggregated resources must be located within the same transmission region, namely southern, central, or northern Finland. Aggregation across different regions is allowed only in cases where the minimum bid requirement could not otherwise be satisfied (Fingrid Oyj, 2025d).

In the markets, the TSO acts as a counterparty to market participants who bid up-regulation or down-regulation capacity and/or energy. Fingrid compensates participants whose bids are accepted, providing the opportunity for an additional revenue stream for companies in the energy sector.

Compared with the day-ahead (spot) market, reserve markets involve substantially lower trading volumes. As a result, market participants possessing high-capacity assets, such as wind farms, electric boilers, or data centers, may significantly influence market price formation. At the same time, generation technologies with near-zero marginal costs are becoming increasingly common. According to the literature, this development shifts the supply curve toward the right and generally reduces average market prices (Pavlík et al., 2025). Consequently, optimal operational decision-making requires accounting for the price impacts associated with market participation. At present, reserve market prices are largely determined using merit order curves. The effect of various factors to the merit order curve is uncertain, which creates challenges for optimization of the bids. In general, the merit order principle refers to the ranking of generation units according to increasing marginal costs in order to determine the sequence in which generation sources are dispatched to satisfy electricity demand (Pavlík et al., 2025).

1.2 Research questions

The objective of this project is to study the bid curves of the mFRR up and down capacity markets. The shape and size of the curves are assumed to be dynamic, varying with the time of day, seasonally, and over time as additional renewable generation and flexible assets are integrated into the grid. The main task of the project is to reconstruct bid curves for future time steps based on historical data. The initial research questions were stated as follows:

1. What is the shape of the whole bid curve in different time periods?
2. What is the clearing price in different time periods?
3. Is there a static component to these bid curves?
4. How do underlying variables such as day-ahead price, weather, seasonality affect these curves?
5. How have the bid curves changed over time as the electricity system develops?

In the course of the project, an extra research objective was added:

6. How the bid curve would change if an extra offer is injected into the market?

This aims to answer the questions of how sensitive the bid curve is and how additional capacity affects the bid curve. In this project, interactions related to different offer zones, transmission areas, and the mFRR energy markets are not modelled separately in the context of the mFRR capacity market.

2 Literature review

This section gives an overview of the relevant literature related to reconstruction of merit order curves in capacity markets. The studies presented below use a wide array of different methodologies to solve problems within the volatile energy industry whether that be predicting spot prices or mFRR trading volumes. The goal of this literature review is to uncover the use cases of different methodologies and choose a relevant method in support of our project.

A relevant study in the field of mFRR balancing volume prediction is a Master's Thesis from the University of Uppsala, where the use of machine learning was researched in predicting mFRR activation volumes across Sweden's bidding zones (Azarang and Edling, 2025). The thesis had inherently the same problem context as our project, as the energy market's in the Nordics are experiencing increasing complexity from the adoption of renewable energy sources, such as wind, hydropower and solar (Satymov et al., 2025). The machine learning algorithm applied to the dataset was Long Short-Term Memory (LSTM), which is particularly well suited for sequential training data and is robust in capturing complex temporal dependencies. The suitability of this method for our project is amplified by the fact that the data used in this thesis is gathered from the same Transparency Platform (ENTSO-E, 2015), which is the central repository for electricity market data in Europe. The relevant limitations of this thesis were the focus solely on downward activations and mFRR activation volumes, rather than the whole bid curve. The prediction algorithm also exclusively focused on one-hour-ahead markets, making it inherently a more simplified model. The results show variations in model performance across bidding zones, likely attributable to differences in hydropower concentration — the dominant source of balancing energy in Sweden — as well as varying shares of zero-activation hours.

Våle et al. (2025) researched the uses of machine learning in the case of mFRR activation price prediction in Norway's bidding zones by comparing two ML methods, XGBoost and EBM to see if a less "black-box" model could result in sufficient predictions while remaining explainable and traceable. The study also focused strongly on price drivers to see what relevant parameters are related to the construction of the activation price. The study identified day-ahead spot price, hydropower production, and heating demand as the most influential predictors across all bidding zones, listed in order of significance. The limitations recognized by the study relate to these predictors, as the model developed performed poorly when the mFRR activation price deviated strongly from the day-ahead spot price. While this study serves as a useful foundation for the feature engineering process, the multi-zone bidding structure of Norway may not translate directly to Finland's single-zone market.

Mitridati and Pinson (2018) attempt to reconstruct European day-ahead electricity supply curves, using solely data of historical clearing prices and total traded electricity values. The authors propose a hidden Markov model to model the clearing price mechanism, and a Bayesian modeling framework to predict the supply curve and its uncertainty, in even-sized generation blocks. Using synthetically created bid-level data, the authors note that the model is able to predict the shape of the price curve with reasonable accuracy but struggles in reconstructing the lower and higher ends, since the clearing price rarely takes on values in those regions.

Ziel and Steinert (2016) implement a vector autoregression model for day-ahead electricity markets in Germany to construct the true electricity price using the constructed supply and demand curves formed by market participants. The process of vector autoregression in this paper consists of three steps, where first the space of possible bids is deconstructed into 16 price classes for both curves. Then each price class bid volume is forecasted using a stochastic model after which the whole supply/demand curves are reassembled using the individually forecasted bid volumes. The stochastic model is a high-dimensional sparse VAR estimated via LASSO, where for each of the 24 auction hours a separate system of equations is estimated, with each bid volume process regressed on lagged values of all other price class volumes. The most notable features used by the model are the lagged bid volume processes of all 32 price classes, which form the core autoregressive structure, alongside lagged market clearing prices and volumes from previous auctions. Forward-looking planned generation data for conventional thermal, wind, and solar power is also incorporated, capturing the effect of renewable energy sources on the merit-order curve. Seasonality is handled both through the deep autoregressive lag structure, reaching back up to 36 days to capture weekly periodic patterns, and through explicit weekday dummy variables accounting for the systematically different bidding behavior observed across days of the week. The open challenges of this method are the necessary approximations done by the discretization of the merit-curve into price classes.

To eliminate the limitations arising from the coarse discretization of the X-Model, Koechlin et al. (2025) introduce a novel framework for forecasting day-ahead prices in the Italian spot market by applying functional principal component analysis (FPCA) to historical merit-order curves. By treating supply and demand curves as continuous functional objects and projecting them onto a compact set of eigenfunctions, the approach eliminates discretization errors entirely and preserves the true functional form of the curves. The resulting FPCA scores for supply and demand are concatenated into a joint vector representation, which is then modeled using a sparse VARX estimated with LASSO and BIC-selected penalty. Four model variants are compared with respect to how they handle cross-hour and cross-score dependencies, with the best performing model being the semi-full VARX achieving a mean absolute error of 8.28 €/MWh. The most important exogenous features are day-ahead forecasts of national load, wind and solar generation, and transfer capacities with neighboring market zones, alongside day-type dummy variables. The limitations of this model are the relatively shallow lag structure not allowing for a long-term dependencies to be modelled and the choice of using FPC fundamentally assumes that the form of the curves are somewhat stationary.

As some highly complex machine learning models often exhibit a "black-box" nature as discussed above by Våle et al. (2025), the interpretability and reproducability of the model is increasingly important. There is therefore a need for interpretability models that aim to decompose the prediction into the sum of individual contributions of the different features in the model. In the context of electricity markets, this interpretability of a XGBoost model was researched by Shimomura et al. (2024) who aimed to predict hourly spot prices from a set of features, determine the importance of said features and examine how renewable energy effects the electricity market price formation. This was supported with a SHAP analysis, which ranks the input features in order of importance with respect to a particular prediction (Lundberg and Lee, 2017) such that the effect of the different features on the hourly spot price could be examined. The results suggest that solar capacity injection suppresses day-ahead spot prices during daylight hours, while simultaneously introducing price volatility during periods of low solar output. In the context of this project, this study can be used in the model validation and exploration phase.

Table 1 presents a summary of the reviewed literature with methods and use-cases.

Table 1: Summary of related literature on electricity market forecasting methods.

Paper	**Context**	**Method**
Azarang and Edling (2025)	mFRR balancing volume prediction across Swedish bidding zones (ENTSO-E data)	LSTM neural network for one-hour-ahead downward activation volume prediction
Våle et al. (2025)	mFRR activation price prediction across Norwegian bidding zones, with focus on price driver identification	XGBoost and EBM comparison, with day-ahead spot price, hydropower production and heating demand as key predictors
Mitridati and Pinson (2018)	Reconstruction of European day-ahead supply curves from historical clearing prices and traded volumes	Hidden Markov Model combined with a Bayesian framework to predict supply curve shape in even-sized generation blocks
Ziel and Steinert (2016)	Day-ahead price forecasting in the German-Austrian spot market via supply and demand curve modeling	Sparse VAR with LASSO on 32 discretized bid volume price classes across 24 hourly systems, with renewable generation and weekday seasonality features
Koechlin et al. (2025)	Day-ahead price forecasting in the Italian spot market via continuous merit-order curve modeling	Functional PCA on historical curves combined with sparse VARX estimated via LASSO, using load, renewable generation and transfer capacity forecasts
Shimomura et al. (2024)	Hourly spot-price prediction to examine the effect on increasing renewable energy sources in the market	XGBoost spot price forecasting with SHAP decomposition of market driver contributions

3 Data description

3.1 Data sources

Fingrid release public data about the Finnish electricity system and markets via their Open Data platform (Fingrid, n.d.). Available data includes for example electricity consumption and generation forecasts (including wind and solar power generation separately), real-time information about nuclear and hydro power generation, as well as procurement forecasts and past clearing prices for balancing capacity markets.

ENTSO-E Transparency Platform (ENTSO-E, 2015) is a pan-European publishing platform for electricity market data, launched in 2015 as a response to the EU transparency regulation 543/2013 (European Commission, 2013), which requires TSOs to publish information about electricity markets in their region. Since 2023, accepted bids on the Finnish mFRR markets have been available via the transparency platform.

Based on the accepted bids, it is possible to perfectly reconstruct the merit order curve up until the clearing price. However, it is important to note that we do not have information about bids that were not accepted (higher bid price than the clearing price). This affects further analysis and the methodological choices for predicting future bid curves.

Finally, we utilize a proprietary dataset of spot electricity price predictions (Anonymous, 2026). This third dataset is used as a static document.

3.2 Data pipeline

Fingrid and ENTSO-E platforms have application programming interfaces (API) that allow downloading datasets via HTTP GET requests. Data gathering is done using tailored scripts that allow downloading latest available data continuously.

We group the ENTSO-E bid data by hour. This means that in case of block bids (bids that offer identical capacity and price for multiple hours at the same time), we create a row of data for each row it applies to. In the end, this unifies the format such that the data can be indexed by hour and each hour value can be used to search for all accepted bids that are applicable to that hour.

3.3 Datasets

One of the primary goals of the project is constructing a real-time prediction model that uses latest available information to predict the mFRR capacity market bid curves for the next day. Since bidding for mFRR capacity market closes the previous day by 8:30 AM, an essential property of any dataset is its publication timeline. The prediction model can only use data that is published by the time bidding closes. For the purposes of this project, we settle for 8:00 AM as a cutoff time for data gathering, as it provides a 30-minute margin.

We note that in reality, initial bidding closes already at 7:30 AM, and on a case by case basis, it may be extended until 8:00 or 8:30 AM (Fingrid, 2024). However, this difference does not change the availability of the datasets.

Table 2 lists Fingrid datasets that offer forecasts about electricity consumption, production, or balancing capacity procurement. The listed datasets are such that complete forecasts for each calendar day are available by 8 AM the previous day.

We note that for consumption and production, Fingrid also publishes more accurate, 24-hour forecasts once a day, but since they are published at 12 PM or 6 PM, they are not suitable for our purpose.

Similarly, the spot price forecasts are also used such that only forecasts available before the 8 AM cutoff are taken into account.

Table 2: Fingrid forecast datasets considered for analysis (criterion: next-day forecasts available by 8 AM). ID refers to the dataset ID in Fingrid Open Data platform.

ID	Dataset name	Additional information
166	Electricity consumption forecast - updated every 15 minutes	72-hour forecast
334	Balancing Capacity (mFRR), up, hourly market, procurement forecast	Forecast for next day
335	Balancing Capacity (mFRR), down, hourly market, procurement forecast	Forecast for next day
241	Electricity production prediction - updated every 15 minutes	1-week forecast
245	Wind power generation forecast - updated every 15 minutes	72-hour forecast
248	Solar power generation forecast - updated every 15 minutes	72-hour forecast

In addition to the forecast datasets, we fetch multiple Fingrid datasets that update in real time. These are listed in Table 3. For all datasets whose data period is shorter than one hour, only the values from even hours are taken into account, to match the data period of mFRR clearing information (datasets 327–330).

All data is gathered from January 1st 2025 to February 28th 2026. This provides a 14-month period for training and validation. As our initial analysis suggests (see Section 4), the electricity market has changed substantially in recent years, which means that older data may not reflect the current market dynamics well. Furthermore, because the data resolution is short (once per hour), the amount of data points is very large even for a 14-month period.

However, market clearing data (datasets 327–330) is gathered already starting from 2024 to allow analyzing the changes in mFRR capacity markets over a longer time.

We note that the historical nuclear power and hydro power production data are available via the Fingrid API in a context window of only approximately 3 months. This severely limits the usability of that data.

3.4 Consistency between Fingrid and ENTSO-E data

Fingrid publishes information about the clearing price and cleared quantity of mFRR balancing capacity (datasets 327–330 in Table 3). Because the clearing capacity determines the bids that are accepted, and subsequently the clearing price, conversely, it is possible to determine the clearing price and clearing quantity from the accepted bids. For any hour, the clearing price is the maximum price and the clearing capacity is the sum of capacity among the accepted bids.

Table 3: Fingrid real-time datasets considered for analysis. ID refers to the dataset ID in Fingrid Open Data platform.

ID	**Dataset name**	**Additional information**
331	Balancing Capacity (mFRR), down, hourly market, bids	Total offered quantity (MW)
332	Balancing Capacity (mFRR), up, hourly market, bids	Total offered quantity (MW)
188	Nuclear power production - real-time data	Updated every 3 minutes
191	Hydro power production - real-time data	Updated every 3 minutes
375	Activated mFRR balancing regulation upward sum	15-minute periods
376	Activated mFRR balancing regulation downward sum	15-minute periods
391	Activated mFRR special regulation upward sum	15-minute periods
392	Activated mFRR special regulation downward sum	15-minute periods
327	Balancing Capacity (mFRR), up, hourly market, procured volume	Updates once in 24 hours
328	Balancing Capacity (mFRR), down, hourly market, procured volume	Updates once in 24 hours
329	Balancing Capacity Market (mFRR), up, hourly market, price	Updates once in 24 hours
330	Balancing Capacity (mFRR), down, hourly market, price	Updates once in 24 hours

Proceeding as described, we can determine the theoretical clearing price and quantity from the bid data. By comparing these to the figures published by Fingrid, we can verify the completeness and accuracy of the transparency platform data. We compute the hour-level differences as $q_E - q_F$, where q_E is the theoretical clearing quantity/price and q_F is the published clearing quantity/price. Table 4 presents a summary of the hour-level differences while Table 5 shows the counts by direction of discrepancy.

The summary shows 10176 individual hours of up-capacity information and 10130 hours of down-capacity information. In total, the data period contains 46 hours where down-capacity demand was 0, and additionally 24 hours where both down and up-capacity demand was 0. For simplicity, these are not included in the summaries since there are no accepted bids for these hours.

For both down and up-capacity, cleared quantities are consistently higher when calculated from the bid data. A possible explanation is that bids can be non-divisible, which means the bid may not be accepted partially (Fingrid, 2024). In these cases, the cumulative quantity might exceed the minimum procured quantity. In fewer than 0.3% of the rows, the cumulative quantity from ENTSO-E data is smaller than the true procured quantity, which indicates missing data. This proportion is so small that its effect to modeling is likely negligible.

The clearing prices are mostly consistent between ENTSO-E and Fingrid datasets. The prices match exactly in over 84% (up-capacity) and 92% (down-capacity) of the cases. The minor but noticeable prevalence of mismatches may stem from how block bids are handled. The specifics of price formation can be slightly different in the presence of block bids.

Table 4: Summary statistics of difference between clearing information (quantity and price) calculated from ENTSO-E bid data and Fingrid published clearing information (ENTSO-E – Fingrid).

	Up-capacity		**Down-capacity**	
Metric	**Quant. (MW)**	**Price (€/MW)**	**Quant. (MW)**	**Price (€/MW)**
Count	10176	10176	10130	10130
Mean	7.69	-0.02	9.23	0.01
Std. Dev.	25.80	9.27	38.33	2.20
Min	-342.00	-121.00	-440.00	-44.00
25%	1.00	0.00	1.00	0.00
50%	4.00	0.00	5.00	0.00
75%	9.00	0.00	10.00	0.00
Max	475.00	494.06	717.00	27.00

Table 5: Distribution of discrepancy directions (ENTSO-E – Fingrid).

	Up-capacity		**Down-capacity**	
Sign of difference	**Quantity**	**Price**	**Quantity**	**Price**
Positive (ENTSO-E > Fingrid)	7,800	431	8,396	319
Zero (ENTSO-E = Fingrid)	2,351	8,574	1,710	9,373
Negative (ENTSO-E < Fingrid)	25	1,171	24	438

4 Historical analysis of market clearing

In order to illustrate how the mFRR market dynamics have changed in the recent years, it is useful to analyze the historical prices and supply and demand quantities of both mFRR up- and down-capacity markets. The dataset used in this analysis spans from 1.1.2024 to 19.5.2026 (Fingrid (n.d.)). This chapter presents an analysis of the development of the mFRR capacity market prices.

Since the mFRR capacity market is cleared for each hour, a descriptive way to understand the market dynamics is to visualize all cleared hours as a price-quantity scatter plot. Figure 2 displays the yearly scatters of the mFRR up-capacity market clearing datapoints.

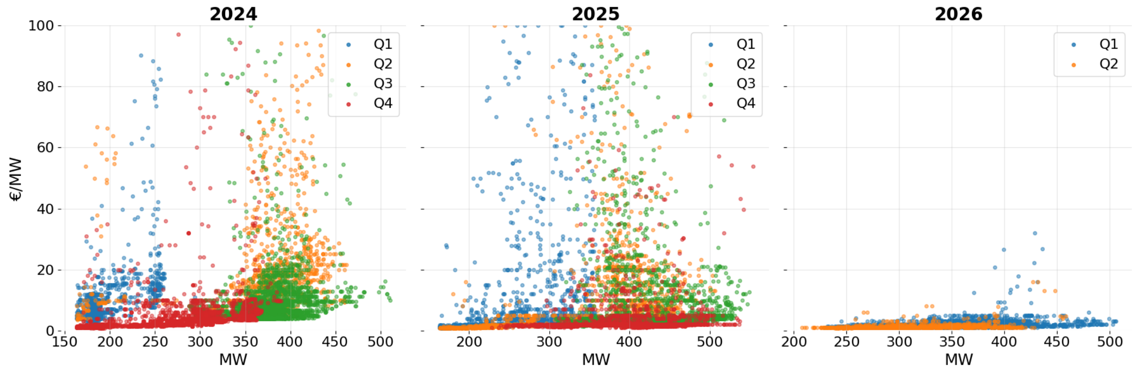

Figure 2: mFRR up-capacity price-quantity pairs of all cleared hours since 2024.

Figure 2 shows that the clearing prices in Q1, Q2 of 2026 have collapsed in comparison to the first quarters of previous years. The monthly average prices have diminished to well below 5 €/MW, as visualized in Figure 3.

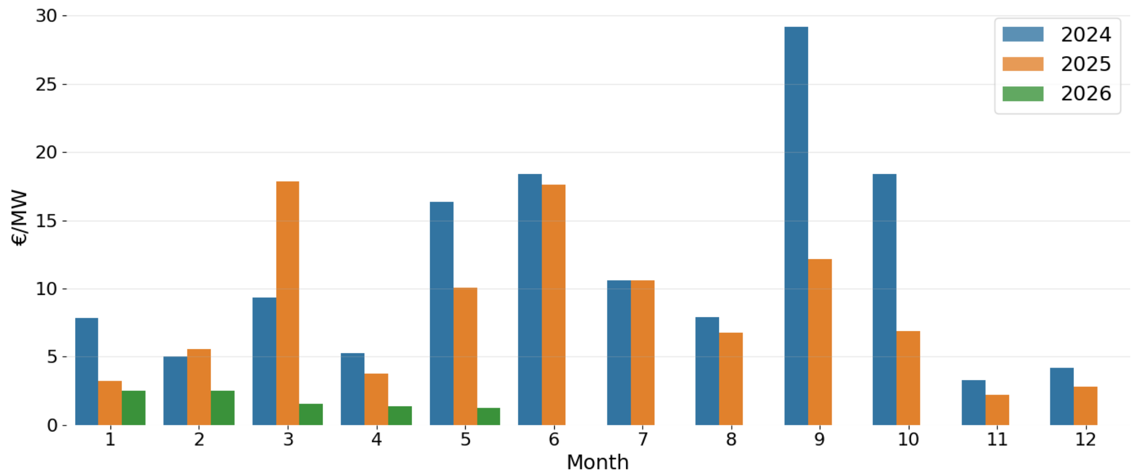

Figure 3: mFRR up-capacity monthly average prices.

While the mFRR up- and down-capacity markets seem similar, they often attract distinct market participants. For instance, wind producers may bid in the down-capacity market by having the option to curtail production during high-wind periods. However, it is not sensible for wind producers to participate in the up-capacity markets. Therefore, the market dynamics between the mFRR up- and down-capacity markets are likely to differ exhibit different dynamics. Figures 4 and 5 show the same visualizations for the mFRR down-capacity market.

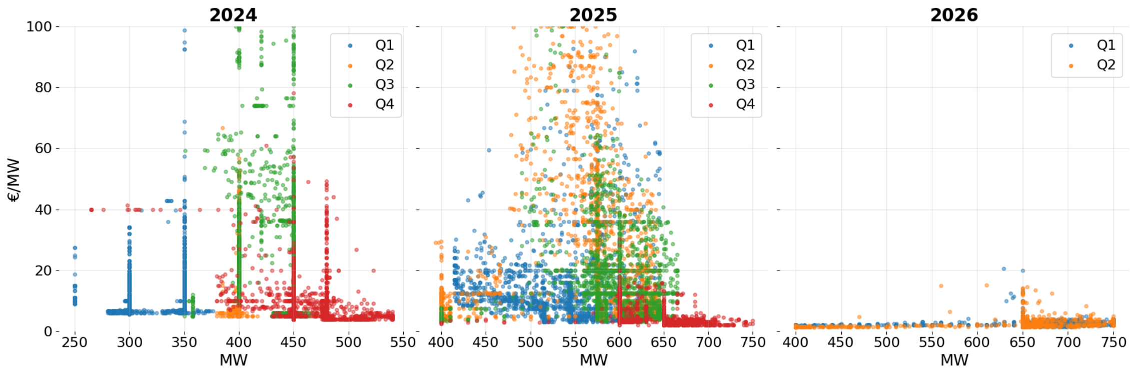

Figure 4: mFRR down-capacity price-quantity pairs of all cleared hours since 2024.

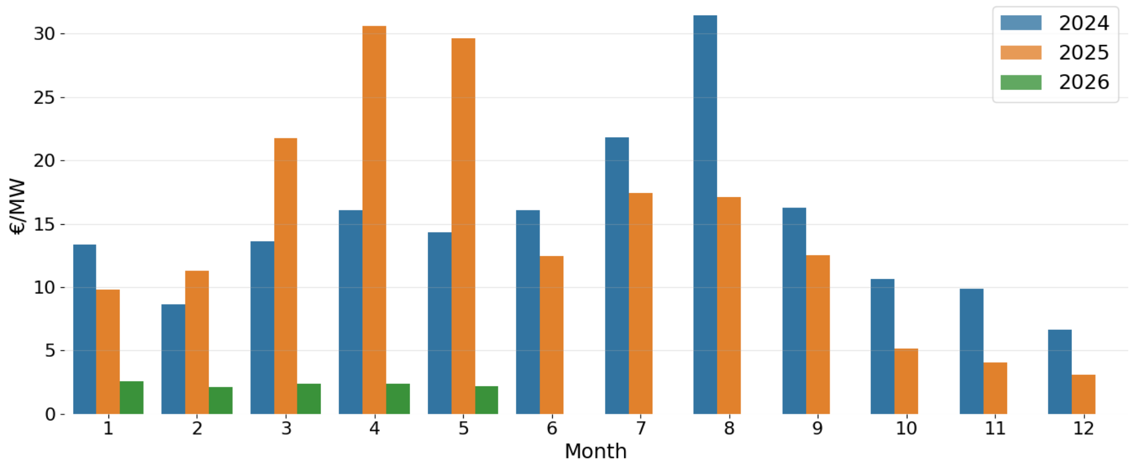

Figure 5: mFRR down-capacity monthly average prices.

The figures show that the mFRR down-capacity market has developed in a way similar to the up-capacity market. Since the prices are directly a result of the intersection of supply and demand bid curves, the change in prices is likely either from increased supply or decreased demand. The open data source by Fingrid (n.d.) also provides data on the quantity of all supply bids in MW, as well as the demand, which is set for each hour by Fingrid. Figure 6 displays the time series of supply and demand quantities since 2024.

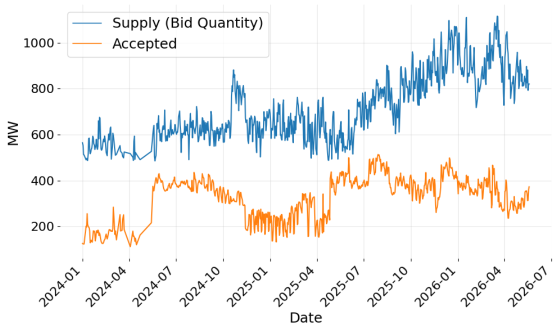

(a) mFRR up-capacity

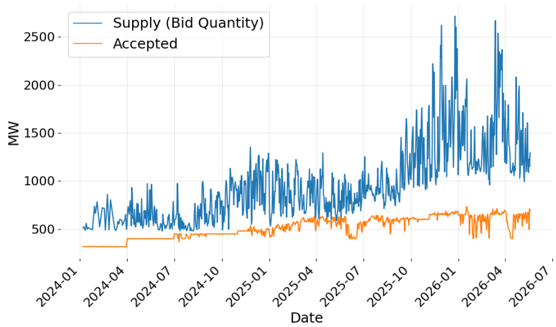

(b) mFRR down-capacity

Figure 6: mFRR up- and down-capacity bid quantities (supply) and accepted bids (demand), in MW.

Figure 6 shows that both markets have become oversupplied since the beginning of 2026. The sudden shift in supply is a result of Battery Electricity Storage Systems (BESS) entering the Finnish grid quickly. On 19.5.2026, the BESS capacity in Finland exceeded 1290 MW, which is a significant increase from 2025 (Fingrid, n.d.).

5 Methodology

5.1 Feature engineering

The feature engineering pipeline for mFRR capacity price forecasting involves transforming accepted bid data into a comprehensive set of predictor variables. The approach addresses the unique temporal structure of mFRR markets, where bids for all 24 hours of the following day must be submitted simultaneously before an 8:30 morning deadline. This constraint means that careful consideration in feature engineering is needed to ensure that data leakage from the future is not exhibited in the model.

5.1.1 Bid binning

Accepted bid data from ENTSO-E is organized into price bins to capture the distribution of market offers across different price ranges. Inspiration is drawn loosely from Ziel and Steinert (2016). Compared to using the bid prices as continuous variables, categorization allows using simpler analysis methods. For each bin, direction (Up or Down), and hour, two metrics are computed: the total number of bids (`bin_total_bids`) and the total capacity offered (`bin_total_quantity`)

We aim to select the bin boundaries such that interpretable insights concerning the bidding behaviour can be made. Unnecessarily sparse binning leads to information loss, whereas too dense bins may be prone to overfitting and computational inefficiency. The bid data reveals that mFRR capacity bids are most often priced at round figures, such as 0.5 or 2.0 euros. Figure 7 shows cumulative distributions of the number of bids made at different prices, in up and down capacity markets separately. The sharp, almost vertical jumps of the curves visualize how the bid prices are usually concentrated at specific values. The highlighed values correspond to cumulative counts at prices 0.5, 1.0, 1.5, 2.0, 3.0, 5.0 and 10.0 euros. These values are selected as the bin boundaries as they offer fairly even distribution of bid counts. For simplicity, we use the same boundaries for up and down-market bins.

The upper end point of the bins is inclusive. For example, from Figure 7a we see that during the data period, there are 70555 up-capacity bids at $[0.0, 0.5]$ euros, and $197524 - 70555 = 126969$ up-capacity bids at prices $(0.5, 1.0]$ euros.

Figures 7c–7d show that for both up and down markets, procurement prices over 10 euros have very rarely been accepted. We note that while the plots are truncated to 100 euros, both datasets contain individual data points over 100 euros. These are deemed outliers that are irrelevant to our research questions project, and thus are not analyzed in detail.

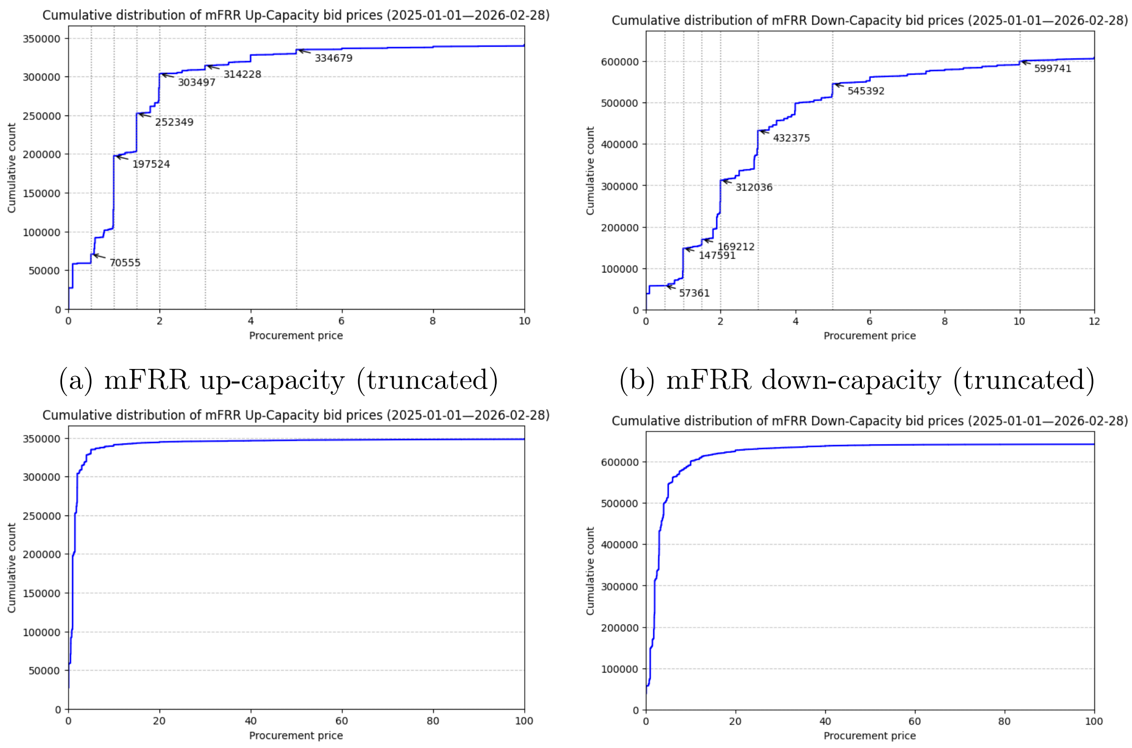

(a) mFRR up-capacity (truncated)

(b) mFRR down-capacity (truncated)

(c) mFRR up-capacity (extended)

(d) mFRR down-capacity (extended)

Figure 7: Empirical cumulative distributions (unnormalized) of mFRR capacity bids at different procurement prices during Jan 2025 – Feb 2026. (a)–(b) truncated x-axis, (c)–(d) extended x-axis.

It is important to note that because the data only includes accepted bids, the bid data is truncated to the realized clearing. Thus, if the clearing price falls within the range of a bin or below it, the binned variables for that hour are set to NULL, as partial bin data would not reliably represent the full distribution of bids inside that bin. For example, if the clearing price was 1.8 euros, the bin $(1.5, 2.0]$ euros does not contain possible bids made at 1.8 to 2.0 euros and is thus an incomplete representation of the number of bids and the offered capacity at that price range. Setting incomplete bins to NULL allows a clear distinction between structural zeros and missing data.

To be able to utilize as much of the data as possible, we extend the bins slightly, adding four more boundary points: 20.0, 30.0 and 50.0. Thus, any bid higher than

50 euros is not included in the analyses by construction. The final bins are:

$$[0.0, 0.5], (0.5, 1.0], (1.0, 1.5], (1.5, 2.0], (2.0, 3.0], (3.0, 5.0],$$
$$(5.0, 10.0], (10.0, 20.0], (20.0, 30.0], (30.0, 50.0]$$

Figure 8 shows the proportion of NULL values in the binned variables. Down capacity prices are higher on average, which is reflected by the slower incline of the null proportions as compared to up-capacity.

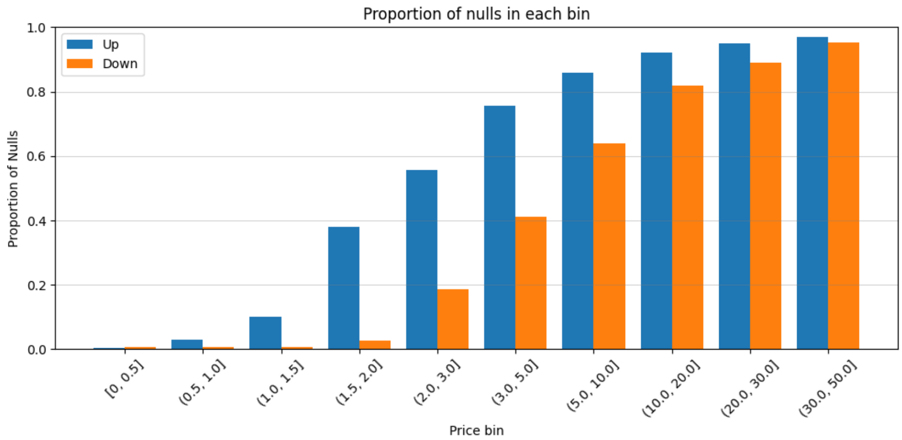

Figure 8: Proportion of null values (clearing price less than the upper bin boundary) in each price bin, during Jan 2025 – Feb 2026.

Figure 9 shows the time series of the binned bid quantity in MW. Daily averages have been computed to reduce plotting noise and allow better spotting temporal shift. Up-capacity quantities are plotted on the positive y-axis, while down-capacity quantities are on the negative y-axis. From the plot, it is apparent that the amount of up-capacity offered at the lowest price category (at most 0.5 euros), has increased from near zero values during the first half of 2025, and stayed relatively constant since. This supports the hypothesis that there might be a static component to the bid curves, resulting from the increased battery capacity in electricity grid.

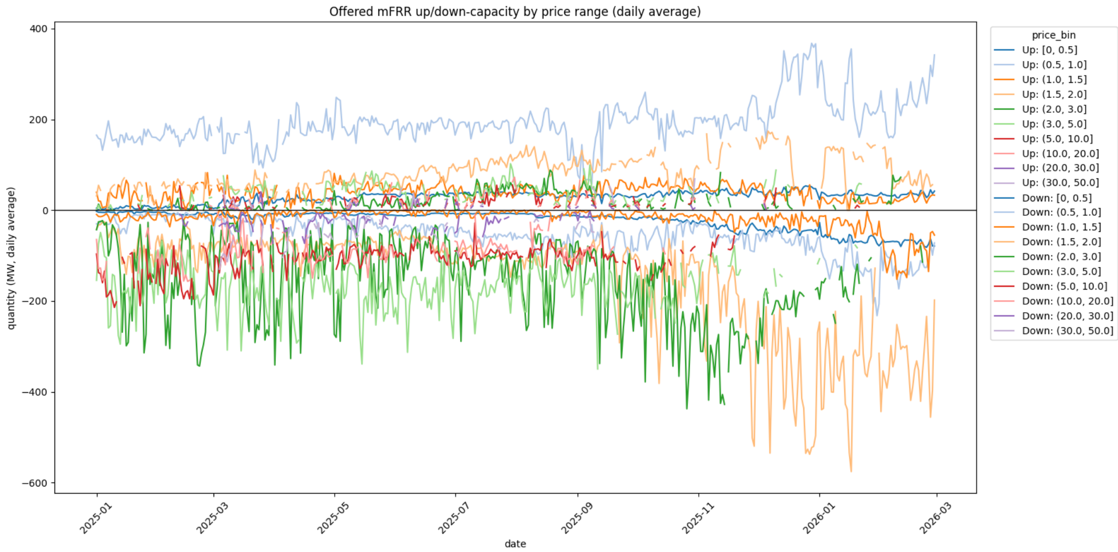

Figure 9: Amount of capacity (MW) offered at each price bin during Jan 2025 – Feb 2026.

5.1.2 Lagged features

Given the 24-hour day-ahead bidding structure, classical autoregressive features require careful design. Four classes of lagged features are considered, though the final model utilizes dynamic lags exclusively.

Dynamic lags compute values from a fixed number of hours in the past relative to the target hour. Four lag horizons are employed: 24 hours (previous day, same hour), 48 hours (two days prior), 72 hours (three days prior), and 168 hours (one week prior). These lags capture weekly seasonality and recent market trends while respecting the constraint that same-day information is unavailable at bidding time. Dynamic lags are constructed for all binned bid metrics (10 bins $\times$ 2 directions $\times$ 2 metrics $=$ 40 variables) and clearing prices and quantities (4 variables), resulting in a total of 176 lagged features (44 base variables $\times$ 4 lag horizons).

Other lagged features we consider are lagged values from fixed hours of the previous day (for example, using the values from a peak consumption hour as a predictor for all hours of the next day), mean values of multiple fixed hour lags (for example, the mean of all 24 hours from the previous day), as well as log-differences of the values between various time points. Due to increased model complexity, the lack of resources to thoroughly evaluate the feature selection for these types of variables, and the lack of noticeable improvement in model performance based on initial testing, these are not implemented for the final model.

5.1.3 Real-time variables

Real-time operational variables are extracted from the early morning hours of the previous day (before the 8:30 AM bidding deadline) to ensure no data leakage from the future. These variables include realized nuclear and hydro generation (`nuc_act`, `hydro_act`), total mFRR bid quantities for both up- and down-directions, activated balancing and special regulation volumes, and forecasts for consumption, generation, wind, and solar production. Since these values are published before the bidding deadline, they represent information available to market participants when placing bids.

5.1.4 Forecast variables

Day-ahead forecasts for the target hours are incorporated directly without lagging, as these forecasts are published before the bidding deadline. Seven forecast variables are included: consumption (`consumption_updated_fct`), total generation (`total_generation_updated_fct`), wind generation (`wind_updated_fct`), solar generation (`solar_updated_fct`), spot price (`spot_price_fct`), and mFRR procurement forecasts for both directions (`mfrr_up_procurement_fct`, `mfrr_down_procurement_fct`).

5.1.5 Calendar features

A single binary calendar feature (`is_weekend_or_holiday`) indicates whether the target hour falls on a weekend or Finnish public holiday. This feature captures systematic differences in electricity consumption and market behavior during non-working days, although it exhibits some multicollinearity with consumption.

5.1.6 Final feature set

The complete feature matrix comprises 194 features organized into four classes (Table 6).

The response variables consist of the binned bid metrics (40 variables) and clearing prices and quantities (4 variables) for the target hour, providing a representation of market outcomes across the price distribution. Forecasting the binned bids accurately would enable constructing a forecast for the total merit order curves.

Table 6: Overview of the feature matrix

Feature class	Variable / Sub-category	Count	Details / Horizon
1. Lagged features	Binned bid counts (up direction)	40	10 bins × 4 lags (24h, 48h, 72h, 168h)
	Binned bid quantities (up direction)	40	10 bins × 4 lags (24h, 48h, 72h, 168h)
	Binned bid counts (down direction)	40	10 bins × 4 lags (24h, 48h, 72h, 168h)
	Binned bid quantities (down direction)	40	10 bins × 4 lags (24h, 48h, 72h, 168h)
	Clearing prices and quantities (up/down)	16	4 variables × 4 lags (24h, 48h, 72h, 168h)
2. Real-time variables	`nuc_act`	1	Actual nuclear generation
	`hydro_act`	1	Actual hydro generation
	`mfrr_up_all_bids_quantity`	1	Total mFRR up-regulating bids quantity
	`mfrr_down_all_bids_quantity`	1	Total mFRR down-regulating bids quantity
	`mfrr_up_activated_balancing_volume`	1	Activated mFRR up balancing volume
	`mfrr_down_activated_balancing_volume`	1	Activated mFRR down balancing volume
	`mfrr_up_activated_special_volume`	1	Activated mFRR up special volume
	`mfrr_down_activated_special_volume`	1	Activated mFRR down special volume
	`consumption_updated_fct_realtime`	1	Real-time updated consumption forecast
	`total_generation_updated_fct_realtime`	1	Real-time updated generation forecast
3. Forecast variables	`consumption_updated_fct`	1	Consumption forecast
	`mfrr_up_procurement_fct`	1	mFRR up procurement forecast
	`mfrr_down_procurement_fct`	1	mFRR down procurement forecast
	`total_generation_updated_fct`	1	Total generation forecast
	`wind_updated_fct`	1	Wind generation forecast
	`solar_updated_fct`	1	Solar generation forecast
	`spot_price_fct`	1	Day-ahead spot price forecast
4. Calendar features	`is_weekend_or_holiday`	1	Indicator for weekends and public holidays
Total		**194**	

5.2 Prediction model for mFRR price curves

The prediction task involves forecasting the quantity of bids across different price bins for both mFRR up- and down-capacity markets. Rather than predicting a single clearing price or aggregate capacity, the model constructs complete bid curves using independent price bin forecasts. This approach enables the reconstruction of full merit order bid curves, thus supporting decision-making in energy markets. To predict the price bin forecasts by the features mentioned in Section 5.1.6, we employ XGBoost model from XGBoost Python package. Using simpler and more interpretable linear regression model is restricted due to missing values (nulls) in price bin quantities, while XGBoost has support for missing values by default (Chen et al., 2026).

To prevent feature instability, the prediction model does not utilize all of the lagged features. Since each price bin is predicted as a separate response variable, we can perform targeted variable selection. We include the 48h, 72h and 168h lags as predictors only for the response variable. For all other bins, and the clearing prices and quantities, we only include the 24h lagged variable as a predictor in the model. Thus, each model is fitted using 63 predictor variables instead of all 194.

5.3 Model evaluation

5.3.1 Evaluation metrics

The results of the model are evaluated using the following metrics.

Mean Absolute error (MAE) measures the average absolute difference between the actual values and the predicted values, defined as

$$MAE = \frac{\sum_t |y_t - \hat{y}_t|}{T},$$

where y_t is the actual value, $\hat{y}_t$ is the forecast value at time t and T is the total number of predicted time steps. MAE is scale-dependent and expressed in the same units as predicted values.

Weighted Average Percentage Error (WAPE) was originally introduced by Kolassa and Schütz (2007) and is defined as

$$WAPE = \frac{\sum_t |y_t - \hat{y}_t|}{\sum_t |y_t|},$$

where y_t is the actual value and $\hat{y}_t$ is the forecast value at time t. Compared to mean absolute percentage error, WAPE performs better when actual values y_t are close to zero. WAPE is a more suitable measure than MAE, when comparing errors where scales are not consistent.

Root Mean Square Error (RMSE) measures the square root of the mean squared error and is defined as

$$RMSE = \sqrt{\frac{\sum_t (y_t - \hat{y}_t)^2}{T}},$$

where y_t is the actual value, $\hat{y}_t$ is the forecast value at time t and T is total number of predicted time steps. RMSE penalizes larger errors while keeping the measurement in the original units, making the metric easy to interpret and useful when large errors are especially undesirable.

5.3.2 SHAP values (feature contribution)

In addition to evaluating the performance of the model, we interpret the contribution of the different features to the prediction value of the model. As mentioned in Section 2, interpretability of "black-box" models is important. We seek to obtain this understanding with SHAP analysis by using SHAP Python package.

SHAP can explain specific predictions by assigning SHAP values to each feature; these values represent the effect of the features on the predicted value. More precisely SHAP value is the effect of feature on difference from a base value. The base value is defined as the expected value of the prediction without any knowledge of feature values (Lundberg and Lee, 2017). Now a single prediction can be expressed as

$$\text{prediction} = \text{base value} + \text{sum of SHAP values}.$$

Note that SHAP values are of the same scale and units as prediction values, making interpretation easier.

While SHAP values are local approximations of the effect of features, some understanding of the global importance and effects of the features can be examined by combining SHAP values over multiple or all predictions. This is later done to provide insight into whether high or low feature values were associated with increases or decreases in predicted values.

5.3.3 Expanding window validation

Model validation is performed using expanding window validation to assess how well our XGBoost model generalizes to unseen data. In the expanding window validation method, the training set expands over time, while the test set is always the next time period. Once data is included in one training period, it is included in all future training windows. Expanding window validation maintains the chronological order of data and future is always predicted on past training data. The method also keeps all past data and the last training window uses all data available for training. Our first window includes training data from start of the year 2025 until end of August 2025 and September 2025 as a validation data. After the first window, each new training window is expanded by one calendar month. Expanding window validation with our selected window splits is illustrated in figure 10.

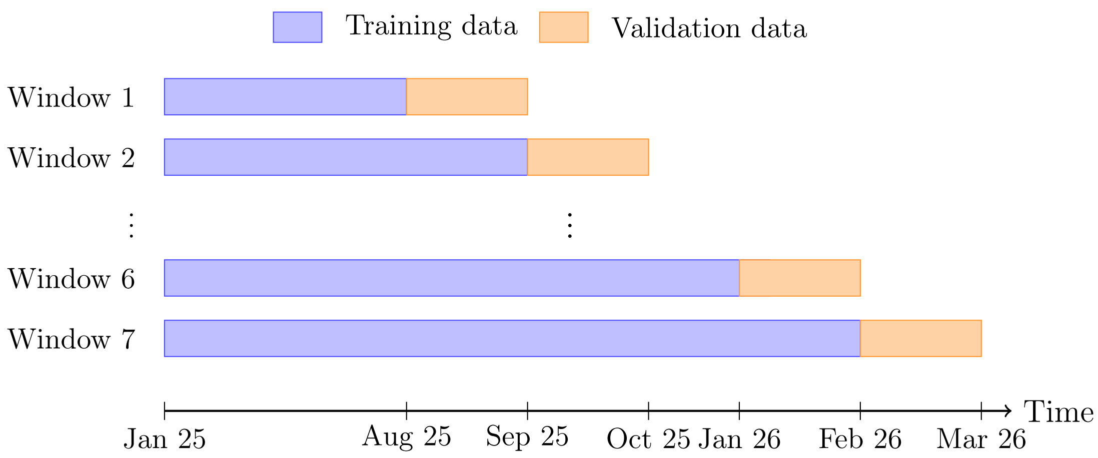

Figure 10: Illustration of expanding window validation.

5.4 Price impact of bid injection

For an electricity provider, it is especially important to know how the market would behave given a certain bid, including its corresponding offered capacity and price. This is a question of uncertain payoffs for the provider as generally a provider seeks to maximize their profit for the participation in the market. As our project aimed to reconstruct the merit order curve from historical data, this is essentially the first part of the uncertainty problem. Fundamentally, given an electricity provider with a prediction model of a future time periods reconstructed merit order curve, it is of interest to know how the curve behaves given a new injection of capacity at a certain price. As our model aims to predict the order curve up to the clearing price, introducing an additional bid with a price below the predicted clearing price will inherently reduce the clearing price. We check the performance of our model by predicting a merit order curve, injecting the market with 50 MW of additional capacity and zero price, and comparing how this performs against a naive 24h-lag model.

Figure 11: Price drop distributions for an injection of 50MW

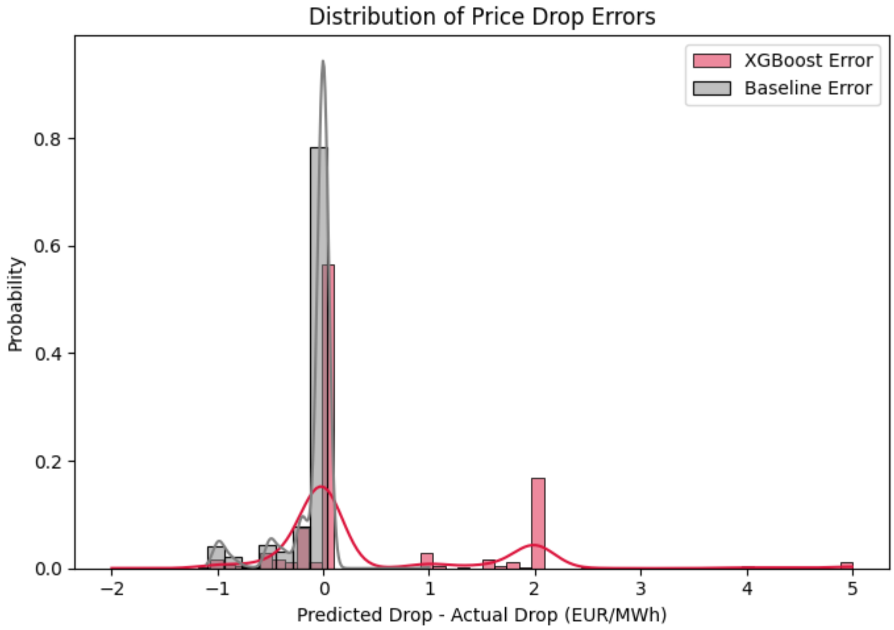

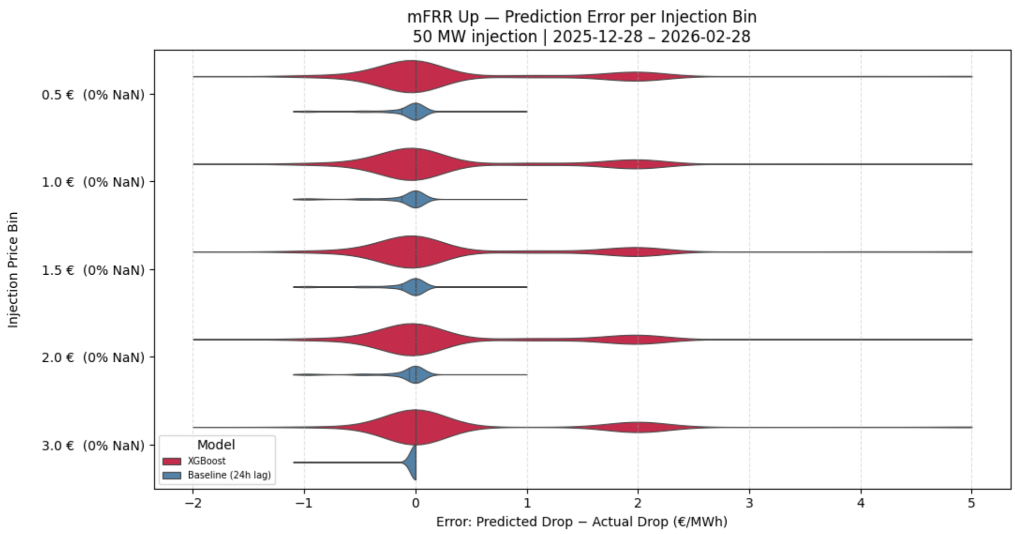

As we see, our XGBoost model has a persistent positive bias where the difference in price given an 50 MW injection is evaluated to be too large. This is likely stemming from the binning principle in this project as the model registers all bids to be on the upper end of the corresponding bin. This means that a injection of additional capacity either does not cause any change in price as the addition is not sufficient to lower the clearing price to the next bin, or it has moves the clearing price to another bin causing a probability spike in the 2 EUR/MWH section.

As our model cannot effectively predict the drop in price given an injection, we decided to inspect historical data from Fingrid and ENTSO-E and use these findings to extrapolate insights into the prediction model. This can be done by taking past merit order information and simulating novel bids into and averaging the change in clearing price caused by the new bid. The matrix below quantifies the price impact of a novel capacity injection using the modern state of the mFRR market. Each cell shows the mean drop in clearing price resulting from injecting a given volume at a fraction of the original clearing price:

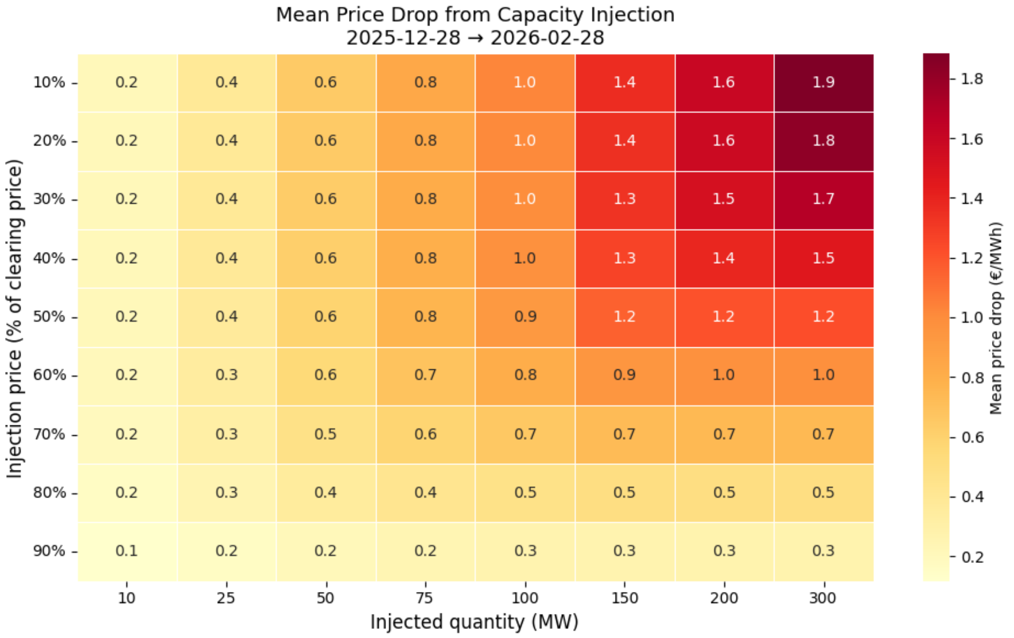

Injection price (% of clearing price)	10	25	50	75	100	150	200	300
10%	0.2	0.4	0.6	0.8	1.0	1.4	1.6	1.9
20%	0.2	0.4	0.6	0.8	1.0	1.4	1.6	1.8
30%	0.2	0.4	0.6	0.8	1.0	1.3	1.5	1.7
40%	0.2	0.4	0.6	0.8	1.0	1.3	1.4	1.5
50%	0.2	0.4	0.6	0.8	0.9	1.2	1.2	1.2
60%	0.2	0.3	0.6	0.7	0.8	0.9	1.0	1.0
70%	0.2	0.3	0.5	0.6	0.7	0.7	0.7	0.7
80%	0.2	0.3	0.4	0.4	0.5	0.5	0.5	0.5
90%	0.1	0.2	0.2	0.2	0.3	0.3	0.3	0.3

Figure 12: Mean drop in clearing price given price and injected quantity

From this heatmap we see a clear relationship between the injection price and quantity as large injections at a small price lowers the price substantially pricing out several competitors outside of the clearing price. Given additional time and resources to develop our choice in methodology, this effect on clearing price given an injection could be further developed to include a novel solution to a optimization problem for the market participant to optimize their bid price and volume.

6 Results and model validation

6.1 Results and validation metrics

Figures 13a, 13b, 13c and 13d illustrate examples of up and down capacity bid merit order curves, which are constructed from predicted price bin quantities. Along with the curves, the plots include true clearing prices and quantities as well as true merit order curves of accepted capacity bids. The procurement forecast from Fingrid (datasets 334 and 335 in Table 2) is also plotted for reference.

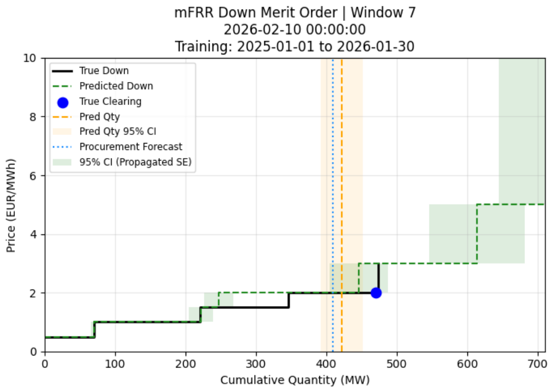

(a)

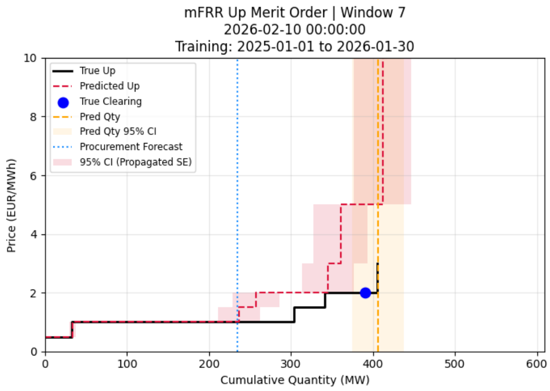

(b)

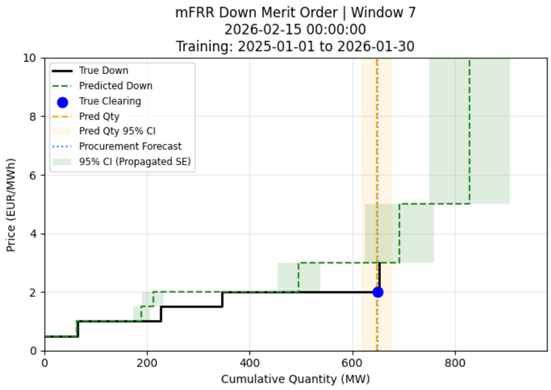

(c)

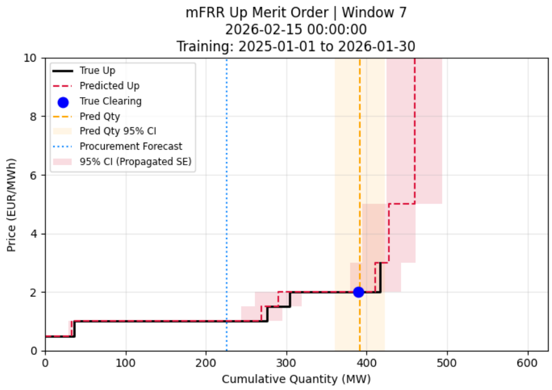

(d)

Figure 13: Examples of predicted merit order curves of up- and down-capacity bids.

Visualizations of the model performance based on MAE, WAPE and RMSE can be seen in figures 14, 15 and 16. In these plots, x-axis represents the expanding training windows 1-7 from section 5.3.3 and y-axis represents the predicted values (response variable). The values of the metrics are indicated by heatmap coloring. Additionally we compared MAE and RMSE values against a baseline predictions constructed using naive predictions equal to the 24-hour lagged values. Differences in the metrics between the baseline and our XGBoost model are presented in Figures 17 and 18. Compared to the baseline, there is no clear consistency of our model performing better on some response variables or in specific validation windows.

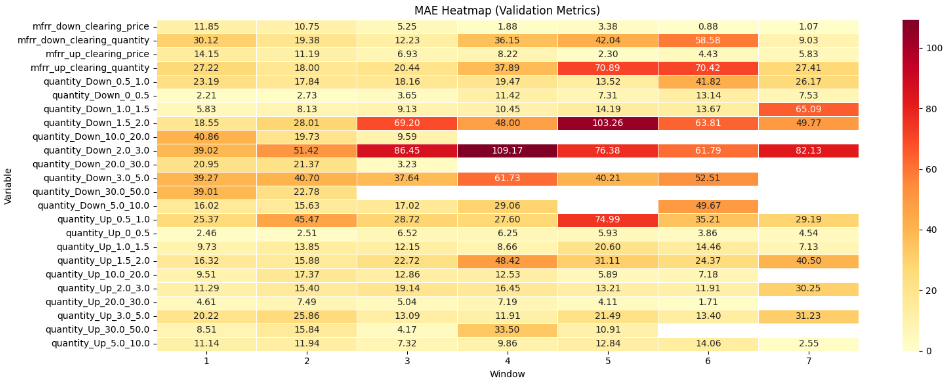

Figure 14: MAE values for all the validation windows and predicted values.

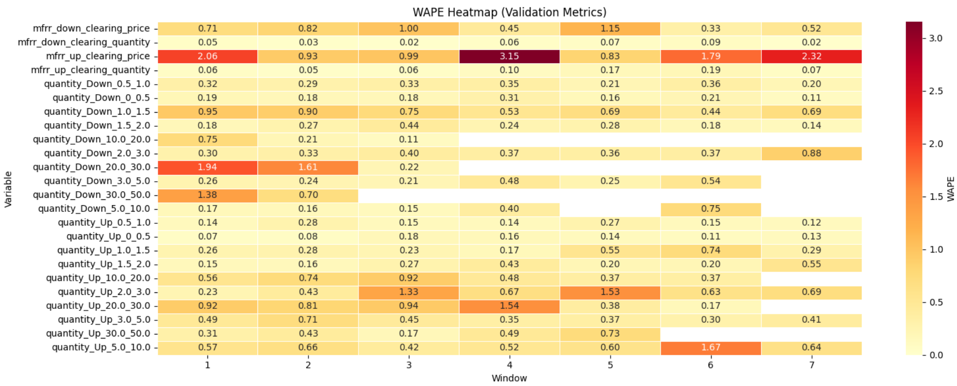

Figure 15: WAPE values for all the validation windows and predicted values.

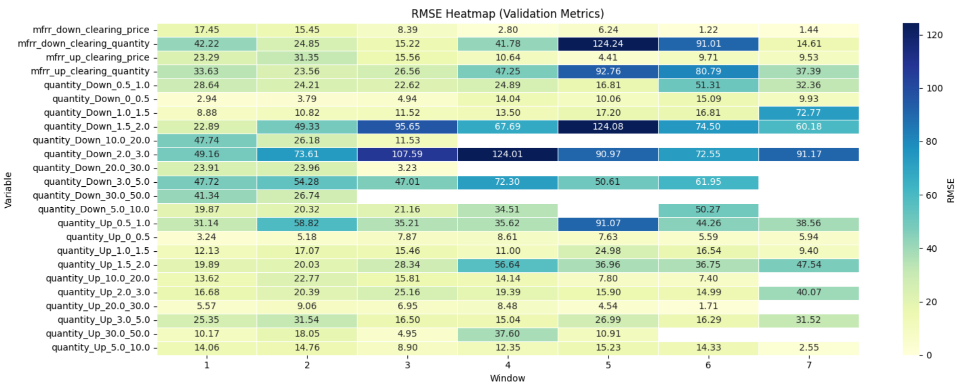

Figure 16: RMSE values for all the validation windows and predicted values.

It is apparent that our XGBoost model is best suited for predicting the clearing quantity of a given hour; with a mean error 9.03 MW in down-capacity (2% WAPE), and 27.41 MW in up-capacity (7% WAPE). This is likely largely due to the fact that Fingrid publishes the procurement forecast, which is expected to offer an accurate baseline for predicting it. Interestingly, based on our findings, the procurement forecast systematically underestimates the realized procurement in the up-markets, which is why our model rarely uses the original prediction, but instead tends to predict higher (see e.g. Figure 13b, 13d). It is possible that Fingrid's forecast often

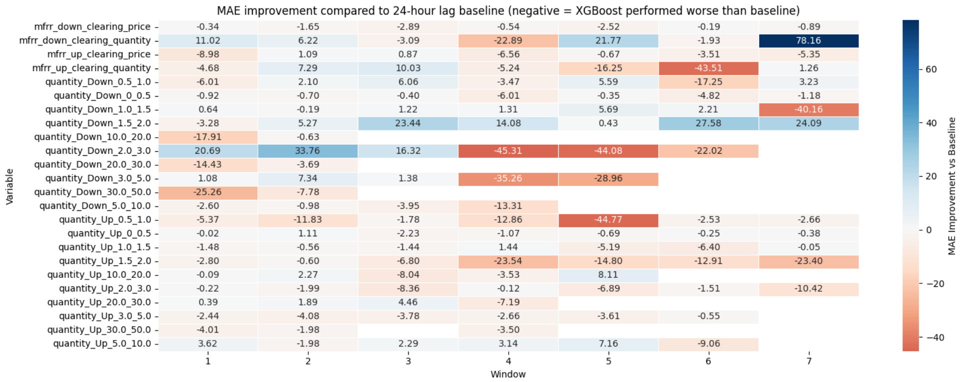

Figure 17: Comparison of MAE values between baseline and XGBoost model.

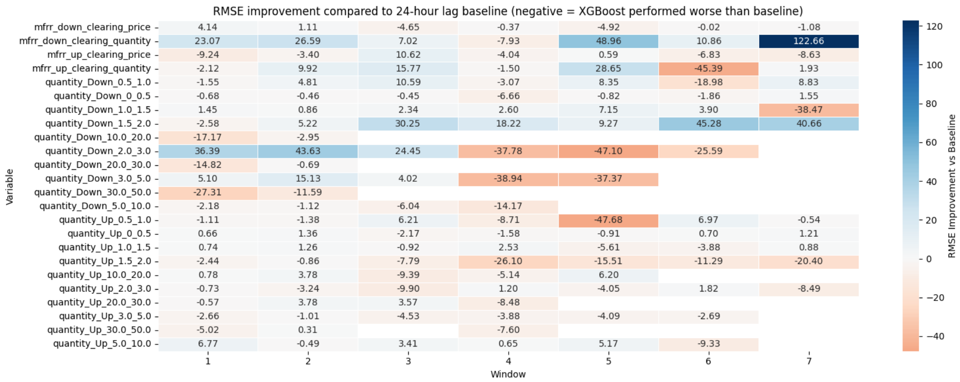

Figure 18: Comparison of RMSE values between baseline and XGBoost model.

underestimates the procurement by a constant amount, which would explain the small validation error.

The validation errors of bin-specific quantities are consistently higher than the errors of the clearing quantity. As an example, the down-quantity bin of 1–1.5 euros has 69% WAPE, and the up-quantity bin of 1.5–2 euros has 55% WAPE (see Figure 15). Additionally, it is important to note that the predicted merit order curve is subject to cumulating errors from multiple different price bins.

Figures 17–18 show that in terms of the validation metrics, the baseline forecast performs better than the XGBoost model in most of the binned variables. This indicates that the shape of the merit order curve is subject to variation not accounted for by our predictive features; taking the merit order curve from 24 hours prior often yields similar or better predictive power. Furthermore, comparing the performance in different training/validation windows reveals no systematic improvement in model performance from expanding the training window. This might support the claim that the market behaviour in mFRR capacity markets is largely unpredictable, and due to evolving electricity grid structure, data loses its relevance for prediction purposes quickly.

Based on the results, the predictive factors used in this project offer limited utility for reconstructing the future mFRR price curves. A possible contributor for large validation errors is that separate models are used to predict the bids for each price bin, independently of each other. Future research should aim for modeling the interactions between the price bins in greater detail. In addition, since the clearing quantity could be predicted with a very small mean error, utilizing the clearing price predictions to guide the formation of the price curve is another potential technique worth exploring.

6.2 SHAP analysis

Figure 19 shows a waterfall plot of SHAP values for the prediction of up-capacity quantity in price bin $(0.5, 1.0]$ in window 7. The plot illustrates the individual impacts of 19 features and the combined impact of the remaining 44 features. In the window average prediction (base value) is 191 MW and the specific prediction reaches 228 MW, from the plot it can be been that spot price prediction, consumption forecast and 24 hour lag of the predicted value were the most influential features.

SHAP value summary plot, shown in Figure 20, visualizes how high or low feature values are associated with increases or decreases in predicted values. The plot visualizes SHAP values for the prediction of the up-capacity quantity in the price bin $(0.5, 1.0]$ in window 7, colors of the dots represent the value of the features, while the distribution density of the values is indicated by the area of the dots. The plot reveals that for example high value of 24 hour lag of the predicted value tends to increase the predicted value. On the other hand, higher values of spot price prediction seem to lower the predicted value.

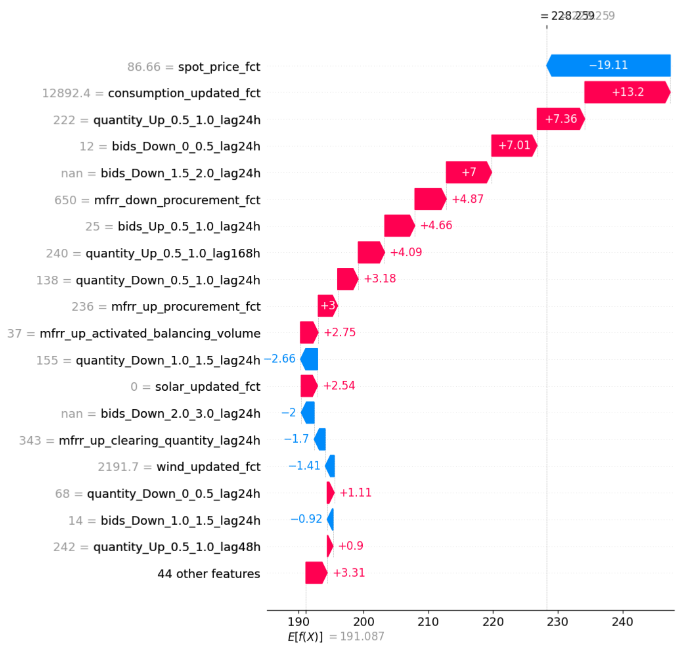

Figure 19: SHAP-values of features for target hour 2026-02-15 15:00.

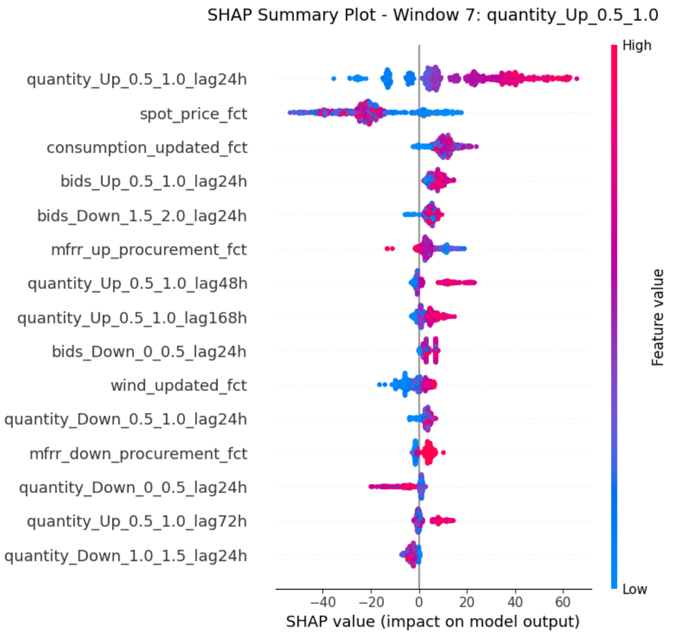

Figure 20: Summary plot of SHAP values in window 7 for the prediction of up-capacity quantity in price bin $(0.5, 1.0]$.

7 Discussion and future research

As energy markets become increasingly complex with new technologies and evolving regulations, the field is full of suitable problems for operations research related problems. As our project had several goals with the overarching theme of reconstructing the merit order curve from historical data and using this reconstruction in support of decision-making for market participants, the project proved to be difficult yet rewarding. Our results mostly focused on the reconstructions of the merit order curves while the downstream problem of optimal bid placement under uncertainty was left for future work due to perceived complexity and time constraints. Herein we include some discussion on future research and open problems.

As discussed in Section 4.4, the binning methodology chosen early in the project to discretize the problem proved to be problematic in the validation phase of bid injection. The binning methodology has been used in literature by Ziel and Steinert (2016), who refer to bins as classes with the main difference being the class bound computation. In our model the limits were approximated by an even distribution of bid counts among bins and general readability, while Ziel and Steinert (2016) computed the limits by approximating equal volumes of trading activity in MW between the price classes. This could be researched further as our choice in binning methodology worked well for the prediction model while suffering in the price injection validation phase. Further research could focus on the method of computing bin bounds as well as the optimal amount of bins to ensure feasible computation time without drastic overfitting while capturing as much of the order curve as reasonable.

The bid injection problem could be studied further with a multitude of different directions in terms of model validation or novel optimization solutions. In terms of validation, an interesting avenue is to use game-theoretic based validation of results in simple example cases. This could bring additional insights into the strategic behaviour of the market participants as researched by Puiu and Hauser (2022) and

Abate et al. (2025). As for actually optimizing bidding behaviour, methods such as stochastic MILP (Herding et al., 2023) or reinforcement learning (Zhu et al., 2023) could be interesting novel solutions to implement into Finnish energy markets.

8 Conclusions

This project studied and reconstructed Finnish mFRR up- and down-capacity market merit order curves using historical accepted bids and forecast variables. The results show that the market has changed significantly, especially in early 2026. The XGBoost model combined with our solution to constructing curves from price bin forecasts was able to produce plausible bid curve forecasts. A key limitation of the study is that only accepted bids are publicly available, meaning that the model aims to predict curves only up to the clearing price. Future work should focus on improving the bid curve representation, for example through optimizing amount of bins and bin bounds. Further research could also extend the framework toward optimal bidding under uncertainty, where reconstructed merit order curves are combined with decision-making models for market participants.

9 Self Assessment

1) How closely did the actual implementation of the project follow the initial project plan? Were there any major departures and, if so, what? The actual implementation followed the initial project plan to a reasonable extent, especially regarding the overall modelling approach and market data analysis. However, there were some departures from the original plan. The most significant issue was that studying the injection problem had not been planned and was considered relatively late in the project.

2) In what regard was the project successful? The construction of curves based on bins was successful and supported by relevant literature. This provided a structured approach for analysing the history data as well as constructing the curves. Additionally, we were able to interpret local and global effects of feature values of the XGBoost model via the SHAP analysis. Finally, despite the complexity of the topic area, every member of the team was able to contribute to the project.

3) In what regard was it less so? The project was less successful in addressing the injection problem, since there was not enough time to research it properly.

4) What could have been done better, in hindsight? Project management could have been improved in some areas and some practicalities could have been agreed on better. Some practicalities and responsibilities could have been agreed upon more clearly at the beginning of the project.

The initial project problem could perhaps have been scoped slightly more narrowly. The open scope was valuable because it allowed the group to explore independent ideas and approach the topic creatively. However, a somewhat more clearly defined starting point might have helped the team in the early phases of the project.

The course is overall very well organized and from a students perspective it was clear what the completion of the course requires. One small improvement point could be that students would have more time to decide on which project topic they decide on.

References

Arega Getaneh Abate, Dogan Keles, Salim Hassi, Xiufeng Liu, and Xiao-Bing Zhang. Bidding strategies for energy storage players in 100% renewable electricity market: A game-theoretical approach, 2025. URL https://arxiv.org/abs/2509.26568.

Anonymous. Dataset of spot price predictions, 2026. Proprietary dataset obtained under a confidentiality agreement. The source and specific model architecture are confidential.

Dena Azarang and Carl Edling. Machine learning-based prediction and key drivers of mFRR activations: A Swedish balancing market study. Bachelor's/master's thesis, Uppsala University, 2025. URL http://www.diva-portal.org/smash/get/diva2:1987054/FULLTEXT01.pdf. Accessed: March 2, 2026.

Tianqi Chen, Tong He, Michael Benesty, Vadim Khotilovich, Yuan Tang, Hyunsu Cho, Kailong Chen, Rory Mitchell, Ignacio Cano, Tianyi Zhou, Mu Li, Junyuan Xie, Min Lin, Yifeng Geng, Yutian Li, Jiaming Yuan, and David Cortes. *xgboost: Extreme Gradient Boosting*, 2026. URL https://github.com/dmlc/xgboost. R package version 3.3.0.0.

ENTSO-E. ENTSO-E Transparency Platform. https://transparency.entsoe.eu/, 2015. Accessed: March 2, 2026.

European Commission. Commission Regulation (EU) No 543/2013 of 14 June 2013 on submission and publication of data in electricity markets and amending Annex I to Regulation (EC) No 714/2009 of the European Parliament and of the Council. Official Journal of the European Union, L 163, 2013. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32013R0543.

Fingrid. Ehdot ja edellytykset manuaalisen taajuuden palautusreservin (mFRR) toimittajalle, 2024. URL https://www.fingrid.fi/globalassets/dokumentit/fi/sahkomarkkinat/reservit/liite-1-mfrr-ehdot-ja-edellytykset-voimaantulo-12.11.2024.pdf. Accessed: February 9, 2026.

Fingrid. Avoin data -palvelu, n.d. URL https://https://data.fingrid.fi/. Accessed: February 9, 2026.

Fingrid Oyj. Reservimarkkinapäivä 2025 – Esitysmateriaalit. PDF document, 2025a. URL https://www.fingrid.fi/globalassets/dokumentit/fi/sahkomarkkinat/reservit/reservimarkkinapaiva-2025-esitysmateriaalit.pdf. Accessed: March 3, 2026.

Fingrid Oyj. Sähköjärjestelmän reservit, 2025b. URL https://www.fingrid.fi/sahkomarkkinat/reservit/sahkojarjestelman-reservit. Accessed: February 9, 2026.

Fingrid Oyj. mFRR, manuaalinen taajuuden palautusreservi, 2025c. URL https://www.fingrid.fi/sahkomarkkinat/reservit/reservituotteet-ja-markkinoille-osallistuminen/

mfrr-manuaalinen-taajuuden-palautusreservi/. Accessed: February 9, 2026.

Fingrid Oyj. Ehdot ja edellytykset manuaalisen taajuuden palautusreservin (mFRR) toimittajalle. PDF document, 2025d. URL https://www.fingrid.fi/globalassets/dokumentit/fi/sahkomarkkinat/reservit/liite-1-mfrr-ehdot-ja-edellytykset-15-min-mtu-voimaantulo-19.3.2025.pdf. Accessed: February 9, 2026.

Fingrid Oyj. Sähköjärjestelmän reservit, 2026a. URL https://www.fingrid.fi/sahkomarkkinat/reservit/sahkojarjestelman-reservit. Accessed: March 3, 2026.

Fingrid Oyj. Johdanto sähkömarkkinoihin, 2026b. URL https://www.fingrid.fi/sahkomarkkinat/markkinoiden-yhtenaisyys/johdanto-sahkomarkkinoihin/#sahkomarkkinoiden-aikajana. Accessed: March 3, 2026.

Fingrid Oyj. mFRR – manuaalinen taajuuden palautusreservi, 2026c. URL https://www.fingrid.fi/sahkomarkkinat/reservit/reservituotteet-ja-markkinoille-osallistuminen/mfrr-manuaalinen-taajuuden-palautusreservi/. Accessed: March 3, 2026.

Robert Herding, Emma Ross, Wayne R. Jones, Vassilis M. Charitopoulos, and Lazaros G. Papageorgiou. Stochastic programming approach for optimal day-ahead market bidding curves of a microgrid. *Applied Energy*, 336:120847, 2023. ISSN 0306-2619. doi:https://doi.org/10.1016/j.apenergy.2023.120847.

Guillaume Koechlin, Filippo Bovera, and Piercesare Secchi. Forecasting electricity spot market merit-order curves with functional time series modeling. In *2025 21st International Conference on the European Energy Market (EEM)*, pages 1–9, 2025. doi:https://doi.org/10.1109/EEM64765.2025.11050354.

Stephan Kolassa and Wolfgang Schütz. Advantages of the MAD/mean ratio over the MAPE. *Foresight: The International Journal of Applied Forecasting*, (6):40–43, 2007. URL https://EconPapers.repec.org/RePEc:for:ijafaa:y:2007:i:6:p:40-43.

Sami Lieskoski, Ossi Koskinen, Jessica Tuuf, and Margareta Björklund-Sänkiaho. A review of the current status of energy storage in Finland and future development prospects. *Journal of Energy Storage*, 93:112327, 2024. doi:https://doi.org/10.1016/j.est.2024.112327.

Scott Lundberg and Su-In Lee. A unified approach to interpreting model predictions, 2017. URL https://arxiv.org/abs/1705.07874.

Jens Lundgren. Market liberalization and market integration: Essays on the Nordic electricity market. (852), 2012. URL https://www.diva-portal.org/smash/get/diva2:570731/FULLTEXT01.pdf. Doctoral dissertation.

Lesia Mitridati and Pierre Pinson. A Bayesian inference approach to unveil supply curves in electricity markets. *IEEE Transactions on Power Systems*, 33(3):2610–2620, 2018. doi:https://doi.org/10.1109/TPWRS.2017.2757980.

Nord Pool. Price formation in the day-ahead market, 2025. URL https://www.nordpoolgroup.com/en/the-power-market/Day-ahead-market/Price-formation/. Accessed: September 17, 2025.

Nord Pool. The power market, 2026. URL https://www.nordpoolgroup.com/en/the-power-market/. Accessed: March 3, 2026.

M. Pavlík, F. Kurimský, and K. Ševc. Renewable energy and price stability: An analysis of volatility and market shifts in the European electricity sector (2015–2025). *Applied Sciences*, 15(12):6397, 2025. doi:https://doi.org/10.3390/app15126397.

Ioan Alexandru Puiu and Raphael Andreas Hauser. A fundamental game theoretic model and approximate global Nash equilibria computation for European spot power markets. 2022. doi:https://doi.org/10.48550/arXiv.2208.14164.

Rasul Satymov, Dmitrii Bogdanov, Tansu Galimova, and Christian Breyer. Energy and industry transition to carbon-neutrality in nordic conditions via local renewable sources, electrification, sector coupling, and power-to-X. *Energy*, 319:134888, 2025. ISSN 0360-5442. doi:https://doi.org/10.1016/j.energy.2025.134888.

Mizue Shimomura, Alexander Ryota Keeley, Ken'ichi Matsumoto, Kenta Tanaka, and Shunsuke Managi. Beyond the merit order effect: Impact of the rapid expansion of renewable energy on electricity market price. *Renewable and Sustainable Energy Reviews*, 189:114037, 2024. ISSN 1364-0321. doi:https://doi.org/10.1016/j.rser.2023.114037.

Ville Sihvonen. *System-level Flexibility Resources in a Decarbonized Finnish Energy System: Demand and Availability*. PhD thesis, Lappeenranta–Lahti University of Technology LUT, Lappeenranta, Finland, 2025. URL https://lutpub.lut.fi/bitstream/handle/10024/169240/Ville%20Sihvonen_A4.pdf. Doctoral dissertation.

Oskar Våle, Shiliang Zhang, Sabita Maharjan, and Gro Klæboe. Exploring the interpretability of forecasting models for energy balancing market. *Artificial Intelligence Science and Engineering*, 1(4):295–306, 2025. ISSN 2097-5104. doi:https://doi.org/10.23919/AISE.2025.000020.

Ziqing Zhu, Ze Hu, Ka Wing Chan, Siqi Bu, Bin Zhou, and Shiwei Xia. Reinforcement learning in deregulated energy market: A comprehensive review. *Applied Energy*, 329:120212, 2023. ISSN 0306-2619. doi:https://doi.org/10.1016/j.apenergy.2022.120212.

Florian Ziel and Rick Steinert. Electricity price forecasting using sale and purchase curves: The X-model. *Energy Economics*, 59:435–454, 2016. ISSN 0140-9883. doi:https://doi.org/10.1016/j.eneco.2016.08.008.