Aalto University
School of Science

Master's Programme in Mathematics and Operations Research

Flow-Based Market Coupling in the Nordic Electricity Market: Evaluating Market Dynamics and Machine Learning Methods

Joel Niemi

Master's Thesis
2026

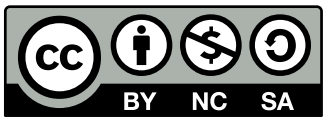

Aalto University
School of Science

Author Joel Niemi

Title Flow-Based Market Coupling in the Nordic Electricity Market: Evaluating Market Dynamics and Machine Learning Methods

Degree programme Mathematics and Operations Research

Major Systems and Operations Research

Supervisor Prof. Ahti Salo

Advisor D.Sc. (Tech) Jari Järvi

Collaborative partner UPM Energy

Date 25.5.2026 **Number of pages** 72 **Language** English

Abstract

This thesis investigates the transition from Net Transfer Capacity (NTC) calculation method to flow-based market coupling (FBMC) in the Nordic electricity market, which was implemented in October 2024. Specifically, this research encompasses two primary objectives, including an empirical analysis of the shifting market dynamics post-FBMC launch, and the development of a machine learning framework to forecast essential flow-based capacity parameters, with a specific focus on the Finnish price area. The empirical analysis reveals that FBMC has fundamentally changed price dynamics between Finland and the SE3 price area, decreasing price convergence, shifting the price premium to SE3, and increasing the frequency of counter-intuitive flows. Crucially, the research identifies a single internal Finnish transmission line that is the primary bottleneck constraining the Finnish grid. Because the Remaining Available Margin (RAM) of this internal element strongly correlates with Finnish SPOT price spikes, accurate forecasting is necessary. Consequently, eXtreme Gradient Boosting (XGBoost) and NeuralProphet models were developed to forecast the week-ahead hourly RAM. Model evaluations demonstrate that a recursive XGBoost approach achieves the highest overall forecast accuracy, while NeuralProphet outperforms XGBoost at capturing the low-RAM extreme events that drive transmission congestion and consequently price spreads. Interpretability analysis confirms that Finnish wind generation and electricity consumption are the primary fundamentals driving RAM variance. Ultimately, forecasting RAM and thus transmission constraints provides significant strategic value for electricity trading, allowing market participants to better forecast price spreads and optimize decision-making in the FBMC environment.

Keywords Nordic Electricity Market, Flow-Based Market Coupling, Remaining Available Margin, Machine Learning, Critical Network Element, Congestion

Aalto-yliopisto
Perustieteiden
korkeakoulu

Tekijä Joel Niemi

Työn nimi Flow-Based Market Coupling pohjoismaisessa sähkömarkkinassa: markkinadynamiikkojen sekä koneoppimismenetelmien arviointi

Koulutusohjelma Mathematics and Operations Research

Pääaine Systems and Operations Research

Työn valvoja Prof. Ahti Salo

Työn ohjaaja TkT Jari Järvi

Yhteistyötaho UPM Energia

Päivämäärä 25.5.2026 **Sivumäärä** 72 **Kieli** englanti

Tiivistelmä

Tässä diplomityössä tutkitaan Pohjoismaiden sähkömarkkinoiden siirtymistä nettosiirtokapasiteetin (NTC) laskentamenetelmästä flow-based market coupling -menetelmään (FBMC), joka otettiin käyttöön lokakuussa 2024. Tutkimuksella on kaksi pääasiallista tavoitetta: empiirinen analyysi markkinadynamiikan muutoksista FBMC:n käyttöönoton jälkeen sekä koneoppimiskehyksen kehittäminen keskeisen FBMC-parametrin ennustamiseksi, painottaen erityisesti Suomen hinta-aluetta. Empiirinen analyysi paljastaa, että FBMC on muuttanut perustavanlaatuisesti hintadynamiikkaa Suomen ja SE3:n hinta-alueen välillä. FBMC on vähentänyt hintojen konvergoitumista, siirtänyt hintapreemion SE3:een ja lisännyt epäintuitiivisen sähkönsiirron yleisyyttä. Keskeistä on, että tutkimuksessa tunnistetaan yksi ainoa Suomen sisäinen siirtolinja, joka on Suomen verkkoa rajoittava pullonkaula. Koska tämän sisäisen elementin Remaining Available Margin (RAM) korreloi voimakkaasti Suomen SPOT-hintapiikkien kanssa, sen tarkka ennustaminen on välttämätöntä. Tämän vuoksi diplomityössä kehitetään eXtreme Gradient Boosting (XGBoost) ja NeuralProphet -mallit ennustamaan seuraavan viikon RAM-parametria tuntitasolla. Mallien arviointi osoittaa, että rekursiivisella XGBoost-menetelmällä saadaan paras kokonaisennustetarkkuus, kun taas NeuralProphet suoriutuu XGBoostia paremmin matalien RAM-piikkien, ja siten hintaeroja aiheuttavien tilanteiden ennustamisessa. Mallien sertifiointi vahvistaa, että Suomen tuulivoimatuotanto ja sähkönkulutus ovat tärkeimmät tekijät, jotka vaikuttavat RAM-vaihteluun. Isossa kuvassa RAM:n ja siten siirtokapasiteetin rajoitusten ennustaminen tarjoaa merkittävää strategista arvoa sähkökauppaan, sillä se voi parantaa markkinaosapuolien hintaennustetarkkuutta ja optimoida päätöksentekoa FBMC-ympäristössä.

Avainsanat Pohjoismainen sähkömarkkina, Flow-Based Market Coupling, Remaining Available Margin, koneoppiminen, kriittinen verkkoelementti, siirtoverkon pullonkaula

Preface

With deep gratitude, I want to thank the nurses, doctors, and Ville Nurminen at HUS, without whom none of this work would have been possible. I am also grateful for my partner, family, and friends, who have been supporting me along my journey.

I want to thank my employer, UPM Energy, for providing an opportunity to conduct my thesis on such an interesting and relevant topic. Specifically, I want to thank Jari Järvi for providing advisory throughout the writing process. Furthermore, I want to thank my supervisor Ahti Salo, as well as all teaching staff at Aalto University.

Helsinki, 5 May 2026

Joel Niemi

Use of AI

Large language models, mostly Microsoft Copilot and Google Gemini, were used in this thesis to proofread text and refactor isolated sentences. Gemini was also used to assist in model development, specifically in code debugging and the implementation of some Python libraries.

Contents

Abbreviations

AAC	Already Allocated Capacity
ACF	Autocorrelation Function
aFRR	Automatic Frequency Restoration Reserve
AMR	Adjustment for Minimum RAM
AR	Auto-regression
ATC	Available Transfer Capacities
BC	Base Case
BRP	Balance Responsible Parties
BSP	Balancing Service Providers
CC	Congestion Costs
CET	Central European Time
CHP	Combined Heat-and-Power
CNE	Critical Network Elements
CNEC	Critical Network Element and Contingency
CR	Congestion Rent
CS	Consumer Surplus
CWE	Central-Western Europe
DA	Day-Ahead Market
EPAD	Electricity Price Area Differentials
EPR	External Parallel Runs
EU	European Union
EUPHEMIA	Pan-European Hybrid Electricity Market Integration Algorithm
FAV	Final Adjustment Value
FB	Flow-Based
FBMC	Flow-Based Market Coupling
FCR	Frequency Containment Reserve
FFT	Fast Fourier Transform
FRM	Flow Reliability Margin
GCT	Gate Closure Time
JAO	Joint Allocation Office
KDE	Kernel Density Estimation
MAE	Mean Absolute Error
mFRR	Manual Frequency Restoration Reserve
MIQP	Mixed-Integer Quadratic Program
ML	Machine Learning
MSE	Mean Squared Error
NP	Net Position
NTC	Net Transfer Capacity method
PACF	Partial Autocorrelation Function
PPA	Power Purchase Agreements
PS	Producer Surplus
PTDF	Power Transfer Distribution Factor

RA	Remedial Actions
RAM	Remaining Available Margin
RCC	Regional Coordination Centre
RES	Renewable Energy Sources
RMSE	Root Mean Square Error
SD	Standard Deviation
SHAP	SHapley Additive exPlanations
SRMC	Short-Run Marginal Cost
SW	Social Welfare
SYS	System Price
TSO	Transmission System Operators
WRMSE	Weighted Root Mean Square Error
XGBoost	eXtreme Gradient Boosting

1 Introduction

The Nordic electricity market is structured to enable regional market coupling between price areas. The market operates under a liberalized framework that facilitates cross-border trade and competition among electricity consumers, producers, and traders. With a mainly emission-free energy mix, driven by hydro, wind, and nuclear power, the Nordic market effectively integrates renewable energy sources while ensuring system stability and competitive pricing.

The Nordic electricity market relies on Nord Pool, a power exchange managing physical electricity trade in the region. In the Day-Ahead market (DA), Nord Pool receives bids from suppliers and consumers to trade electricity in the next day, and clears the market to determine optimal electricity flows and the corresponding area-specific prices. The market clearing is constrained by transfer capacity constraints, which are translated from the physical grid to a mathematical form using capacity calculation. Until October 2024, capacity calculation in the Nordics was performed with the Net Transfer Capacity method (NTC), which is a relatively simplified approach that lacks the ability to model the actual physical electricity flow in the grid.

Flow-based (FB) capacity calculation method was developed to address the shortcomings of the NTC method, and has been implemented in the Nordics since October 2024 after a few years of parallel testing. While the External Parallel Runs (EPR) have provided some data on the behaviour of the method, at the time of writing there is only 17 months of data on how the FB implementation has affected actual prices and physical flows within the Nordic region. By reviewing the literature, it is evident that there exists a gap in research on the impacts of FB implementation in the Nordics.

FB is fundamentally a different approach to capacity calculation, as it models the physical grid through Critical Network Elements (CNE). For each CNE, specific parameters define the transfer capacity constraints used in market clearing, thus affecting physical flows and electricity prices. Being able to forecast electricity prices accurately can often provide value for electricity market participants, such as producers or major industrial consumers, whose revenue can depend on the price of electricity. Forecasting methods in electricity markets are often based on fundamental models, although machine learning (ML) methods have gained popularity in recent years. Given the recency of FB in the Nordics, little to no open research has been conducted on the performance of ML methods in the context, making it an intriguing and relevant topic of research. Accordingly, the two main objectives of this thesis are defined to address the identified research gaps:

1. Empirical analysis of the impacts of flow-based market coupling (FBMC) on the Nordic electricity market dynamics, with a focus on Finland.

2. Development and evaluation of a machine learning framework for forecasting flow-based capacity parameters.

While this thesis introduces the Nordic electricity market broadly, the empirical analysis and ML framework are mostly restricted to the Finnish price area. Additionally,

the scope is restricted strictly to the day-ahead (DA) market. Furthermore, rather than forecasting electricity prices directly, the ML framework specifically targets a flow-based capacity parameter of an identified internal Finnish grid bottleneck.

To achieve the objectives, it is necessary to thoroughly understand the fundamentals of the Nordic electricity market, which are presented in Chapter 2. Additionally, this section discusses the old NTC method and its drawbacks. Chapter 3 introduces the concept of FBMC and the theory behind it. The section also covers the fundamental FB parameters, Power Transfer Distribution Factor (PTDF) and Remaining Available Margin (RAM), and how they affect electricity price formation. Chapter 4 addresses the first objective by discussing the identified empirical findings 17 months after the implementation of FBMC in the Nordics. The most important finding is that there is only one CNE constraining the electricity flow within the borders of Finland. Chapter 5 focuses on the second objective by introducing a ML framework with three distinct methods for forecasting the week-ahead hourly RAM of this particular Finnish element. The results and interpretations of the model forecasts are presented in Chapter 6, while Chapter 7 discusses the implications for electricity trading, model limitations, and future outlook.

2 Fundamentals of the Nordic Electricity Market

The Nordic wholesale electricity market is a deregulated and unbundled market, the main objective of which is to balance the power system efficiently. The market consists of both financial and physical markets, from which the latter can be further separated into the day-ahead, intraday, and balancing markets. Within the context of this thesis, the day-ahead market is the most relevant, however it is also important to understand other markets and the incentivizes of market participants. [1]

2.1 Market Structure

Electricity markets are structurally complex due to the physical properties of electricity. Unlike other commodities, electricity has to be consumed at the same time as it is produced, which means that electricity grids have to continuously maintain balance within supply and demand. Constant balance is essential for frequency stability and the secure operation of a power system [2]. Although electricity storage technologies could provide flexibility to power system balancing, their economic feasibility remains an issue [3]. In the Nordic region, the balancing responsibility is shared between transmission system operators (TSO) and balance responsible parties (BRP) [1].

2.1.1 Market Participants

The Nordic electricity market consists of suppliers, consumers, and regulating authorities. The market is structured to ensure that electricity is generated, transmitted, and consumed efficiently, while maintaining system stability and enabling fair competition. Different market participants play distinct roles to facilitate both the physical delivery and financial trading of electricity, contributing to the overall functioning and reliability of the power system.

TSOs are national operators, for instance Fingrid in Finland and Statnett in Norway, with the obligation to maintain the balance in power systems in real time [1]. To account for possible system imbalances, TSOs use balancing products as elaborated in Section 2.2.4. According to the EU Target Model [2], TSOs are also responsible for the facilitation of cross-border balancing. The Nordic Regional Coordination Centre (RCC) aims to increase coordination between the four TSOs in the Nordic region, thus potentially increasing market efficiency. RCC is also responsible for the regional Flow-Based capacity calculation. [4].

BRPs are commercial actors, such as generators and consumers, who are required to ensure the balance of their scheduled individual net positions (supply - demand). In case the actual net position deviates from what it is designated to be, imbalances occur, forcing TSOs to use balancing products. However, the financial costs that occur from utilizing the balancing products fall on the BRPs through what is called imbalance settlement. Given that the financial risk of having imbalances is on BRPs, this system incentives them to actively forecast, monitor, and adjust their electricity flows to avoid imbalances in the system. [2]

Generators supply the market with electricity, which will eventually flow to the end consumer, whether it be a major industrial consumer or a typical household. In the short run, generators aim to maximize their profits by bidding in physical markets. Trivially, high electricity prices tend to favor generators, especially those with low marginal cost generation capacity.

Consumers form the demand side for the market. Major industrial consumers might often participate directly to the bidding, and thus they can benefit from low electricity prices. The demand of small consumers, such as households or small private businesses, is aggregated to the market through retail companies, who procure electricity from the market and sell supply contracts to the end consumers. Additionally, some speculators have been historically participating in the financial markets, where they aim to benefit from price movements, while not actually participating in any physical deliveries. [5, 6]

2.1.2 Nordic Market Characteristics

The Nordic electricity market region consists of Denmark, Finland, Norway, and Sweden. The countries are further divided into smaller price areas, except for Finland, which acts as a single uniform price area. In total, there are 12 price areas, with five of them in Norway, four in Sweden, two in Denmark, and one in Finland. Each price area, or bidding zone, has their own electricity price in both the physical and financial markets. Although the market promotes cross-border trade, there are often some price differences between areas. This is because the transfer capacities between areas constrain the flow of electricity. In theory, if all capacity constraints were neglected, prices in all areas would converge to the same price. Known as the system price (SYS), this is often used as the reference price in standard financial contracts.

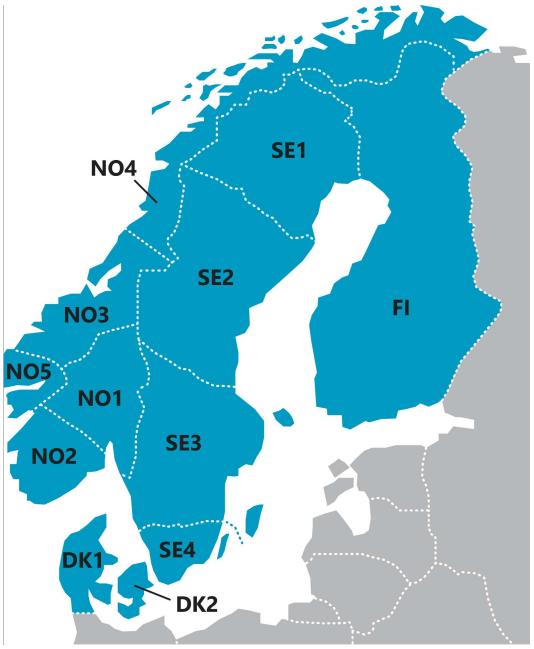

Figure 1: Price areas in the Nordic region. Modified from [7].

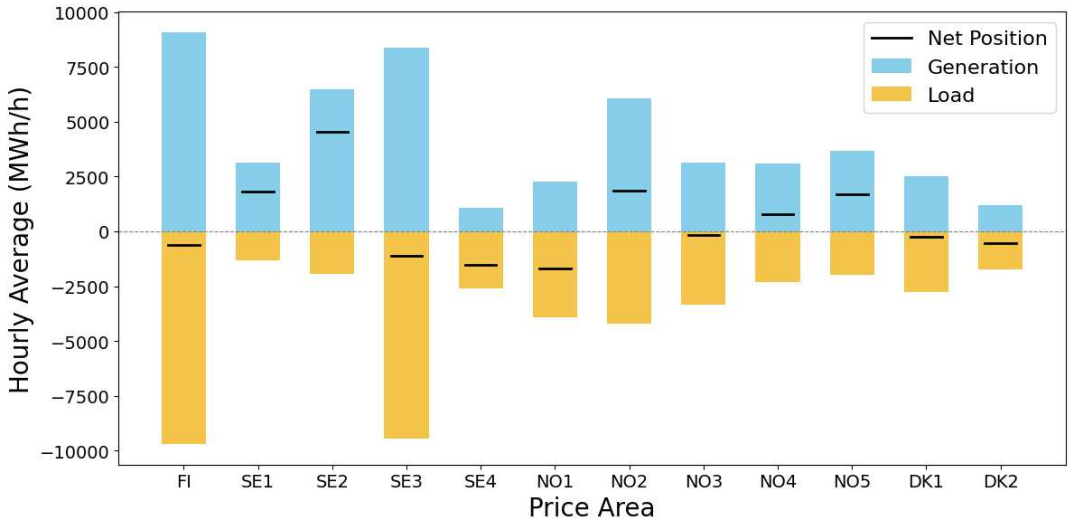

Figure 2: Average hourly generation, load, and net position (in MW) of each price area in 2025. Data retrieved from [8].

In a liberalized market, it is not feasible for all areas to produce electricity just for their own use. Thus, some areas constantly produce more electricity than needed and export surplus electricity to other areas, while others rely on imports. This relationship between generation and load is often referred to as the Net Position (NP) of an area. There are significant differences between the average hourly NPs of the price areas (Figure 2). Areas with substantial industrial demand, such as FI and SE3, tend to import electricity to account for the demand, although they are the only areas that have nuclear generation. Given that approximately half of Nordic electricity generation in 2025 came from hydro [9], areas with hydro-favorable topological characteristics, such as NO2 and SE2, have positioned themselves as net exporters, as seen in Figure 2.

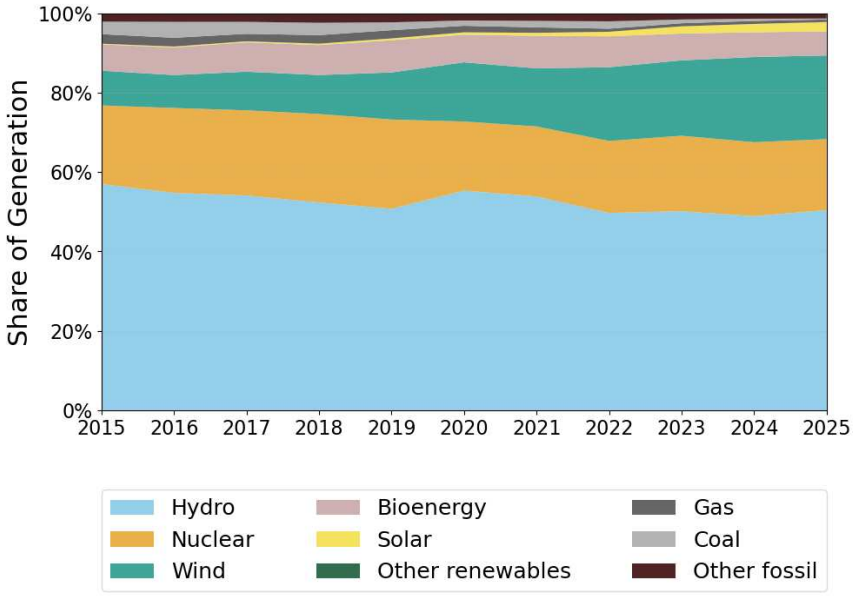

Figure 3: Nordic electricity mix in 2015-2025. Total generation has increased from 402 to 447 TWh. Data retrieved from [9].

In addition to hydro power, the Nordic electricity mix consists mainly of nuclear and wind power, with up to 96.9% of the generation being carbon-free (Figure 3). However, the increased amount of intermittent renewable energy sources (RES), such as wind and solar, has strengthened the relationship between weather conditions and electricity price. Natural intermittency in generation leads to short-term supply shocks, thus increasing price volatility and extreme prices. This dynamic can be seen in daily System SPOT price, its 1-year centered rolling average price, and the standard deviation (SD) of a rolling 30-day window (Figure 4). The chart clearly shows the effects of the energy crisis in 2021-2023 in extreme spikes. In Europe, the crisis accelerated the adoption of renewables, which can be seen as increased volatility in the time series, even though the yearly average price has returned to similar levels as before the crisis. In some price areas, most notably in Finland, the hourly SPOT price has even hit zero or negative values (Figure 5). Negative prices might occur when there is oversupply from low marginal cost generation, and producers have to pay consumers to use that electricity. Often, shutting down production for a short time period can be costlier than running the unit with zero or negative profit. [10]

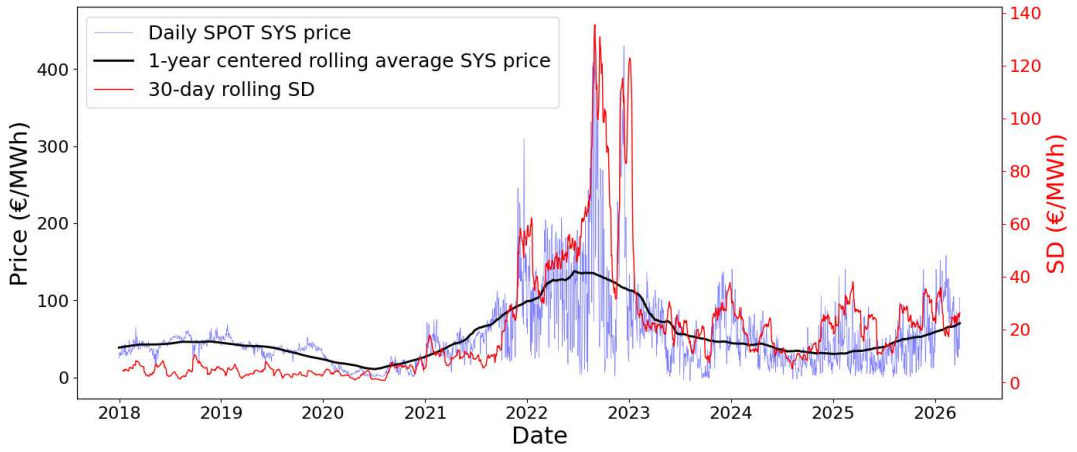

Figure 4: Daily SPOT System price, its 1-year centered rolling average, and the 30-day rolling standard deviation (as of April 2026). Data retrieved from [11].

.

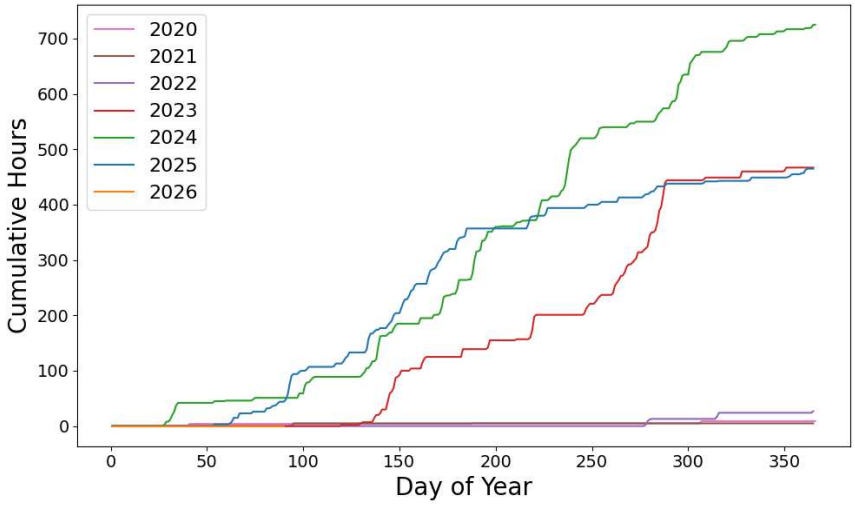

Figure 5: Cumulative count of hours with zero or negative prices in Finland (as of April 2026). Data retrieved from [8].

2.1.3 Market Coupling

Market coupling is a core concept of the EU Target Model, designed to create an integrated European wholesale electricity market by coordinating auctions across different bidding zones. By enabling cross-zonal electricity transmission, it drives price convergence across the Nordic region by ensuring that electricity is generated by the most cost-efficient sources, regardless of bidding zone borders. Its primary objective is to determine zonal electricity prices and allocate cross-border transmission capacity. Market coupling is facilitated by an algorithm aggregating all bids from regional power exchanges, and solving area-specific net positions and prices. This

algorithm is introduced in Section 2.3.2. Market coupling is applied across various electricity trading timeframes, which are introduced in Section 2.2. [2]

2.2 Electricity Trading

Stable and secure electricity systems require functioning trading frameworks and platforms to work. Within the context, it is essential to understand the fundamentals of microeconomics in electricity markets, which are introduced in Section 2.2.1. Electricity trading can be further separated into financial and physical trading, with the latter being more relevant to this thesis. Physical trading can further be divided into the day-ahead (DA), intraday, and balancing markets, each with their own responsibilities as discussed in Sections 2.2.2 - 2.2.4. Financial markets are also briefly introduced in Section 2.2.5. Figure 6 visualizes an overview of different electricity markets in the Nordics and how they locate on a timeline, with varying gate closure times (GCT).

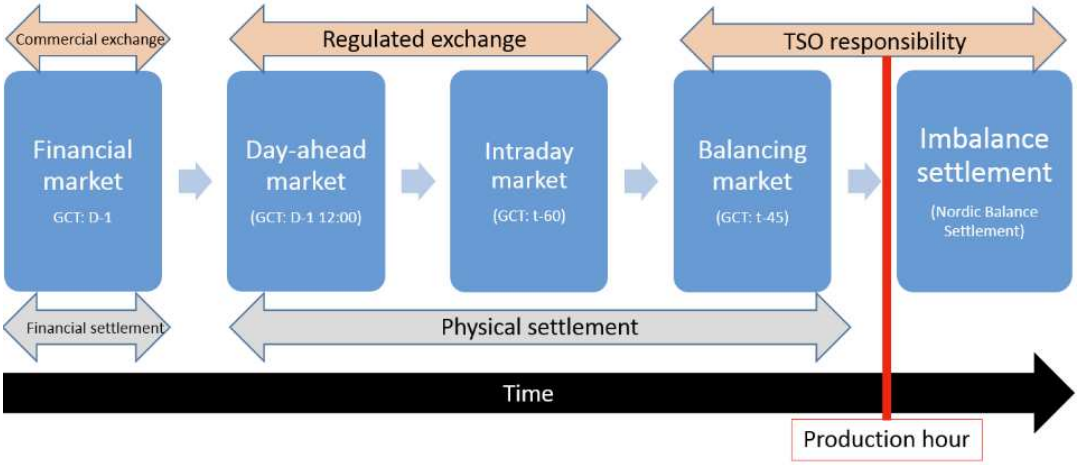

Figure 6: Electricity trading timeline (D = Day of delivery, t = Time of delivery in minutes). Figure from [12].

2.2.1 Electricity Market Economics

The price of electricity on a micro level is determined by the supply and demand. In essence, the market equilibrium at the intersection of supply and demand curves defines the quantity and price of electricity traded. [13] Each market participant in the Nordics can place supply or demand bids which can be aggregated into area-specific supply and demand curves. Figure 7 illustrates SYS bid curves, which essentially represent the aggregated supply and demand in the Nordic region.

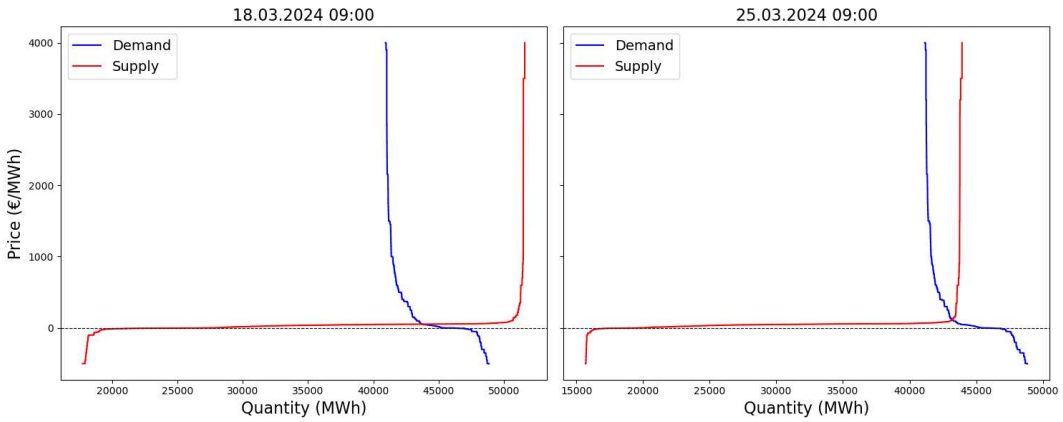

Figure 7: Supply and demand curves on two consecutive Mondays in the Nordic electricity market, with a differing RES output. Data retrieved from [11].

As Figure 7 shows, the supply and demand curves of electricity tend to have unique characteristics. The slope of the demand curve at a specific quantity reflects the marginal benefit for the aggregated consumer, though distinct consumer groups, such as households and industrial consumers, have unique demand curves due to their differing utility functions for electricity. In the short run, the aggregated base demand is highly inelastic, as a trivially constant amount of electricity is needed to support basic necessities, such as heating, lighting, and industrial processes. Beyond the base demand level, the demand becomes increasingly elastic, as consumers can begin to adjust their electricity usage for non-essential consumption. The shapes of Nordic electricity demand curves often exhibit consistent patterns, though the base demand level varies significantly due to factors, such as air temperature, weekday, and seasonal variations.

On the other hand, the slope of the aggregated supply curve reflects the marginal cost of producing electricity at some quantity. The supply curve is shaped by the merit order, which ranks power plants by their short-run marginal cost (SRMC), dispatching the cheapest sources first, such as RES with near-zero SRMC, followed by more expensive sources, such as fossils [14]. As the amount of RES production has increased, the supply curve can be highly elastic at certain quantities, where low-cost RES dominate the merit order. However, the intermittency in weather conditions, such as wind availability, can lead to sudden supply shocks that shift the supply curve, as shown in Figure 7, which demonstrates the aggregated supply and demand curves of two consecutive Mondays in March 2024. The dates are selected to demonstrate a short-run supply shock arising from a difference in the amount of wind at the given time. On the first graph, the market equilibrium is at a point where the price elasticity of supply is high, as low-cost RES dominate the merit order. On the second graph, due to the supply shock reducing RES availability, the equilibrium arises at a point where more expensive generation sources begin to be dispatched, making the supply curve increasingly inelastic. Thus, even small demand shocks, for instance caused by errors in air temperature forecasts, or further supply shocks, can lead to extremely high prices in a RES-led market.

2.2.2 Day-Ahead Market

The day-ahead market (DA) has historically been the most important marketplace in terms of volume and liquidity. In DA, electricity generators can place supply bids and consumers can place demand bids for each hour in the next day. In the Nordics, DA is primarily facilitated by Nord Pool, which essentially aggregates the bids into supply and demand curves, and uses them to determine the hourly equilibrium prices for the next day for each price zone. In 2020, EPEX SPOT launched a competing DA market in the Nordics, ending the monopoly position of Nord Pool. In 2022, Nord Pool facilitated 696 TWh of trade in the Nordic and Baltic DA markets, holding the largest market share, while EPEX SPOT facilitated 47 TWh in the Nordic region alone [11, 15]. The DA auction is arranged every day and closes at 12.00 CET, and the DA prices for each hour in the next day are published shortly after, usually around 12.45 CET. At 10.00 CET, two hours before the auction closes, market participants receive information regarding the transmission capacities, enabling them to make more informed bidding decisions. [1, 16]

After the bids have been submitted, the market is cleared across Europe using a common price coupling algorithm EUPHEMIA [17]. The algorithm considers all bids across most European price areas and the cross-border transfer capacities, and gives area-specific prices and net positions. Thus, calculating the capacities is a crucial part of the price determination, as further discussed in Section 2.3. Considering this thesis, DA stands out as the most relevant marketplace, since FBMC is at the time of writing implemented for the capacity calculation only in DA.

2.2.3 Intraday Market

Trading on the Nordic intraday market opens after DA closes, and continues often until one hour before the delivery begins. Intraday allows market participants to adjust their DA positions to address errors in their demand and supply forecasts. In the Nordics, the intraday market is operated by Nord Pool and EPEX SPOT. Intraday offers products with 15-, 30-, and 60-minute delivery periods, thus helping in balancing the near real-time supply and demand of the system.

Unlike DA, which operates with a predetermined schedule for market clearing, the intraday market operates continuously. Intraday prices are determined on a first-come, first-served basis through ongoing trade clearing, similar to the operational structure of a stock exchange. Intraday prices often vary from DA prices, and this price spread can mainly be explained by wind and demand forecasts errors. The increasing amount of RES in the system has made forecasting more difficult, thus attracting market participants to trade on the intraday markets to avoid imbalance costs. While trading volumes in intraday have historically been smaller than in DA, the intraday volumes have increased in the recent years. [1, 11, 15, 18]

2.2.4 Balancing Market

The primary function of the Nordic balancing market is to ensure supply and demand balance in real-time, thus securing the stable operation of the system. The balancing

market operates after both DA and intraday markets have closed, although in contrast to DA and intraday markets, the balancing market is facilitated by the TSOs instead of a regulated exchange. It can be broadly divided into the reserve market, in which balancing capacity is traded, and the regulating market, which focuses on the actual delivery of electricity. TSOs are responsible for procuring balancing capacity in the reserve market through reserve products, such as Frequency Containment Reserve (FCR), Automatic Frequency Restoration Reserve (aFRR), and Manual Frequency Restoration Reserve (mFRR), primarily differentiated by their activation time. In the simplest case, the supplier to the reserve market is some participant, who is willing to provide on-demand production or consumption capacity to the system. In the regulating market, balancing energy is traded and activated in real-time or near real-time to correct actual imbalances between electricity generation and consumption during the delivery hour. Balancing Service Providers (BSP) can submit bids for up- or down-regulation close to the time of electricity delivery, and TSOs can activate these bids to either increase production or decrease consumption (up-regulation) or vice versa (down-regulation) to maintain a 50 Hz grid frequency. Furthermore, the prices established in the regulating market also serve as the basis for determining the imbalance settlement costs. [1, 19]

2.2.5 Financial Market

In the financial markets, market participants can trade power derivative contracts to hedge against risks associated with price volatility. The fundamental difference between financial and physical markets is that while the electricity traded in physical markets is physically delivered to the buyer, financial contracts are settled with cash payments. In the Nordics, the main platform for power derivative trading is Nasdaq OMX Commodities Europe, providing trading services for both exchange-traded and OTC derivatives. Additionally, forward contracts and Power Purchase Agreements (PPA) can be done bilaterally. Financial contracts are settled against the average realized DA price during the delivery period. System (SYS) price is often used as the reference price in financial contracts, which can be traded up to a decade in advance, though liquidity is generally limited to future contracts expiring within a few years from the time of evaluation. Additionally, area-specific Electricity Price Area Differentials (EPAD) can be used to hedge against price area risk. [20]

2.3 Capacity Calculation and Market Clearing

In a fully liberalized market enabling unlimited cross-border trade, the market price would in theory be the same in all regions. However, supplying electricity from one node or region to another requires physical grid infrastructure, such as transmission lines. In practice, all price spreads between Nordic price areas arise from binding physical grid constraints. When clearing the DA market with EUPHEMIA, the algorithm aims to find a market solution, including area-specific net positions and prices, constrained by the transmission capacities. However, due to the complexity

of the physical grid, simplifications need to be made when translating the grid into a form applicable for EUPHEMIA. This process is referred to as capacity calculation.

2.3.1 Transmission

The transmission infrastructure is a critical component in any electricity market as it enables the physical delivery of electricity across price areas and national borders. The Nordic transmission infrastructure is managed by TSOs, and consists of high-voltage transmission lines, substations, and interconnectors, for instance. The main objective of the transmission system is to transport electricity from generation sites, often located in areas with renewable resources like hydropower in Norway, to consumption centers, such as industrial and urban areas in Finland and Sweden. However, the physical properties of electricity and the finite capacity of transmission lines form constraints that significantly influence market outcomes, including price differences across price areas. [21]

Transmission infrastructure in the Nordic region is particularly important due to the geographical distances between generation and consumption centers. For example, much of the hydro generation capacity is located in the northern parts of Norway and Sweden, while the consumption centers are mostly in the south. This typical power flow from north to south increases congestion on the transmission infrastructure, thus creating bottlenecks and contributing to price differences. These bottlenecks not only affect market efficiency by limiting trade but also raise concerns regarding grid security and even investment decisions.

2.3.2 EUPHEMIA

EUPHEMIA (Pan-European Hybrid Electricity Market Integration Algorithm) serves as the central market clearing mechanism for bids submitted in the Nordic region. Its main objective is to determine the set of executed orders that maximizes total social welfare, which is the sum of consumer surplus, producer surplus, and congestion rent in the region. The objective is subject to physical grid constraints, such as available transfer capacities and ramping restrictions as defined by the TSOs, which are translated to the problem with capacity calculation. The problem can be modelled as a Mixed-Integer Quadratic Program (MIQP), which can be solved using a branch-and-cut approach with sub-problem decomposition. The optimal set of executed orders provides net positions and prices for each area. The mathematical formulations of the problem are introduced in Section 3.1. [17]

2.3.3 Net Transfer Capacity Method

Until October 2024, the capacity calculation in the Nordics was done with the Net Transfer Capacity method (NTC). In the NTC method, the capacities are calculated for bilateral exchange on each price area border. The capacities are given in MW, and are also known as Available Transfer Capacities (ATC). The TSOs define ATC values based on assumptions of the physical flows, and the calculation process is often

relatively non-transparent. Often, the defined ATC values are rather conservative to avoid physical overloadings of transmission lines. Additionally, the ATC value of a cross-border link is independent from others, thus neglecting the actual physical flow in some cases. The lack of information on the physical flow can lead to inefficiencies and unutilized capacity when using the NTC method. To address these shortcomings, FBMC has been implemented in the Central-Western Europe since 2015, and in the Nordics since October 2024. [22, 23]

3 Flow-Based Market Coupling

Flow-Based Market Coupling (FBMC) is a mechanism for coupling electricity markets more efficiently by improving the utilization of the existing grid infrastructure. More specifically, FBMC aims to increase calculated cross-border transmission capacities, thus often leading to relaxed constraints in the market clearing procedure. The capacity calculation with FBMC accounts for the physical flow of electricity and considers constraints set by actual grid elements, thus overcoming some of the limitations of the NTC method. Theoretically, this should lead to increased social welfare (SW) and smaller price spreads. [24]

3.1 Theoretical Background

According to Kirchhoff's laws, electricity generated at a node in area A and transmitted to a node in area B physically flows through all parallel paths in the grid following the path of least resistance, despite of any theoretical area borders. In other words, increased net position (NP) of area A might lead to increased flow between areas B and C. In FBMC, this is modelled using Critical Network Elements (CNE), Power Transfer Distribution Factors (PTDF), and Remaining Available Margin (RAM). [23]

3.1.1 Critical Network Elements

CNEs are transmission infrastructure elements in the grid, such as transmission lines and transformers, that might constrain electricity transmission. For instance, transmission lines between FI and SE1, and between SE1 and SE2, are two distinct CNEs. Furthermore, Critical Network Element and Contingency (CNEC) refers to a CNE that is considered together with a specific contingency or grid outage, such as maintenance. An example of a CNEC would be the FI-SE1 CNE in a scenario, where the SE1-SE2 CNE would have an outage. A specific CNE can occur in multiple CNECs, both as the CNE and as the contingency. CNECs are determined by national TSOs, and the number of CNECs can vary hourly, usually ranging around several hundred elements in the Nordic region. [23, 24, 25]

CNECs are fundamental in defining the grid constraints for the market clearing process in FBMC. Each CNEC has parameters, that are used to numerically form the capacity constraints. These parameters include Power Transfer Distribution Factors (PTDF) and Remaining Available Margin (RAM).

3.1.2 Power Transfer Distribution Factors

PTDFs are sensitivity factors that can be used to characterize the incremental flow through a CNEC that arises from a 1 MW increase in trade between areas A and B. For instance, increased trade between areas A and B might physically flow through a CNEC in area C, in which case the particular CNEC would have a non-zero PTDF value. PTDF values in the Nordic region are defined for each of the 12 price areas as

seen in Figure 1 and CNEC ($n \in \mathcal{N}$, where $\mathcal{N}$ is the set of all CNECs in the Nordic area, and $N = |\mathcal{N}|$ is the total number of CNECs), forming the PTDF matrix:

$$\mathbf{PTDF} = \begin{bmatrix} PTDF_{FI}^{1} & \cdots & PTDF_{DK2}^{1} \\ \vdots & \ddots & \vdots \\ PTDF_{FI}^{N} & \cdots & PTDF_{DK2}^{N} \end{bmatrix},$$

which contains the PTDF values for each CNEC in the Nordic system, with respect to each price area.

PTDFs are typically calculated by considering an injection in a particular area and extraction in a selected slack area. The choice of slack area for the calculations does not have any impact on the final results of the market outcome.

Another way to interpret PTDFs is to visualize the flow through a CNEC. The physical flow is often nonlinear and is linearized for simplicity. The linearized flow is set equal to the physical flow at the common base case (BC) NP of an area, NP_{BC}, and this quantity of flow is referred to as the reference flow F_{ref}. The linearized flow:

$$F_{CNEC} = PTDF_{i}^{n} \times NP_{i} + F_{0},$$

where F_0 is the flow on the CNEC if NP is zero, i is the price area in the set of Nordic price areas $\mathcal{I}$, and PTDF is the slope of the function as illustrated in Figure 8.

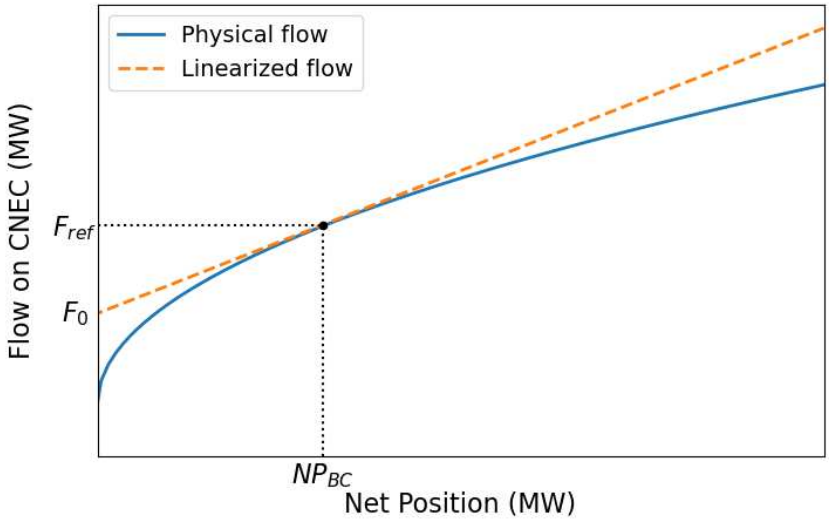

Figure 8: Physical and linearized flow on a CNEC as a function of the net position of an area.

Furthermore, the change in flow through an element n, resulting from a 1 MW increase in exchange from area A to B, can be calculated from the PTDF differential:

$$\Delta F_n = (PTDF_{A}^{n} - PTDF_{B}^{n}).$$

3.1.3 Remaining Available Margin

RAM represents the remaining capacity in MW on some CNEC that can be used for commercial exchanges, specifically within the day-ahead market [23, 25]. It is a fundamental parameter within FBMC, as it together with PTDFs translate physical grid limitations into a numerical form that can be applied by the power exchange.

Like PTDFs, RAMs are determined by the national TSOs for each CNEC in their country. RAM of a CNEC is defined as

$$RAM = F_{max} - F_0 - FRM + RA - FAV - AAC - AMR,$$

where F_{max} is the maximum thermal capacity of the CNEC. F_0 represents the flow through the element when all area net positions are set to zero. Flow Reliability Margin (FRM) is a security margin that covers uncertainties in the capacity calculation process, such as linearization errors or wind forecast deviations. Remedial Actions (RA) account for the additional capacity made available by considering actions, such as re-dispatching generation or topology changes, that can be applied after a contingency. Final Adjustment Value (FAV) allows TSOs to manually adjust the available margin to address risks not captured by the other parameters. Already Allocated Capacity (AAC) represents the capacity reserved for exchanges allocated in previous timeframes, such as long-term transmission rights or bilateral trade. Lastly, AMR indicates the adjustment for minimum RAM by the TSOs. [26] [27]

3.1.4 Flow-Based Domain

PTDFs and RAMs are used to determine the feasible transmission region, or FB domain, for the market clearing [23]. Since the flow through a CNEC is a linear function of area net positions, the FB domain is defined by the inequalities:

$$\sum_{i \in \mathcal{I}} PTDF_i^n \times NP_i \leq RAM_n \quad \forall n \in \mathcal{N},$$

where $NP_i = x_i^S - x_i^D$ is the net position of price area $i \in \mathcal{I}$, n describes a particular CNEC within the set of all CNECs $\mathcal{N}$, and x_i^S, x_i^D represent supply and demand quantities, respectively. This set of inequalities defines an $(I-1)$-dimensional solution space, where $I = |\mathcal{I}|$ is the number of price areas [28]. In the Nordic region, FB increases the number of hourly solution space constraints from 34 (NTC) to roughly 150, thus providing a more accurate representation of grid limitations [29]. This results in a larger solution space, often allowing a better solution, as demonstrated in Figure 9.

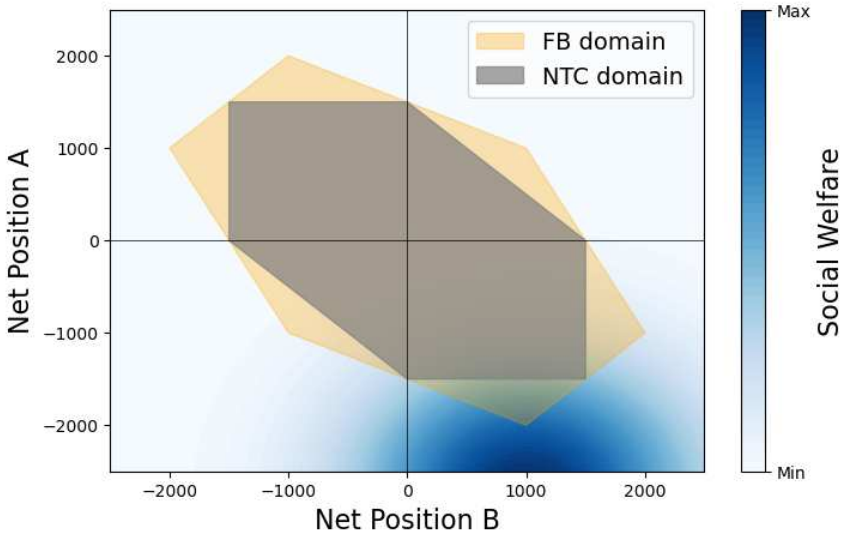

Figure 9: Arbitrary illustration of FB and NTC domains in a two-dimensional system, and how the FB domain allows higher SW market solutions. Author's own visualization.

3.2 Social Welfare and Price Formation

Understanding the impact of FBMC and specifically the RAM and PTDF parameters on electricity pricing requires examining the mathematics behind the price formation, and finding how these parameters fit into it. The primary goal of EUPHEMIA is to maximize the total socio-economic, or social, welfare (SW), while ensuring that the resulting power flows respect the limitations of the grid infrastructure. This section presents the theoretical framework of social welfare maximization and demonstrates how the intersection of physical flow constraints and supply and demand bids mathematically determines both the optimal allocation of resources and the resulting area prices.

3.2.1 Social Welfare

SW in the Nordic electricity market is the sum of consumer surplus (CS), producer surplus (PS), and congestion rent (CR), which is essentially the surplus of the TSOs:

$$SW = CS + PS + CR,$$

where CS, PS are:

$$CS = V_i(x_i^D) - p_i x_i^D$$
$$PS = p_i x_i^S - C_i(x_i^S),$$

where $V_i(\cdot)$ is the utility function for the aggregated consumer of electricity in price area $i \in \mathcal{I}$ for consuming x_i^D amount of electricity, and p_i is the price of one unit. To maximize CS, consumers adjust their consumption to the quantity at which the marginal

utility of consumption equals the marginal cost of consumption. For producers in price area i, $C_i(\cdot)$ is the cost function of producing x_i^S amount of electricity, and the maximum PS is attained when the marginal revenue equals the marginal cost of producing electricity. Thus, at maximum CS and PS,

$$\frac{\partial V_i(x_i^D)}{\partial x_i^D} = \frac{\partial C_i(x_i^S)}{\partial x_i^S} = p_i. \tag{1}$$

CR differs from zero when the trade causes congestion between two price areas in the grid. In case the market is not limited by any transmission constraints, there is no congestion, and CR is zero. When the transmission between price areas i, j is congested, CR is

$$CR_{ij} = (p_j - p_i)Q_{ij},$$

where $p_j - p_i$ is the price spread between areas i and j, and Q_{ij} is the quantity of electricity net transferred from i to j. Although CR is technically a surplus for the TSOs, a non-zero CR causes congestion costs (CC), which reduce the total SW (Figure 10). CC is essentially deadweight loss, indicating market inefficiency.

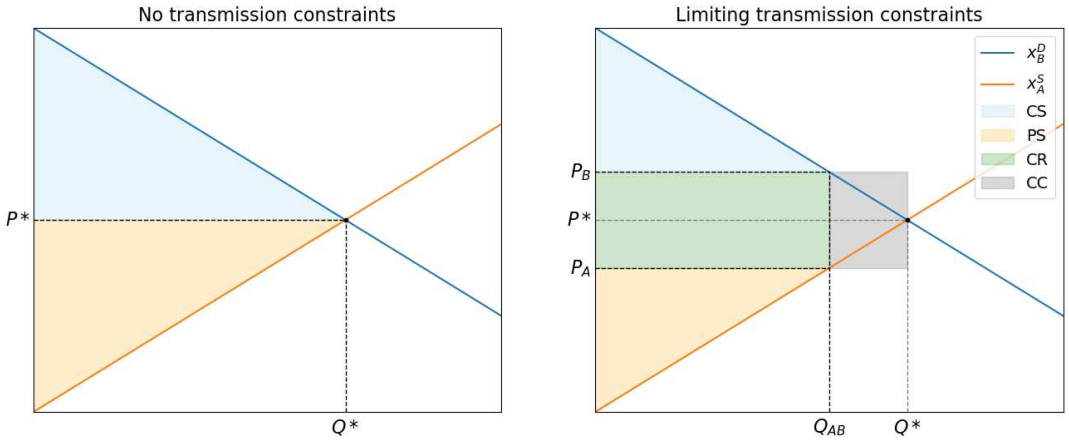

Figure 10: Social welfare formation in a supply-and-demand framework [13]. CS = consumer surplus, PS = producer surplus, CR = congestion rent, CC = congestion cost.

In Figure 10, x_A^S illustrates the net supply curve from surplus electricity generated in price area A, and x_B^D demonstrates the net demand for covering the electricity deficit in price area B. In an unconstrained market, the price at the equilibrium is P^*, which is equal to the SYS in the Nordic market. In this case, SW is maximized. However, if the equilibrium quantity Q^* exceeds the capacity of the transmission between price areas A and B, the prices are determined at Q_{AB}, causing price divergence, CR, and CC. Trivially, higher transmission capacities lead to higher SW and smaller price spreads, and vice versa.

3.2.2 Mathematics of Price Formation

Given that CR is essentially a financial transfer instead of a direct component of SW that should be maximized, the objective function can be formally simplified to

$$\sum_i CS_i + PS_i = \sum_i V(x_i^D) - p_i x_i^D + p_i x_i^S - C_i(x_i^S), \tag{2}$$

where x_i^S, x_i^D are the supply and demand quantities of price area $i \in \mathcal{I}$, with $x_i^S - x_i^D$ being the net position of the area. These are also the decision variables of the objective function. Given that the supply and demand within the market region must match continuously, the power balance constraint can be formed,

$$\sum_i (x_i^S - x_i^D) = 0. \tag{3}$$

With the power balance constraint, the price terms in Equation (2) cancel out, and the objective function can be further simplified to

$$\sum_i V(x_i^D) - C_i(x_i^S).$$

Additionally, the transmission constraints in FBMC are defined with the PTDFs and RAMs for each CNEC $n \in \mathcal{N}$,

$$\sum_i PTDF_i^n (x_i^S - x_i^D) \leq RAM_n, \quad \forall n. \tag{4}$$

With Equations (3) - (4), a simplified SW optimization problem can be formulated as

$$\begin{aligned}
\max. \quad & \sum_i V(x_i^D) - C_i(x_i^S) \\
\text{s.t.} \quad & \sum_i PTDF_i^n (x_i^S - x_i^D) \leq RAM_n \quad \forall n && (5) \\
& \sum_i (x_i^S - x_i^D) = 0 && (6) \\
& x_i^S, x_i^D \geq 0.
\end{aligned}$$

The prices can be derived from the dual variables of the transmission and power balance constraints. Lagrangian of the optimization problem is

$$L = \sum_i V_i(x_i^D) - C_i(x_i^S) + \sum_n \mu_n \left[RAM_n - \sum_i PTDF_i^n (x_i^S - x_i^D) \right] + \lambda \sum_i (x_i^S - x_i^D).$$

From the first-order conditions with respect to x_i^D, x_i^S and Equation (1), price at area i can be derived as

$$\frac{\partial V_i(x_i^D)}{\partial x_i^D} = \frac{\partial C_i(x_i^S)}{\partial x_i^S} = p_i = \lambda - \sum_n \mu_n PTDF_i^n,$$

where λ is the dual variable of the power balance constraint in Equation (6), representing the marginal cost of meeting demand in an unconstrained system, or effectively SYS price. The shadow price, μ_n is the dual variable of the transmission constraint in Equation (5) for CNEC n. According to complementary slackness in linear programming, μ_n is strictly non-zero when the transmission constraint is binding, or in other words, flow on the CNEC equals its RAM.

Consequently, the price spread between two areas is determined by the shadow prices of the binding elements and the difference in their PTDFs,

$$p_j - p_i = \sum_n \mu_n (PTDF_i^n - PTDF_j^n).$$

This result indicates that in FBMC, price spreads are a direct result of physical bottlenecks in the grid. The magnitude of the price spread is proportional to the CNEC congestion severity, as well as the relative difference in price area sensitivity to the flow through the CNEC.

If no network elements are congested, all shadow prices are zero, and the market clears at a single uniform price λ across the Nordic region. While this unconstrained outcome is identical to the NTC method, the mechanism for price separation differs fundamentally. Under FBMC, price spreads are driven by the shadow prices of specific physical grid elements and their PTDFs, rather than simple cross-zonal transmission capacities. Consequently, price spreads can occur in FBMC even without binding cross-zonal transmission lines.

4 Finnish Market Dynamics Under FBMC

While FBMC is implemented across the whole Nordic system, this thesis focuses explicitly on the Finnish price area. Due to its isolated geographical location relative to other price areas, Finland shares major transmission capacity with only three other price areas: SE1 (approximately 2000 MW), SE3 (1200 MW), NO4 (100 MW), EE (Estonia, 1000 MW) [11]. Although EE is not a part of the Nordic system, it is an essential link for electricity flow from the Nordics towards Central Europe. Due to their relevance for the FI SPOT price, only SE1, SE3, and SYS are included in this analysis.

Empirical findings on FBMC impacts on FI SPOT price dynamics and spreads with neighboring price areas are introduced in Section 4.1. Also the changes in Finnish NP and transmission flows between FI and SE3 are discussed in Section 4.2. Finally, Section 4.3, introduces an empirical finding on a Finnish grid bottleneck CNEC and justifies the importance of forecasting the RAM of this particular CNEC.

4.1 SPOT Prices

In theory, FBMC should generally lead to lower prices and increased price convergence, given the larger solution space. At the time of writing, there is 17 months of data since the FBMC launch on 30.10.2024, providing insight into its effects. The most obvious variable of interest is the direct SPOT price of electricity in Finland. It is worth noting that the data reflects structural changes in the Nordic electricity system and the 2022 energy crisis.

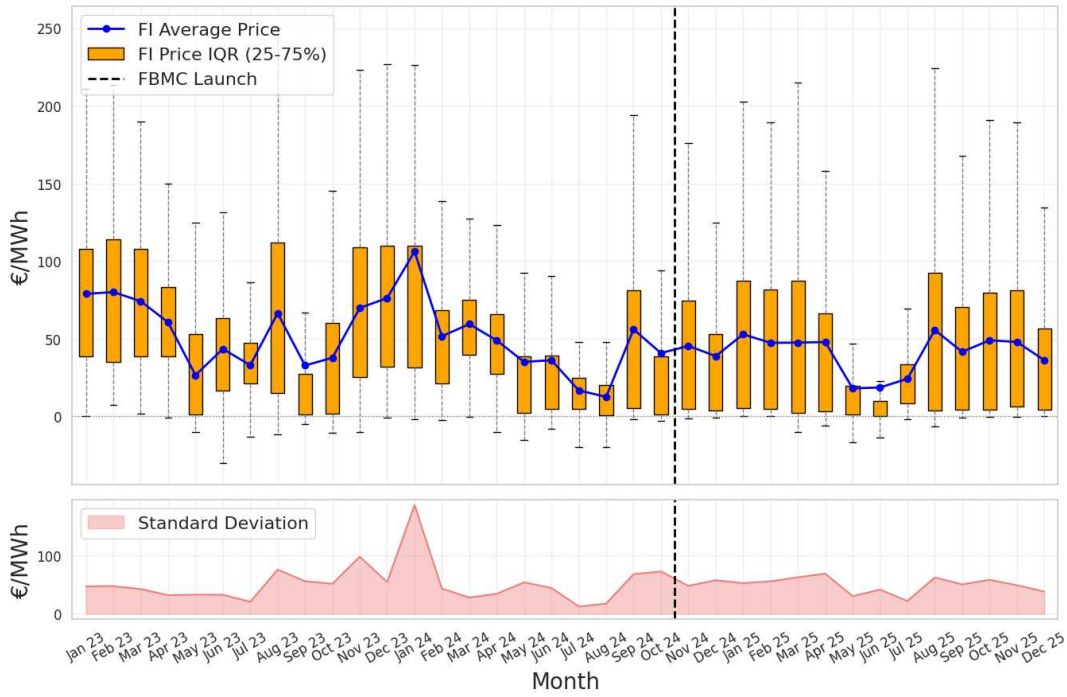

Figure 11: FI average price and price distribution by month (2023-2025). Extreme outliers not displayed. Data retrieved from [8].

Figure 11 displays slightly lower bottom quartiles for FI prices since the launch of FBMC, however this is likely due to the higher-than-average hydrological balance in the Nordics in the first half of 2025 (H1/2023, H1/2024 are well below historical average) [8]. Consequently, it is difficult to make incontestable assumptions based on this observation alone. Given that the hydro balance impacts the broader Nordic market, it is essential to evaluate the price spreads between Finland and its neighboring price areas, SE1 and SE3.

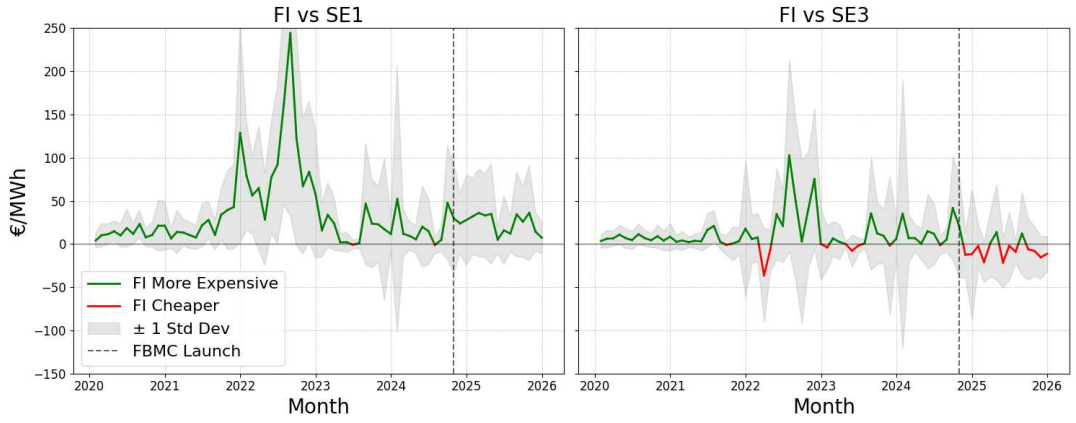

Figure 12: Monthly average SPOT price spreads between FI and SE1, SE3 (2020-2025). Positive value indicates higher FI. Data retrieved from [8].

Historically, FI has been the more expensive area when comparing to SE1, SE3, on average (Figure 12). Price spreads between FI and SE1 seem to be at similar levels pre- and post-FBMC. In November 2025, the transmission capacity between FI and SE1 increased significantly with the commissioning of Aurora Line [30], thus possibly reducing the spread in the future. However, this is excluded from the analysis due to a small sample size.

Additionally, Figure 12 shows that since the launch of FBMC, the price spread between FI and SE3 has turned negative, indicating cheaper FI in comparison to SE3. However, monthly average price spread is not the most robust estimator. Looking into the number of hours with price premiums for both areas provides more detailed insight into the FI-SE3 price dynamics pre- and post-FBMC.

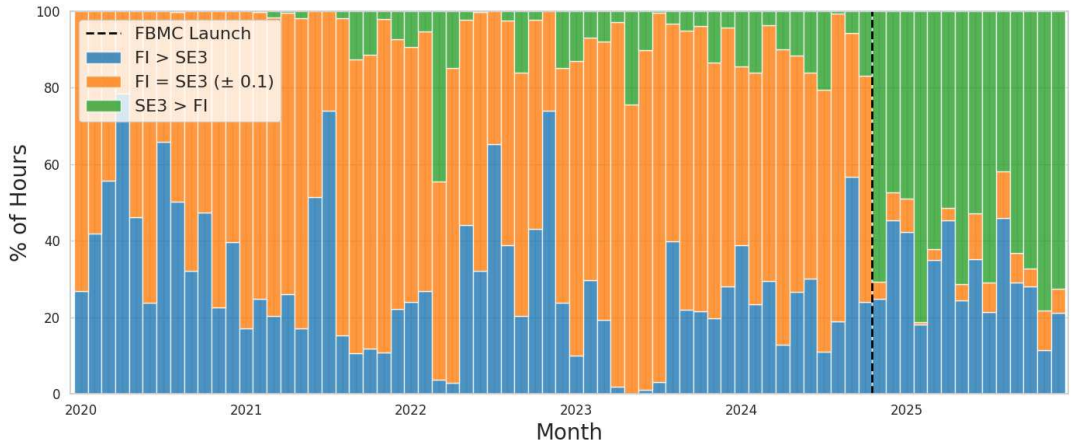

Figure 13: Percentage of hourly price premiums between FI and SE3 by month. Data retrieved from [8].

Figure 13 suggests that there has been a significant change in FI-SE3 price dynamics due to FBMC. Since the launch, the hours where SE3 is more expensive than FI (SE3 price premium) has become the standard, often occurring over 50% of the time. However, these hours are not reduced from the hours with FI price premium, but rather the hours where the prices are equal (by ± 0.1 €/MWh margin). The percentage of equal-prices hours has decreased to often well below 10%, which is a significant change compared to the pre-FBMC market. The suddenness and the permanence of the change at the time of the FBMC launch indicates that it is a consequence of FBMC. To further examine the root causes of the change, it is important to analyze the actual transmission flows between FI and SE3.

4.2 Transmission Flows

When discussing transmission flows, it is important to distinguish the concepts of physical and commercial flows. Physical flow means the actual flow of electricity, following Kirchhoff's laws and the physical properties of the grid. On the other hand,

commercial flow represents the computational market perspective on the electricity flow, often assuming direct electricity flow from the producer to the consumer. [31]

Out of the two, inspecting physical flow is more relevant, since FBMC explicitly accounts for the physical flow and grid constraints. Examining the hourly net position of FI provides insight on the overall import and export dynamics of the price area.

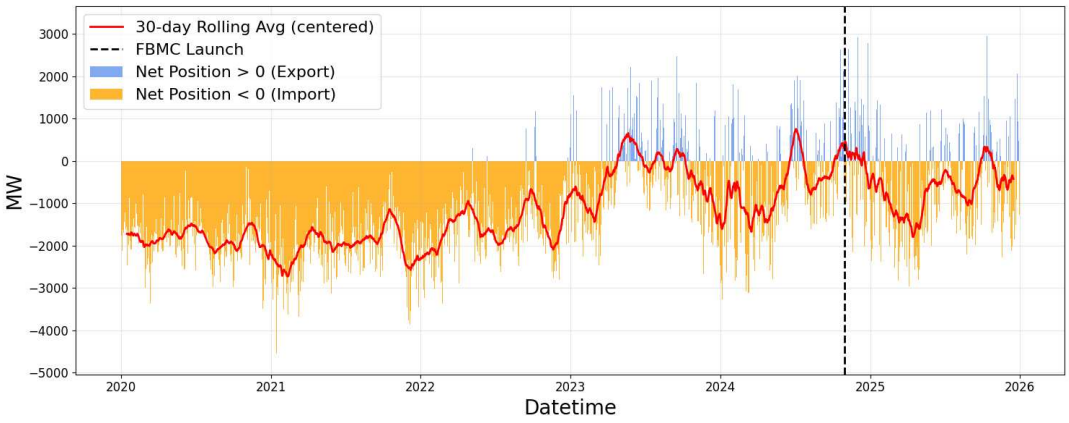

Figure 14: Hourly physical flow net position of Finland. Data retrieved from [8].

Figure 14 shows that FI has begun exporting more electricity post-crisis, compared to pre-crisis. However, this is due to non-FBMC factors, such as increased wind power generation capacity, and the effects of FBMC on FI net position seem negligible.

Electricity transmission between Finland and SE3 is facilitated through two HVDC lines, Fennoskan 1 and 2. The maximum capacity is 1200 MW, but is often limited for maintenance reasons, for instance. Visualizing the Kernel Density Estimation (KDE) distribution of the hourly flows enables examining the flow through a single interconnector (in this case, Fenno-Skan 1 + 2).

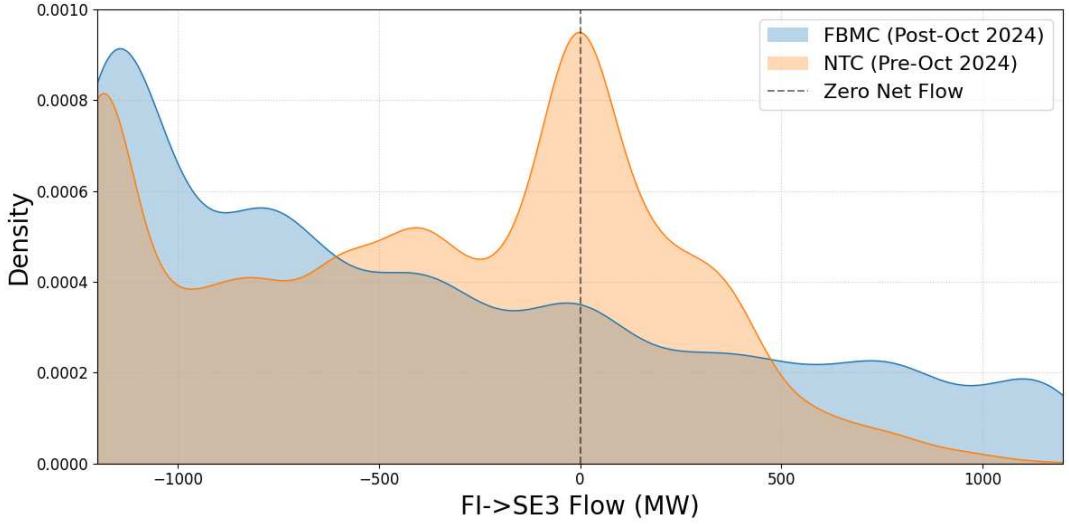

Figure 15: KDE distribution of hourly FI->SE3 flow (2023-2025). Positive value on x-axis indicates export from FI. Data retrieved from [8].

The flow distribution in Figure 15 seems to have changed slightly since the FBMC launch. Compared to pre-FBMC, the distribution has shifted from zero towards both extremes (-1200, 1200 MW). This suggests higher utilization rate of the existing transmission line, indicating that FBMC is successfully unlocking latent grid capacity. More specifically, the congestion frequency (hourly flow above 98% of the capacity) of the FI-SE3 transmission has increased from 26.9% to 28.6%. However, this is contradicting to the price convergence hypothesis of FBMC, since congestion leads to shadow prices and consequently price spreads.

One finding on the price dynamics between FI and SE3 is that the occurrence of counterintuitive flows has become more common (Figure 16). Counterintuitive flow means that electricity is exported from a higher-priced area to a cheaper area during a specific hour. A counterintuitive flow on a transmission line occurs when it relieves congestion in some other line, increasing the total SW in the Nordic system. In other words, the benefit of relieving congestion on a binding CNEC is greater than the loss of welfare caused by the counterintuitive flow.

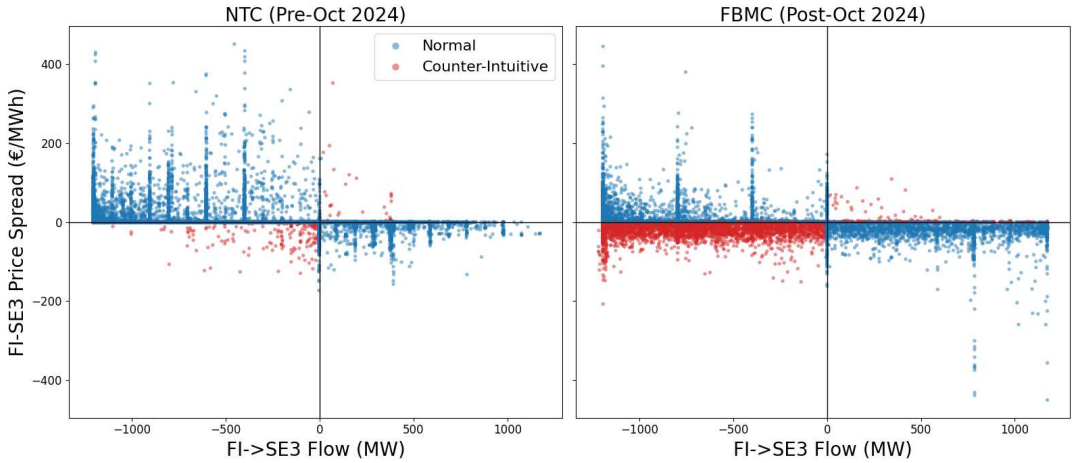

Figure 16: counterintuitive flows (hourly) between FI and SE3. Data retrieved from [8].

4.3 Bottleneck of the Finnish Grid

To further understand how FBMC impacts FI prices, it is instructive to inspect the specific CNECs constraining the market solution. This can be achieved using shadow prices, which are published daily by the Joint Allocation Office (JAO) [32]. In the data, there are 229 CNECs governed by Fingrid. Filtering for the CNECs that have had at least one hour with a non-zero shadow price narrows the query result to just ten distinct CNECs. Out of these, nine are essentially cross-border transmission lines with different contingencies. Interestingly, there is exactly one constraining transmission line segment, which is completely within the borders of Finland. This CNEC, specified as "FI_P1_VUOLIJOKI-ALAPITKA_PETAJAVESI-PYSAYSPERA" in the dataset, is located on the easternmost north-south high-voltage transmission line, as highlighted in Figure 17. More specifically, this CNEC refers to a transmission line segment

from Vuolijoki to Alapitkä, under a contingency on a transmission line segment from Petäjävesi to Pysäysperä. In other words, the values of this CNEC represent a scenario where the Petäjävesi-Pysäysperä transmission line segment has an outage, and electricity has to flow through other parallel lines. In this paper, this CNEC is referred to as Vuolijoki CNEC.

Figure 17: Vuolijoki CNEC on the Finnish grid. Figure from [33].

Typically, Finland imports electricity from the generation-heavy Northern Sweden. This electricity is transmitted from the north to the demand-heavy Southern Finland using four main lines, as displayed in Figure 17. When the amount of electricity flowing through the four parallel main lines gets large enough, Vuolijoki CNEC is the first element to get congested, causing a bottleneck in the Finnish system. This

leads to non-zero shadow prices, which in turn are heavily related to price spreads, as demonstrated in Section 3.2.2.

4.3.1 Shadow Price

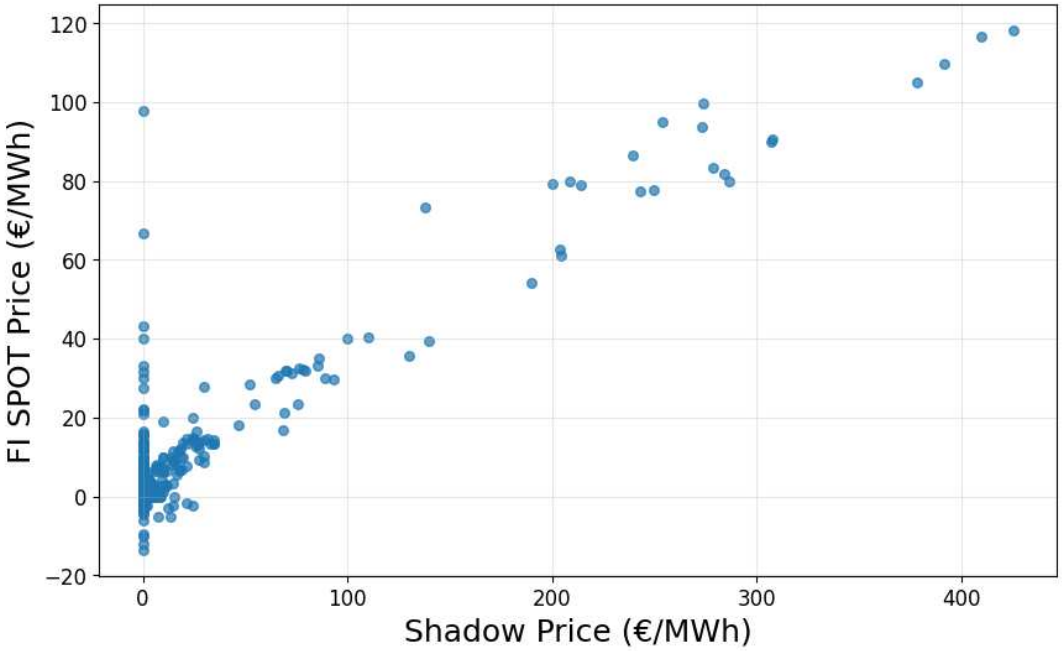

Figure 18: Hourly Vuolijoki CNEC shadow prices against FI SPOT price (Nov 2024 - Dec 2025). Data retrieved from [32].

Figure 18 clearly demonstrates a relationship between the shadow price of Vuolijoki CNEC and the FI SPOT price. The figure also includes the hours where the shadow price is zero. While the SPOT prices can spike even without Vuolijoki CNEC constraining the solution, whenever it is binding, its shadow price has a significantly strong relationship with FI SPOT price. Although Vuolijoki CNEC has been binding for only approximately 1.85% of the hours, it has to be taken into account when forecasting FI SPOT prices.

4.3.2 Power Transfer Distribution Factors

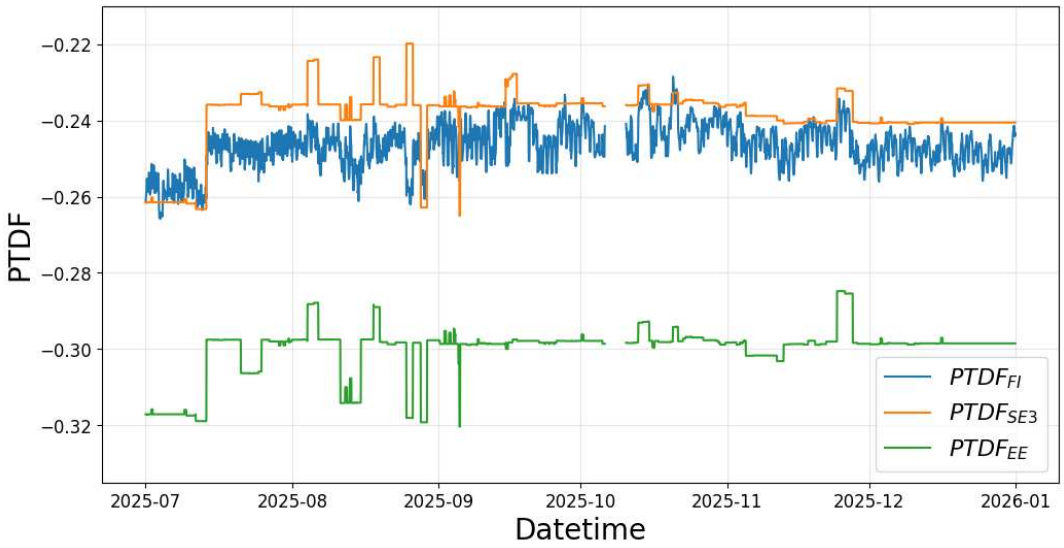

Figure 19: Hourly Vuolijoki CNEC PDTF values (H2/2025). 4 days of data missing in October. Data retrieved from [32].

Vuolijoki CNEC has three distinct PTDFs with significance, one for FI, SE3, EE each. In practice, these PTDF values indicate how much of the electricity injected in that area flows through the Vuolijoki CNEC, as explained in Section 3.1.2. Figure 19 shows that these values are negative, indicating that an injection in any of the three areas would actually decrease the flow through the CNEC.

PTDFs are fundamentally determined by physical grid topology. As Figure 19 shows, PTDF values exhibit stable behavior, remaining near-baseline for extended periods. The step-wise shifts in PTDF values are mostly driven by scheduled line outages or maintenance, instead of market dynamics. Previous research suggests that forecasting PTDF values does not require complex machine learning models, and therefore this thesis focuses on forecasting the RAM. [34]

4.3.3 Remaining Available Margin

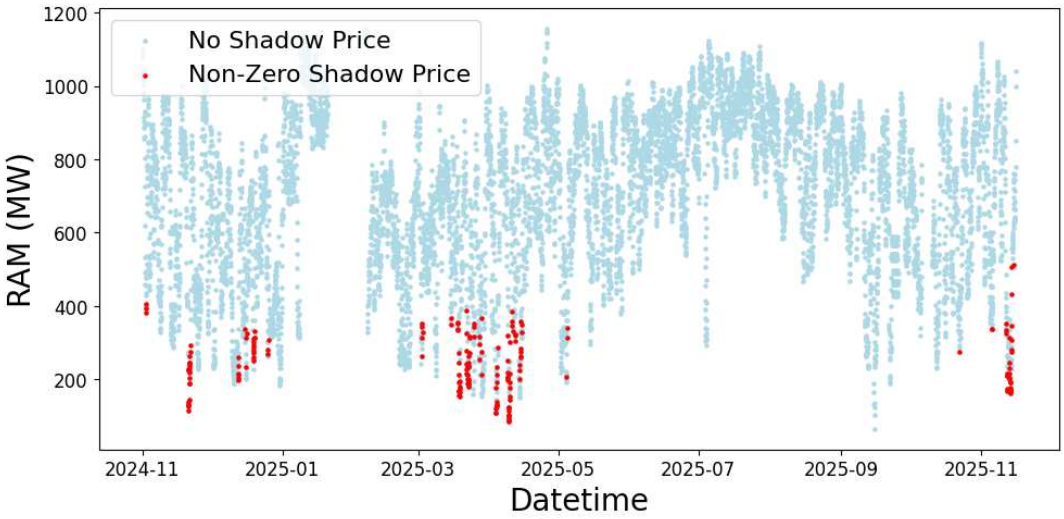

Figure 20: Hourly RAM and shadow price occurrences of Vuolijoki CNEC (12 months from FBMC launch). 14 days of data missing in February. Data retrieved from [32].

The RAM of Vuolijoki CNEC varies mostly between 200 and 1000 MW. As Figure 20 displays, shadow prices on the CNEC occur mostly during hours where RAM is below approximately 400 MW. This indicates that lower RAM increases the likelihood of Vuolijoki CNEC binding the market solution, consequently affecting FI SPOT prices. These results suggest that forecasting RAM accurately could provide real value in improving SPOT price forecasts. To understand the fundamentals explaining RAM, it is necessary to examine the components of RAM as introduced in Section 3.1.3.

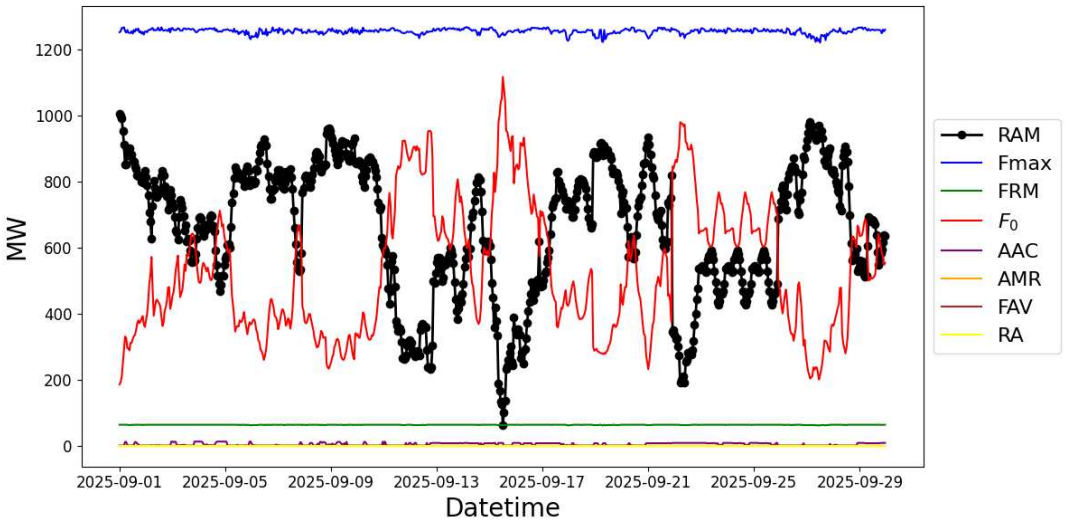

Figure 21: Hourly RAM and its components of Vuolijoki CNEC (September 2025). Data retrieved from [32].

As seen in Figure 21, FRM, AMR, FAV, RA, AMR of Vuolijoki CNEC are relatively insignificant and constant over time. F_{max} is mostly trivially constant, but can exhibit step-wise changes due to maintenance outages. In November 2025, the F_{max} of Vuolijoki CNEC was permanently increased by Fingrid, but the other fundamentals have remained constant.

Essentially the only varying variable driving the RAM of the CNEC in Figure 21 is the zero net position flow, F_0. It demonstrates nearly perfect negative correlation (-0.9985), indicating that a model capable of accurately predicting F_0 would yield a RAM prediction of nearly equivalent precision.

F_0 flow does not strictly necessitate zero cross-border flow between price areas, as long as the sum of exports and imports is zero. Often, there exists non-zero F_0 flow between the areas.

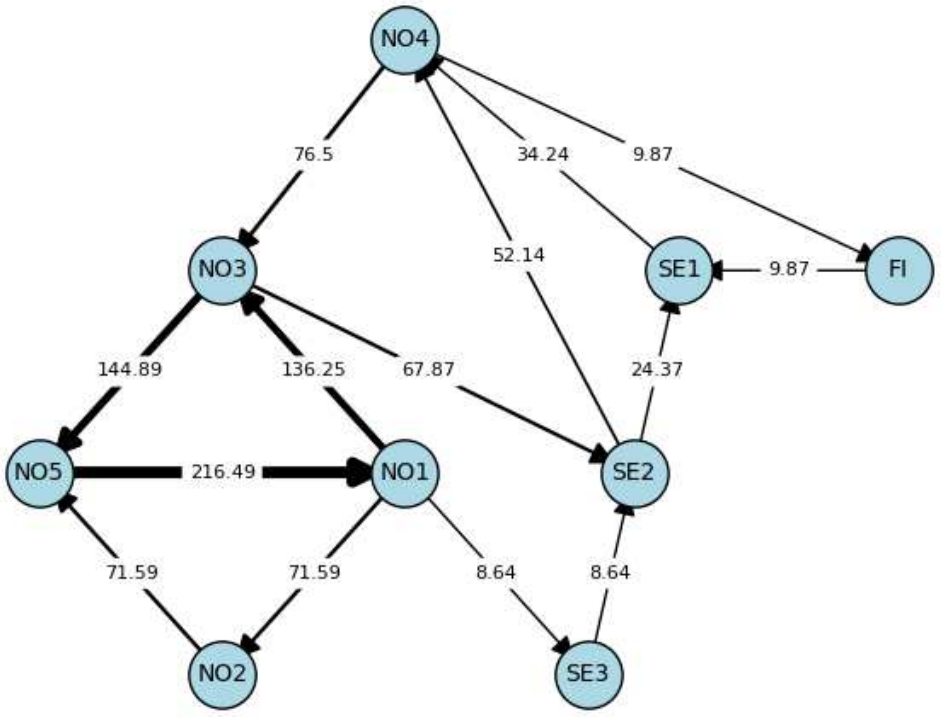

Figure 22: Average F_0 flows in the Nordic system (12 months from FBMC launch). DK1, DK2, and SE4 have zero flows and are excluded from the figure. Data retrieved from [32].

Figure 22 shows a significant F_0 flow loop within Norway. However, the average F_0 flow in and out of Finland is 9.87 MW, which is relatively insignificant. Furthermore, given that the PTDF value for NO4 is zero, the average F_0 flow through the Vuolijoki CNEC resulting from a 1 MW increase in NO4->FI export can be calculated with the PTDF differential:

$$\begin{aligned} \Delta F_{0,Vuolijoki} &= PTDF_{NO4} - PTDF_{FI} \\ &= 0 - (-0.25) = 0.25. \end{aligned}$$

However, since also the PTDF value for SE1 is zero, the F_0 flows through Vuolijoki CNEC occurring from cross-border exchange between FI and SE1, NO4 offset each other. Consequently, this proves that the RAM of Vuolijoki CNEC is not driven by any variables outside of Finland. This finding significantly reduces the number of fundamentals that could be considered as explanatory variables.

To summarize, this empirical analysis justifies the value of forecasting the RAM of the Vuolijoki CNEC. By treating this Finnish bottleneck as a primary use case, ML forecasting methodologies developed in this thesis can potentially be scaled to other CNECs across the Nordic region.

5 Machine Learning Framework for Forecasting RAM

The empirical analysis in Chapter 4 demonstrated the strategic value and technical necessity in forecasting the day-ahead RAM of Vuolijoki CNEC accurately. Furthermore, the analysis proved that the variance in RAM is almost entirely driven by F_0, which was shown to be trivially independent of variables outside of Finland.

This chapter introduces the methodology and development of a machine learning framework designed to forecast the week-ahead RAM of the Vuolijoki CNEC. In Section 5.1, an overview of the data is provided. Section 5.2 discusses the analysis of the feature candidates and their relationships with RAM. In Section 5.3, the ML methods used in this thesis are introduced. Sections 5.4 and 5.5 respectively elaborate the feature engineering process and the model training and hyperparameter optimization processes. Lastly, the framework for model evaluation and interpretability is introduced in Section 5.6.

5.1 Data Description

The dataset used in this thesis is a time series beginning from the FBMC launch on 1.11.2024, spanning to 31.3.2026. The dataset contains realized RAM values from JAO Publication Tool [32], and realized exogenous variable values from Volue [8]. These exogenous variables include Finnish wind, solar, nuclear, combined heat-and-power (CHP), and hydro generation (in MWh/h), as well as hydro balance (MWh), consumption and residual load (MWh/h), and average temperature (°C). The data is in hourly intervals.

In real-world applications, forecasting future RAM requires exogenous predictions, such as wind and consumption forecasts. For a day-ahead forecast, using historical exogenous forecasts as input features would be justified. However, especially wind forecasts exhibit significant errors in forecast horizons longer than few days, thus leading to model inaccuracy in week-ahead forecasts. In order to validate the model with as small external forecast errors as possible, realized data is used in model training and testing phases.

In the dataset, there are 432 missing rows, each containing one hour, in the source data. However, both models used in this thesis, introduced in Section 5.3, are capable of handling the missing data internally, given that the missing rows are present in the time series as NaN values. Since most variables in the dataset exhibit yearly seasonality, only the first 8760 rows, representing a full, non-leap year, of the dataset are used in Chapter 5.2 to not give too much weight to some months.

5.2 Feature Selection

Given the results in Section 4, the list of plausible explanatory features can be narrowed down to Finnish variables. Since RAM is effectively an inverse of the zero net position flow through Vuolijoki CNEC, it is likely that generation and consumption explain at least some of the variance. Furthermore, the location of the generation and consumption is important, since the bottleneck in the CNEC occurs only for electricity

flowing from north to south. However, locational or nodal level of granularity for generation and consumption would require a detailed grid model. The easiest way to increase the granularity of generation is to distinguish it by generation type. The relationships between RAM and exogenous features are discussed in Section 5.2.1. Moreover, the temporal and seasonal properties of RAM are elaborated in Section 5.2.2.

5.2.1 Exogenous Fundamental Features

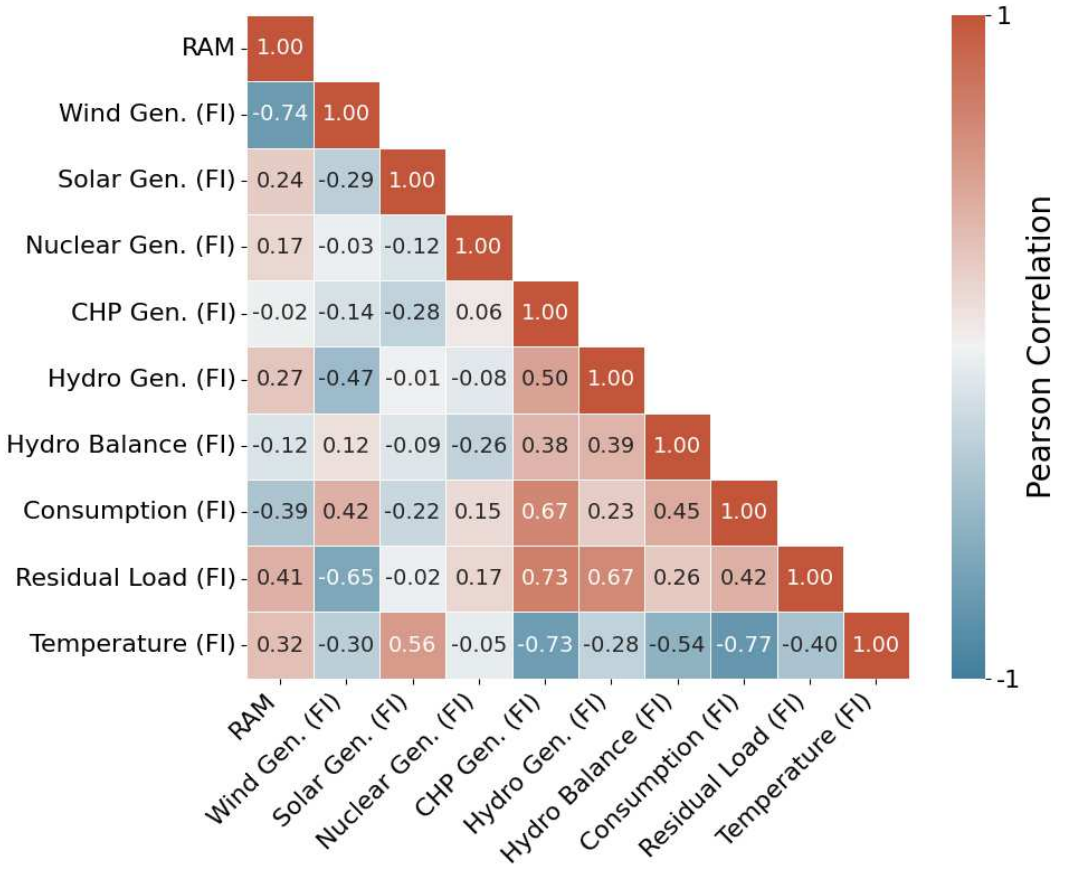

Figure 23: Correlation matrix of RAM and exogenous feature candidates. Data retrieved from [8] [32].

Figure 23 displays the correlation matrix between observed RAM and wind, solar, nuclear, combined heat and power (CHP), and hydro generation (MWh), as well as hydro balance, consumption, residual load (MWh), and temperature ($°C$) values. The matrix shows that wind generation would be the primary feature explaining RAM, with a Pearson correlation of -0.74. Residual load exhibits the second strongest correlation with RAM, however it has significant collinearity with wind, since it is calculated by subtracting variable renewable energy sources from total consumption. Pure consumption has a relatively strong negative relationship with RAM. The

intuition behind these observations is that increased wind generation, typically in the western and northern parts of Finland, directly increases the F_0 flow through Vuolijoki CNEC, consequently reducing RAM. On the other hand, since most of the Finnish consumption is located in the south, decreasing the consumption decreases the need to transmit electricity from the north, thus reducing F_0 and increasing RAM.

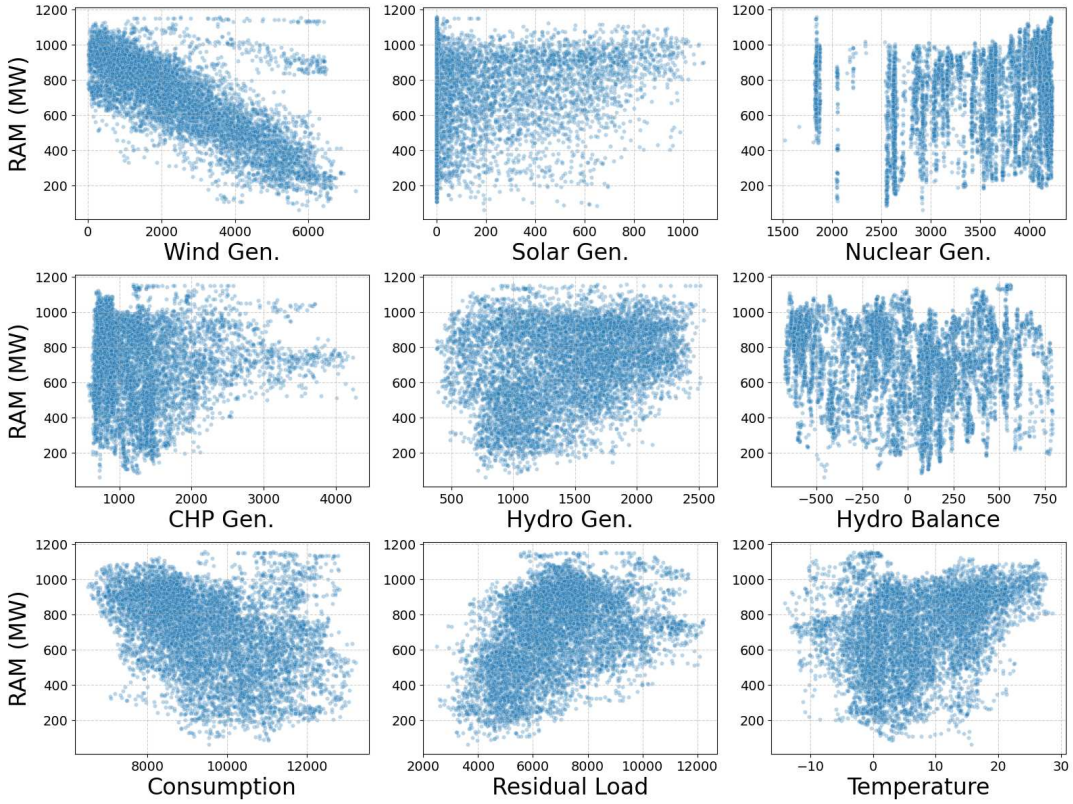

Figure 24: Scatter plots between RAM and relevant Finnish variables. Data retrieved from [8] [32].

Figure 24 confirms the strong relationship between Finnish wind generation and Vuolijoki CNEC RAM. Moreover, the scatters indicate some weaker nonlinear relationships between other variables and RAM.

5.2.2 Temporal and Seasonal Properties

Evaluating temporal and seasonal dependencies in RAM time series is an essential part of methodology selection. While Figure 20 displays clear yearly seasonality in RAM, it is also necessary to inspect the short-term autocorrelation of RAM.

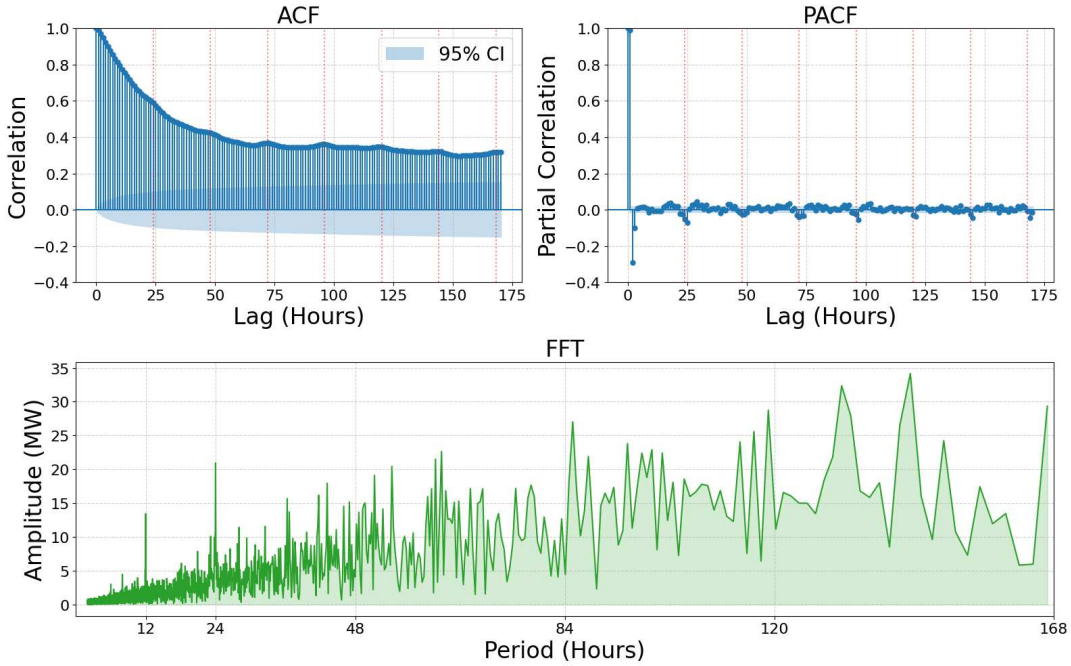

Figure 25: Vuolijoki CNEC RAM autocorrelation (ACF) and partial autocorrelation functions (PACF), and the Fast Fourier Transform (FFT) spectrum. CI = 95% confidence interval. Data retrieved from [32].

Figure 25 visualizes the autocorrelation (ACF) and partial autocorrelation functions (PACF), and the Fast Fourier Transform (FFT) of RAM time series. The slow and smooth decay in the ACF indicates that RAM has long memory, aligning with wind generation, which is effectively a persistent weather phenomenon. PACF displays strong positive partial correlation for the first-order ($t-1$) lag, and modest negative partial correlation for the second-order ($t-2$). However, since the target is to forecast the day-ahead RAM, it is not possible to use short-term lags (less than $t-24$) as inputs. Each of the three plots show a 24-hour periodic component, indicating a daily pattern. However, FFT suggests that the daily seasonality only explains approximately 20 MW, or less than 5%, of the total RAM variance.

These results indicate that while RAM exhibits some temporal dependencies, pure time series models are not likely to perform well. While this does not mean temporal features should be fully excluded, the model must rely primarily on exogenous Finnish features.

5.3 Machine Learning Methods

Based on the feature selection, forecasting RAM accurately requires capturing the relationships between exogenous features, such as Finnish wind generation and consumption, and RAM. While some temporal relationships exist, their effect is not as significant as that of the exogenous variables. Given the context, model interpretability is nearly as important as forecast accuracy, and is taken into account in model selection.

This thesis evaluates and compares two distinct ML methodologies: eXtreme Gradient Boosting (XGBoost) [35] and NeuralProphet [36]. Furthermore, two differently configured XGBoost models are trained, as explained in Section 5.3.1. XGBoost represents the state-of-the-art in tree-based ensemble learning. It is highly efficient and empirically proven to excel at capturing complex, non-linear relationships in tabular data, making it a strong candidate for RAM forecasting. Paired with the feature attribution method SHapley Additive exPlanations (SHAP), it enables high transparency and insight into feature importances [37]. NeuralProphet is evaluated due to its ability to capture time series properties alongside exogenous features and the predictive power of neural networks, while maintaining transparency, as introduced in Section 5.3.2.

5.3.1 eXtreme Gradient Boosting

XGBoost is a tree-based method that has been widely adopted across industries, with evidence of strong performance in energy-related forecast applications. Wu et al. [38] show that XGBoost is able to beat Long Short-Term Network and Support Vector Regression models in power price forecast accuracy. O'Connor et al. [39] show XGBoost outperforming deep learning and ARIMA models in Irish balancing markets. XGBoost is also shown to perform well in other energy market applications, such as load forecasting [40] [41].

At its core, XGBoost utilizes gradient tree boosting, which is an ensemble technique that produces a prediction model using an ensemble of weak learners. These learners are typically regression trees, which unlike decision trees, can contain continuous scores on each leaf. Given input features x_i, a regression tree space $\mathcal{F}$, and K additive functions, a tree ensemble model predicts the output $\hat{y}_i$ by summing the scores

$$\hat{y}_i = \phi(x_i) = \sum_{k=1}^{K} f_k(x_i), \quad f_k \in \mathcal{F},$$

where f_k is an independent regression tree. XGBoost is trained additively, meaning that at each iteration t, a new regression tree f_t is added to the model to minimize the regularized objective function, defined as

$$\mathcal{L}^{(t)} = \sum_{i=1}^{n} l(y_i, \hat{y}_i^{(t-1)} + f_t(x_i)) + \Omega(f_t),$$

where l is a differentiable convex loss function and $\Omega(f_t)$ is a model penalization term, which is given by

$$\Omega(f_t) = \gamma T + \frac{1}{2}\lambda \|w\|^2,$$

where T is the number of leaves in f_t and w represents leaf weights. The parameter γ specifies the minimum loss reduction required to create a new split, acting as a penalty for adding leaves, while λ represents the L_2 regularization term applied to the weights. By penalizing both the number of leaves and the magnitude of the weights, the algorithm tends to select simple and predictive functions, thus preventing overfitting.

While traditional gradient boosting algorithms minimize the loss function using first-order gradients, XGBoost expands the loss optimization with a second-order approximation. With the first and second-order gradients, $g_i = \partial_{\hat{y}^{(t-1)}} l(y_i, \hat{y}^{(t-1)})$ and $h_i = \partial^2_{\hat{y}^{(t-1)}} l(y_i, \hat{y}^{(t-1)})$, the loss function becomes

$$\mathcal{L}^{(t)} \simeq \sum_{i=1}^{n} \left[l(y_i, \hat{y}^{(t-1)}) + g_i f_t(\mathbf{x}_i) + \frac{1}{2} h_i f_t^2(\mathbf{x}_i) \right] + \Omega(f_t).$$

With this loss function, XGBoost captures the curvature of the loss, thus enabling faster convergence towards optimal leaf weights. Furthermore, XGBoost incorporates a novel sparsity-aware split finding algorithm that is advantageous when using real-world datasets. In many practical applications, input data is often sparse due to missing values or artifacts from feature engineering, such as one-hot encoding. [35] Given the adoption and performance in academia, efficiency, and robustness, XGBoost is a justified methodology for forecasting RAM.

Capturing temporal dependencies with XGBoost requires additional feature engineering, essentially adding lagged features, for instance RAM_{t-24}. However, since the forecast horizon is longer than the lag, training a single model with this lag feature would result in data leakage and thus over-accurate and unrealistic test forecasts. To overcome this, two different XGBoost models are trained. First, one direct model without any lagged variables, producing a direct 168 hour forecast. Second, a recursive model is trained, effectively consisting of seven separate models. Out of the seven, the first recursive model forecasts the first 24 hours of the 168 hour forecast horizon using the lagged feature RAM_{t-24}. Consequently, the second model forecasts hours 25-48, using the first model prediction as its lagged input. In total, seven consequent models produce recursive forecasts without any data leakage.

5.3.2 NeuralProphet

In domains such as the electricity market, understanding how features shape the forecast is essential. Unlike most deep learning approaches, NeuralProphet provides this necessary interpretability. Introduced in 2021 by researchers from Stanford and Facebook, it bridges typical time series forecasting properties with deep learning. [36]

Due to being relatively new, research on NeuralProphet is not as extensive as that of XGBoost. Some studies suggest that NeuralProphet often outperforms other deep learning methods in electricity load forecasting [42] [43]. NeuralProphet is also shown to provide more accurate electricity demand forecasts than traditional time series models, such as ARIMA and SARIMA [44]. In the context of forecasting RAM, NeuralProphet allows the comparison of effects from exogenous variables and temporal dynamics.

NeuralProphet is composed of modules that each contribute an additive component to the forecast. These components can be configured as linear regression or as neural networks, allowing the model to fit non-linear dynamics while keeping interpretability. NeuralProphet predicts the output with the model

$$\hat{y}_t = T(t) + S(t) + E(t) + F(t) + A(t) + L(t),$$

where the additive components at time t are defined as

- $T(t)$: Trend
- $S(t)$: Seasonal effects modeled using Fourier terms
- $E(t)$: Event and holiday effects
- $F(t)$: Regression effects for future-known exogenous variables
- $A(t)$: Auto-regression effects based on past observations
- $L(t)$: Regression effects for lagged observations of exogenous variables.

NeuralProphet models the trend as a continuous piecewise linear series, allowing the growth rate to change at specific changepoints. The equation for the trend at time t is defined as

$$T(t) = (\delta_0 + \Gamma(t)^T \delta) \cdot t + (\rho_0 + \Gamma(t)^T \rho),$$

where δ_0 is the initial growth rate, δ is a vector of rate adjustments at each changepoint, and ρ_0, ρ represent offset adjustments. The binary vector $\Gamma(t)$ indicates whether time t has passed each respective changepoint.

Seasonal effects are modeled using Fourier terms to produce smooth, interpretable functions. Additionally, this approach eliminates the need for manual feature engineering for daily cycles, for instance. For a seasonality with periodicity p and k Fourier terms, the effect is

$$S_p(t) = \sum_{j=1}^{k} \left(a_j \cos\left(\frac{2\pi jt}{p}\right) + b_j \sin\left(\frac{2\pi jt}{p}\right) \right),$$

where t is the specific time step being evaluated, j is the harmonic index functioning as a frequency multiplier for the base wave, and a_j and b_j are the coefficients fitted during training. This allows the model to handle non-integer periodicities, such as yearly seasonality in daily data. The total seasonality effect from a set of periodicities $\mathbb{P}$ is

$$S(t) = \sum_{p \in \mathbb{P}} S_p(t).$$

Sporadic events and national holidays are modeled as binary variables $e(t) \in \{0, 1\}$. The total effect from a set of events $\mathbb{E}$ is

$$E(t) = \sum_{e \in \mathbb{E}} z_e e(t),$$

where z_e is the fitted coefficient representing the impact of event e. NeuralProphet handles country-specific holiday event dates automatically.

Future regressors are essentially exogenous variables, the future values of which are known. The effect of all future regressors $\mathbb{F}$ is

$$F(t) = \sum_{f \in \mathbb{F}} d_f f(t),$$

where d_f is the fitted weight for each regressor f.

The auto-regression (AR) component in NeuralProphet is based on AR-Net, which uses a neural network for modelling the AR process [45]. For a forecast horizon h and p past observations, or lags, the predicted AR effects at forecast origin t are

$$A(t), \ldots, A(t + h - 1) = \text{AR-Net}(y_{t-1}, \ldots, y_{t-p}). \tag{7}$$

By default, AR-Net does not contain hidden layers or activation functions. Therefore, it is functionally identical to a classic $AR(p)$ model, but is optimized using deep learning methods. However, by introducing hidden layers, it can be extended to capture non-linear auto-regressive dependencies.

Lagged regressors correlate the target variable with past observations of the exogenous variables $x \in \mathbb{X}$. Each variable is processed through a similar AR-Net structure as in Equation (7), precisely

$$L_x(t), \ldots, L_x(t + h - 1) = \text{AR-Net}(x_{t-1}, \ldots, x_{t-p}).$$

This enables the model to capture the lagged effects of the exogenous variables on the target variable. Particularly in the context of forecasting RAM, the lagged regressors component could capture some dependencies XGBoost is not capable of without severe feature engineering.

To optimize the weights of these additive components, NeuralProphet uses the Huber loss function. It provides robustness for large errors while maintaining precision for smaller errors. The Huber loss is mathematically defined as

$$L_\delta(y, \hat{y}) = \begin{cases} \frac{1}{2}(y - \hat{y})^2 & \text{if } |y - \hat{y}| \leq \delta \\ \delta|y - \hat{y}| - \frac{1}{2}\delta^2 & \text{otherwise,} \end{cases}$$

where y is the realized RAM, $\hat{y}$ is the forecasted RAM, and δ is the threshold parameter (default = 1) separating the exponential and linear penalty regions. [36]

5.4 Feature Engineering

Before feeding the data to the models, it is necessary to clean and transform raw data into a suitable dataset. This dataset contains some missing values, as well as a significant regime shift during the time series. These issues are addressed in data preprocessing, as described in Section 5.4.1. Additionally, three feature candidates are dropped from the dataset due to multicollinearity with other features and absence of a relationship with RAM. Section 5.4.2 explains the engineering of additional features for improving model awareness of temporal dynamics and grid conditions.

5.4.1 Data Preprocessing

The dataset is parsed using UTC timestamps and checked for duplicate index entries. Duplicate timestamps are removed, keeping only the first occurrence. The data is then reindexed to a strict, continuous hourly frequency. This ensures that any missing hours in the original data are represented as NaN rows rather than being silently omitted.

On 15.11.2025 00:00:00 UTC, a significant regime shift occurred in the Finnish grid. Aurora Line, commissioned two days prior, increased the F_{max} of Vuolijoki CNEC by 300 MW. This directly increases the RAM by 300 MW as well. Given that this regime shift is permanent, the RAM values in the dataset prior to 15.11.2025 00:00:00 UTC are manually increased by 300 MW, preventing the models from treating this shift as a systemic signal.

Furthermore, three of the feature candidates displayed in Figure 24 are dropped from the training dataset. Residual load is excluded due to strong multicollinearity with consumption, wind, and solar generation. Additionally, hydro balance does not exhibit a significant relationship with RAM, and the effects of temperature are effectively captured within consumption data.

5.4.2 Additional Features

While NeuralProphet addresses temporal dependencies natively, lagged and cyclical features have to be engineered manually for XGBoost. First, the daily cycles are engineered using the sine and cosine of the hour (0-23) to ensure continuity between hours 23 and 0:

$$hour\_sin = \sin(2\pi \times \frac{hour}{24})$$
$$hour\_cos = \cos(2\pi \times \frac{hour}{24}).$$

As mentioned in Section 5.3.1, two models are trained using XGBoost, direct and recursive. Since the forecast horizon is 168 hours, using a $t-24$ lagged feature is infeasible with the direct model. However, the lagged feature RAM_{t-24} is engineered for the recursive model to use for the first 24 hours. This also requires trimming the first 24 hours of the dataset, since those rows do not have $t-24$ values.

Capturing the extreme peaks of RAM is essential, often more important than achieving the best average accuracy. In the first experiments of this thesis, it was found that NeuralProphet was able to capture the shape of the 168-hour curve, but was too conservative with both high and low peaks. To address this, a binary flag feature, *wind_fi_extreme*, is engineered to flag the hours with extreme wind generation (top and bottom 15%). In XGBoost, extreme wind hours are given larger weights in the loss function (see Section 5.5), and therefore this feature is included only with NeuralProphet.

5.5 Training and Hyperparameter Optimization

The processed dataset is partitioned into training and testing datasets at 01.02.2026 00:00:00 UTC. This split provides a 15-month continuous training horizon, ensuring the models are exposed to complete annual cycles of Finnish load and generation patterns. The final two months of the full dataset, February and March of 2026, are reserved for out-of-sample testing and model performance evaluation. Furthermore, the models can and will be tested using real-time data after the test dataset period.

Specific model configurations are established for both XGBoost and NeuralProphet, with direct and recursive XGBoost using the same configurations. For the XGBoost models, the objective function is set to minimize the symmetric mean squared error (MSE), with an upper bound of 2000 estimators. To ensure that the model captures the extreme RAM peaks, an instance weighting of $w = 5.0$ is applied to the top and bottom 15% of wind generation rows. For the NeuralProphet model, the architecture is configured to utilize a Huber loss function, with 100 training epochs, and apply automatic native Min-Max scaling to the inputs. Furthermore, to capture temporal dynamics, Fourier-based periodic components for yearly, weekly, and daily seasonalities are turned on. Both methodologies are constrained to a 168-hour week-ahead forecast horizon.

Before actual training, model hyperparameters are optimized using Optuna framework, which often outperforms other hyperparameter optimization frameworks while providing visualizations on the optimization process. [46]. Optuna searches hyperparameters using time-series expanding-window cross-validation within the training set. Specifically, the training data is split chronologically into five folds. During each of the 100 optimization trials, Optuna uses a set of hyperparameters from the defined search spaces as specified in Appendix A, and evaluates them by calculating the average MSE, with instance weighting enabled for XGBoost, across the five validation sets. By minimizing this MSE across 100 trials, Optuna converges and finds the optimal set of hyperparameters. Tables A.1 and A.3 summarize the fixed model configurations and hyperparameter search spaces. The complete fixed configurations, hyperparameter search spaces, and the final optimized values, as well as the Optuna visualization for both XGBoost models and NeuralProphet are detailed in Appendix A. For direct and recursive XGBoost, Optuna is run separately, although they share the fixed configurations and search spaces.

5.6 Evaluation and Interpretability Framework

For evaluating model performance, a multi-faceted framework is implemented. Standard quantitative error metrics are used for evaluating the global accuracy of the models, as described in Section 5.6.1. However, as shown in Section 4.3.3, Vuolijoki CNEC shadow prices and thus price spreads occur almost exclusively when RAM drops below 400 MW. Therefore, capturing bottom peaks is considered at least as critical as overall predictive accuracy. To assess this, a visual inspection of test data forecasts against realized data is applied and elaborated in Section 5.6.2. Furthermore, understanding feature importances is critical, as it allows an analyst to critically evaluate and adjust the real-time forecasts. This model interpretability is further discussed in Section 5.6.3.

5.6.1 Quantitative Error Metrics

The baseline predictive accuracy of the models is evaluated using two standard statistical metrics, Mean Absolute Error (MAE) and Root Mean Square Error (RMSE).

MAE provides a highly interpretable measure of the average absolute error between forecasted and realized RAM, represented in MW. It is defined as

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

where n is the number of observations in the test set, y_i is the realized RAM, and $\hat{y}_i$ is the forecasted RAM.

While MAE offers a straightforward view of average model performance, RMSE squares the errors before averaging them, which penalizes larger forecasting errors. It is defined as

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

Furthermore, Weighted Root Mean Square Error (WRMSE) metric is introduced to evaluate model performance during low-RAM hours. WRMSE applies a heavier penalty to forecasting errors in low peaks, and is defined as

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

where w_i is a piecewise weight assigned to each observation based on a realized RAM threshold of 500 MW,

$$w_i = \begin{cases} 5, & \text{if } y_i < 500 \\ 1, & \text{if } y_i \geq 500. \end{cases}$$

These three error metrics provide tools for assessing both the global accuracy and the operational reliability of the forecasting models.

5.6.2 Qualitative Visual Validation

While quantitative metrics like MAE and RMSE provide a standardized measure of global model accuracy, they might not accurately represent the peak capturing ability of the models. To ensure the models are practically viable, four 168-hour forecasts for RAM are plotted against the realized RAM in the out-of-sample test period. This visual inspection allows for the identification of structural forecasting behaviors that metrics are incapable of catching, such as temporal lagging, over-smoothing of peaks, or the failure to capture sudden changes in RAM.

5.6.3 Model Interpretability

For the XGBoost models, model interpretability is achieved using SHAP. SHAP values provide local interpretability by quantifying the exact contribution of each exogenous feature to the final RAM forecast. This allows analysts to verify if the model is weighting fundamentals in a manner consistent with physical grid realities.

On the other hand, NeuralProphet provides interpretability through its additive architecture. The model natively decomposes the final forecast into transparent components, allowing for the isolation of base trends and the specific impact of exogenous features, for instance. By evaluating these interpretability outputs, this thesis assesses not just what the models forecast, but also the most significant drivers of those forecasts.

6 Results

All model training and hyperparameter optimization were executed on a workstation equipped with an Intel Xeon Gold 6144 CPU @ 3.50GHz and 256 GB of RAM, utilizing Python 3.12. Optimized hyperparameters can be found in Appendix A. All results in this chapter refer exclusively to forecasting the RAM of Vuolijoki CNEC.

Five forecast origins with non-overlapping 168-hour forecast horizons were randomly selected from the test dataset, which spans from 1.2.2026 00:00:00 to 31.3.2026 23:00:00 (UTC). The error metrics introduced in Section 5.6.1 were calculated based on these five weeks, as presented in Section 6.1. The section also presents three forecasted weeks against the realized data, with each week having distinct properties for illustrative purposes. Section 6.2 discusses feature importances, along with interpreting the temporal properties.

6.1 Model Performance

Based on MAE and RMSE, XGBoost outperforms NeuralProphet, with minor distinction between direct and recursive XGBoost models. However, NeuralProphet achieves a smaller WRMSE than both XGBoost models as seen in Table 2.

Table 2: Model performance by error metric

Model	MAE	RMSE	WRMSE
XGBoost (Direct)	116.52	142.32	161.91
XGBoost (Recursive)	**113.56**	**136.56**	160.44
NeuralProphet	123.66	155.63	**156.75**

This would indicate that on average, XGBoost outperforms NeuralProphet, while during critical hours, NeuralProphet provides more accurate forecasts. Thus, while NeuralProphet is less accurate at predicting the general variance of the time series, its architecture is effective at capturing the low-RAM events that drive market congestion.

Visual validation of model forecasts against realized data shows promising results. In general, each model is able to produce a forecast that follows a similar path as that of the realized time series. Figure 26 shows an example of a forecast with high daily variance in RAM.

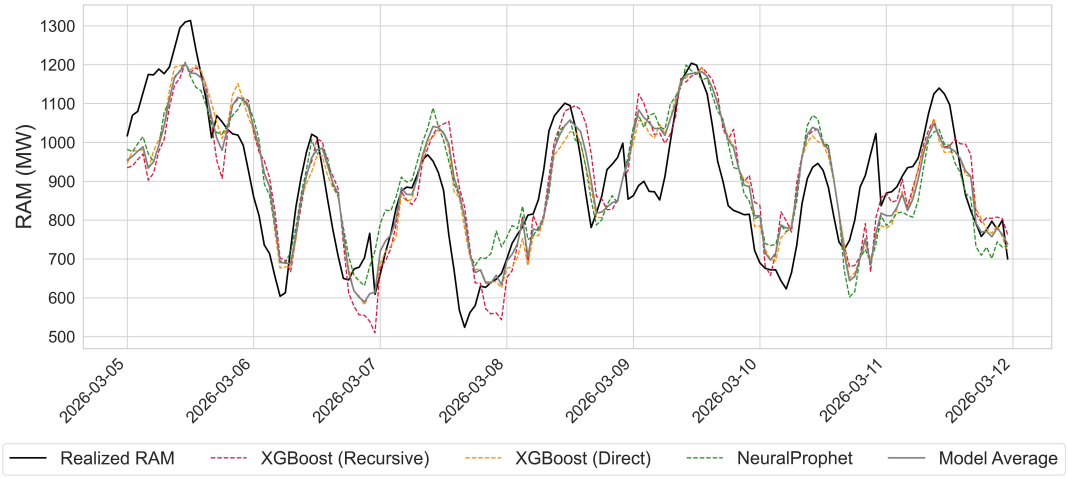

Figure 26: Model forecasts against realized RAM, high variance.

As explained in Section 5.1, the models use realized wind values as input features. While this mitigates the effect of wind forecast errors for the week-ahead RAM forecast, it does not completely eliminate it. For instance, on 28.02.2026 and on 03.03.2026, Finnish wind generation realized nearly 1.5 GW below the day-ahead forecasts [47]. On the other hand, on 25.02.2026, wind realized nearly 1 GW above forecasts. Because the RAM is calculated based on the day-ahead wind forecast, which in this case differs from realized wind generation, the RAM forecast deviates from the realized RAM, as seen in Figure 27.

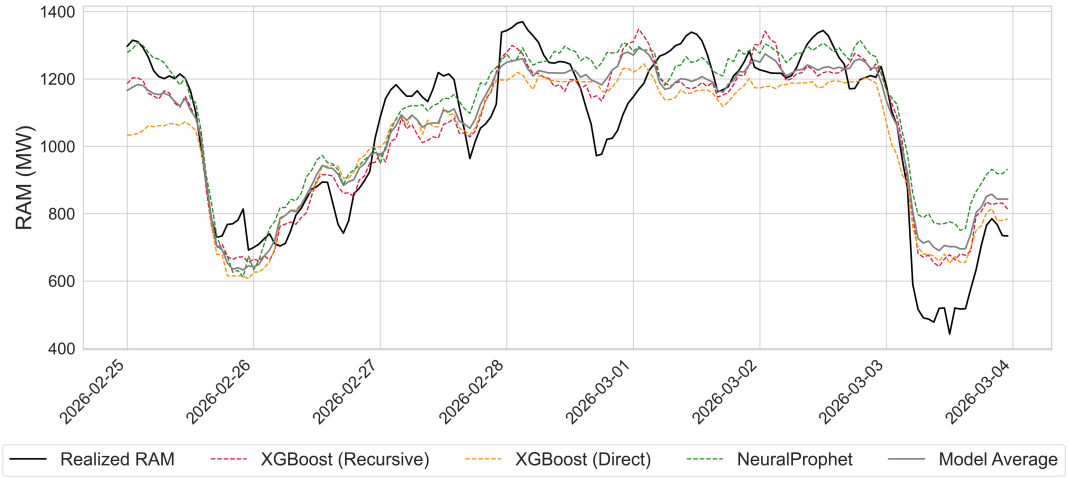

Figure 27: Model forecasts against realized RAM, low peaks with wind forecast errors.

However, while each model is able to capture the timing of the peaks, their magnitudes are often inaccurate due to these wind forecast deviations. Figure 28 visualizes the average RAM forecast error with the day-ahead wind forecast error for the corresponding period.

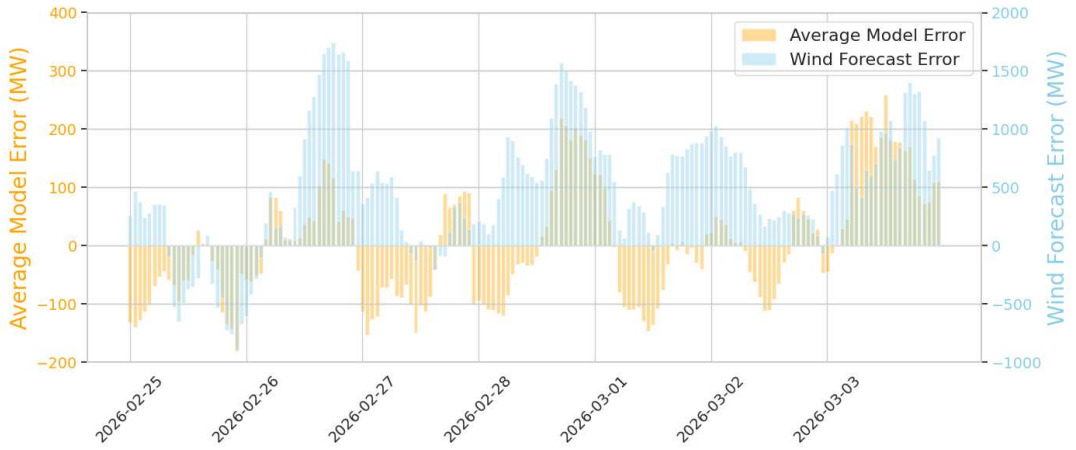

Figure 28: Errors in wind and average model forecasts against realized data. Wind forecast and realized data from Fingrid [47].

As demonstrated in Figure 28, the magnitude and timing of the average model error correlate with that of wind forecast error, exhibiting a correlation of 0.55 during the test period. Specifically, when the average model forecast is above the realized RAM, there are significant wind forecast errors in the same direction.

On 13.03.2026, RAM realized at a historically low level, 237 MW. Given the training data adjustment explained in Section 5.4.1, the training dataset used in this thesis does not include RAM values this low. Consequently, the models are not able to capture this low peak, as seen in Figure 29. In a real-world implementation, regular retraining would enable the model to learn these types of anomalies.

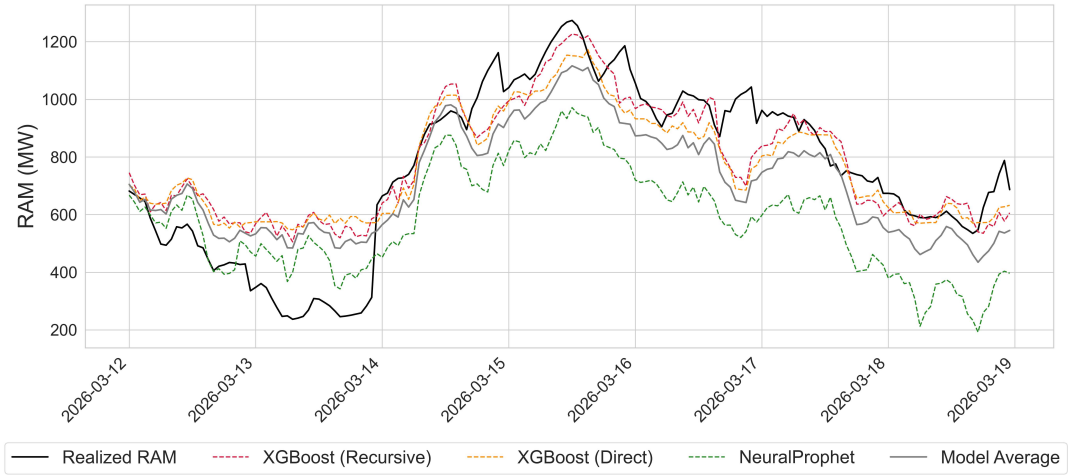

Figure 29: Model forecasts against realized RAM, extreme low peak.

6.2 Feature Importances and Interpretations

As discussed in Section 6.1, errors in forecasting RAM peaks exhibit significant correlation with wind forecast errors. This finding is also prominent in Figure 30,

which confirms the importance of wind generation as a driver for the recursive XGBoost output.

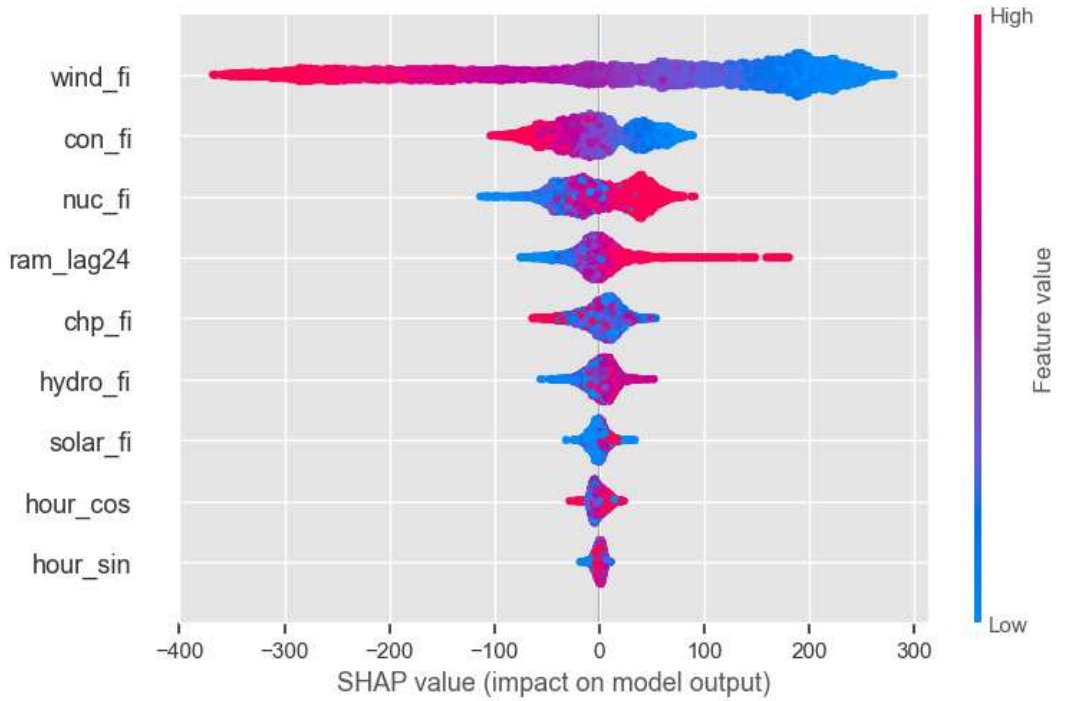

Figure 30: SHAP values of the recursive XGBoost model. Values on x-axis indicate the feature impact on RAM forecast (MW), given high or low feature values. Direct XGBoost exhibits similar SHAP values, without the lagged RAM feature.

As seen in Figure 30, when Finnish wind generation is high, it significantly decreases the Vuolijoki CNEC RAM forecast, and the other way around. This makes sense, since high wind generation, mostly located north from the consumption, increases the F_0 flow from northern Finland to the south, consequently reducing RAM. Finnish consumption, most of which is located in the south, has a similar effect on forecasts, since high consumption in the south increases the north-south flow of electricity. On the other hand, Finnish nuclear generation has a contrary effect, since the nuclear plants are located closer to consumption, thus reducing the need for north-south flow. Lagged RAM from 24 hours prior shows that previous RAM values increase the forecast significantly. CHP, hydro, and solar generation have a relatively minor impact on the model output compared to the first four features. As predicted in Section 5.2.2, the impact of hour of day is almost negligible.

In NeuralProphet, wind generation also stands out as the clear main driver for RAM forecasts, as seen in Figure 31.

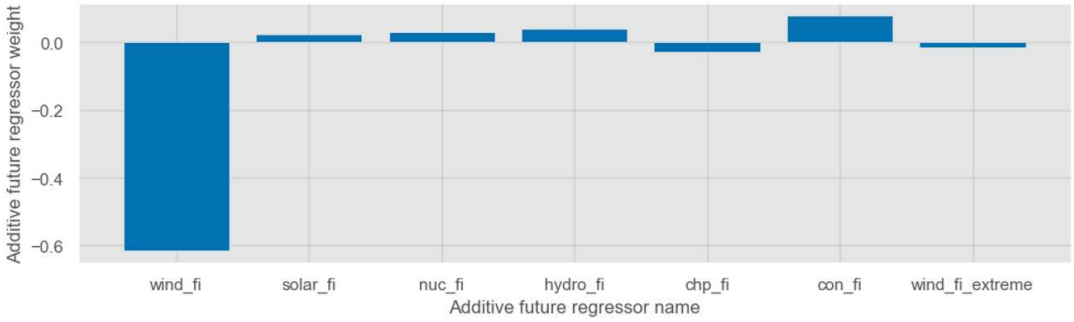

Figure 31: Exogenous regressor importances on NeuralProphet output. Each bar indicates how much the model output changes, given a one unit increase in input feature.

Apart from wind generation, other features, or future regressors, have relatively small impacts on NeuralProphet model forecasts, as seen in Figure 31. Surprisingly, the extreme wind binary flag has the smallest impact on model output. The second-most important feature is consumption. However, contrary to XGBoost, high consumption increases the NeuralProphet RAM forecast. While this is counterintuitive, it is likely explained by the fact that consumption exhibits significant daily, weekly, and yearly patterns, which NeuralProphet natively handles. These seasonality effects are visualized in Figures 32, 33, and 34.

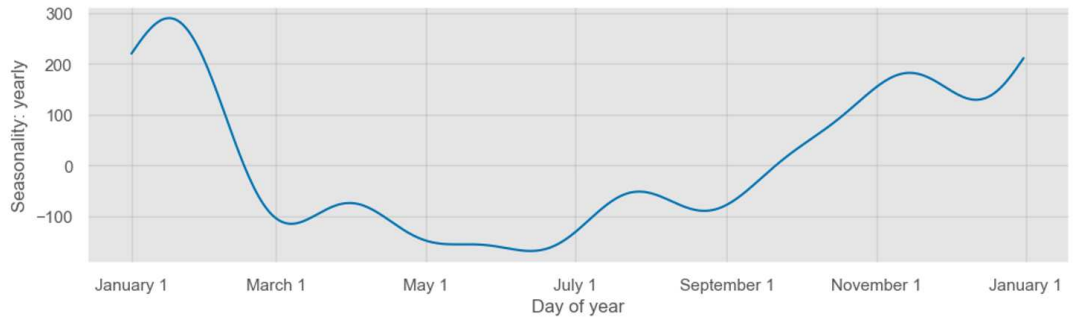

Figure 32: Yearly seasonality in NeuralProphet output.

The yearly seasonality component in Figure 32 follows an opposite pattern compared to the yearly time series of RAM in Figure 20. During summer, RAM is often relatively high due to low consumption and low wind, but the yearly seasonality component counterintuitively decreases the model output during summertime. Conversely, the yearly seasonality increases the model output during winter months.

The counterintuitive logic lies in the additive architecture of NeuralProphet, specifically in how the model optimizes the sum of its components rather than each individually. The $F(t)$ component in Figure 31 continuously overestimates the effect of wind generation on RAM, and the $S(t)$ component in Figure 32 uses the yearly seasonality to adjust the output for this error. Consequently, the sum of the components outputs the most accurate RAM forecast.

On the other hand, weekly seasonality in Figure 33 exhibits an intuitive pattern, where NeuralProphet increases the output during weekends, specifically on Sundays, when the increase is approximately 30 MW.

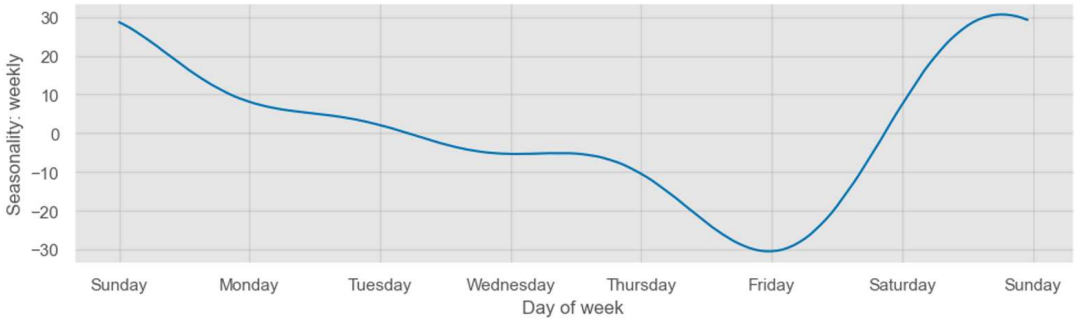

Figure 33: Weekly seasonality in NeuralProphet output.

This weekly seasonality aligns with weekly consumption, which is typically lower during weekends, leading to higher RAM. Weather-driven generation, such as wind, does not depend on the day of week, which is why the weekly seasonality effects are relatively minor.

The daily seasonality component in Figure 34 impacts the NeuralProphet output slightly more than the weekly component.

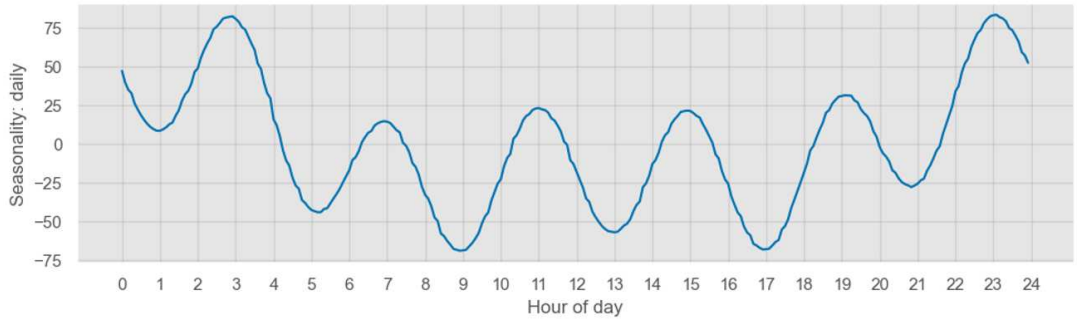

Figure 34: Daily seasonality in NeuralProphet output.

The daily seasonality component successfully captures the intraday dynamics of RAM. Typically, consumption peaks in the morning and in the evening, as seen on the 9th and 17th hours in Figure 34. Again, this results in a lower RAM forecast during those times. Conversely, consumption is lower and RAM is higher during nighttime. The wave-like characteristic of the curve is a direct consequence of the model fundamentally relying on Fourier terms to capture periodic effects.

NeuralProphet also visualizes the impact of its AR-Net on model output.

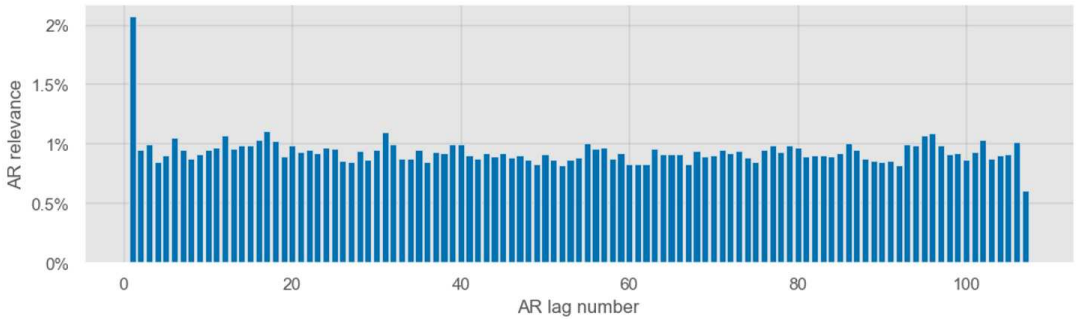

Figure 35: AR lag relevance in NeuralProphet output. 107 lags is the optimized NeuralProphet hyperparameter.

Figure 35 demonstrates the relative importance of the AR lags in the NeuralProphet model. The visualization shows a distinct spike at the $t - 1$ lag, and a uniformly flat distribution for all older observations. This confirms that the immediate previous hour is the primary autoregressive driver, aligning with the PACF shown in Figure 25.

7 Discussion and Outlook

This chapter synthesizes the empirical findings in the Finnish price area regarding the implementation of FBMC and discusses the proposed machine learning forecasting framework. First, the observed shifts in market dynamics, particularly the changes in price spreads and transmission flows between FI and SE3, as well as a grid bottleneck in FI, are discussed in Section 7.1. Furthermore, Section 7.2 discusses the commercial implications of these findings, highlighting the strategic value of accurate RAM forecasting and model interpretability for electricity trading. Finally, the limitations of the developed models are addressed in Section 7.3, and an outlook for future research for enhancing RAM forecasting capabilities is discussed in Section 7.4.

7.1 Interpretation of Market Dynamics

In theory, FBMC is designed to increase total SW by expanding the solution space for EUPHEMIA. However, the empirical analysis of the Finnish price area in Chapter 4 demonstrates that the real outcomes have been counterintuitive to the price convergence hypothesis. Specifically, the price dynamics between FI and SE3 price areas have shifted since the launch of FBMC. The historical norm of an FI price premium has been replaced by an SE3 price premium, and the occurrence of hours with equal prices between the two areas has dropped to below 10%. The analysis shows that the utilization rate of the FI-SE3 interconnectors has increased, with congestion frequency rising from 26.9% in pre-FBMC to 28.6% post-FBMC. Furthermore, counterintuitive flows between FI and SE3 have emerged, with electricity often flowing from the higher-priced area to the lower-priced area. While inefficient locally, these flows increase the total SW in the Nordic region under FBMC, as explained in Section 4.2.

Crucially, the analysis shows that there exists only one CNEC, referred to as Vuolijoki CNEC, that has been constraining the internal flow in Finland. The shadow price of this CNEC, whenever non-zero, exhibits strong correlation with the FI SPOT price, making it a particularly interesting CNEC to inspect. However, it is important to acknowledge that future grid infrastructure improvements are designed to increase transmission capacity and may mitigate the severity of this specific Vuolijoki CNEC as a bottleneck in the future. For instance, the commissioning of the Aurora Line in November 2025 permanently increased the RAM of the Vuolijoki CNEC by 300 MW.

7.2 Implications for Electricity Trading

The RAM and PTDFs of all CNECs in the Nordic system effectively define the power system constraints and thus the solution space for EUPHEMIA. While the PTDF values of Vuolijoki CNEC are often relatively stable, its RAM exhibits significant variance. Consequently, the solution space for EUPHEMIA changes along the shifts in its RAM.

Forecasting near-term SPOT prices is a crucial part of the performance of nearly any company participating in electricity trading. This forecasting is often integrated with fundamental market simulation tools, which typically use realistic power system

constraints. However, the RAM and thus accurate transmission constraints are not published before the SPOT auction. Therefore, being able to forecast RAM accurately could improve SPOT price forecasting accuracy, and thus would be valuable for any industrial electricity producer or consumer. For instance, by forecasting low RAM and subsequent price spikes, a flexible power producer can optimize its power dispatch schedule to maximize revenue.

Chapter 6 shows promising results for forecasting the week-ahead hourly RAM of Vuolijoki CNEC. The models demonstrate strong performance in forecasting the timing and direction of RAM fluctuations. While the recursive XGBoost provides the highest baseline accuracy across the general forecast horizon, NeuralProphet achieves the lowest WRMSE, indicating better performance during low peaks. However, the absolute magnitude of the RAM peaks is still often miscalculated, often caused by exogenous wind forecast errors.

The feature importances elaborated in Section 6.2 make the models highly interpretable. Consequently, an analyst can clearly identify the main drivers of a RAM forecast, enabling them to assess the outputs critically. This transparency, enabled by SHAP and the native architecture of NeuralProphet, gives these models an advantage over black-box neural networks. For instance, it allows an analyst to understand how the RAM forecast would shift in response to a change in wind generation forecasts.

While a RAM forecast itself provides little insight into Nordic price formation, it can be given as an input for a simulation model, which implements transmission constraints similar to Equation (4). Consequently, accurate transmission constraints are likely to improve SPOT price forecasting accuracy.

7.3 Model Limitations

As explained in Chapter 6, the models exhibit some limitations. The models were trained on realized wind generation data to account for the significant wind forecast errors in the week-ahead forecast horizon. However, TSOs define the day-ahead RAM based on the day-ahead wind generation forecast, from which the realized data might differ. In a real-life implementation, the models should use wind forecasts as input features, since realized data is not known at the time of forecast. This would likely increase forecast accuracy for the first day, but cause larger errors further in the forecast horizon. Depending on the exogenous wind forecast accuracy, it is likely that only the first few days would have RAM forecasts that can be trusted.

One limitation of the NeuralProphet model lies in its internal representation of physical grid dynamics. NeuralProphet optimizes the global loss of the sum of its components, which may result in internal weights that mathematically balance each other rather than representing physical grid properties. Consequently, while the final forecast is often feasible, the individual additive components, such as yearly seasonality, exhibit behaviours counterintuitive to actual Nordic market fundamentals. While the transparency is an advantage, this is a operational limitation for an analyst trying to understand how grid fundamentals affect the forecast.

While XGBoost is highly interpretable and aligns with grid dynamics, it does not come without any limitations. Specifically, the direct XGBoost lacks regime awareness

NeuralProphet has, and would not be able to adjust to future regime shifts without retraining the model. Furthermore, including lagged features in XGBoost requires manual feature engineering, and the recursive forecasting framework can get complex with the seven separate models.

7.4 Outlook and Future Research

The machine learning framework developed in this thesis demonstrates significant potential for forecasting the RAM of Vuolijoki CNEC. However, flow-constraining CNECs exist also in Sweden, Norway, and Denmark. Since the whole Nordic area has to be considered when forecasting FI SPOT price, RAM forecasts for CNECs other than Vuolijoki CNEC are needed. Therefore, future research must focus on model scalability. While the RAM of Vuolijoki CNEC is driven by exogenous Finnish fundamental variables, CNECs in other price areas have distinct characteristics and features. For instance, the RAM of the CNECs in the SE3 price area are largely driven by nuclear generation and consumption in the area. In contrast, CNECs in Norway are highly sensitive to the hydrological balance.

Furthermore, because week-ahead RAM forecasts are highly sensitive to errors in exogenous weather predictions, future research could integrate weather scenarios as inputs. With this approach, models could output a scenario-based distribution for RAM rather than a single absolute value. Using the lowest RAM scenario, an analyst could assess how much a given CNEC would constrain the flow and consequently drive price spreads in a tail scenario.

From a commercial viewpoint, the key objective would be to identify the most constraining CNECs and the fundamental features driving their RAM values in other Nordic price areas, and use the framework introduced in this thesis to forecast their RAM. Furthermore, integrating these RAM forecasts into a market simulation tool would enable assessing if the RAM forecasts would improve SPOT price forecasting accuracy.

8 Conclusions

The implementation of FBMC in the Nordic electricity market marked a fundamental shift in capacity calculation and market clearing. This thesis evaluated the empirical impacts of this transition on the Finnish price area and developed a machine learning framework for forecasting RAM, addressing the research gap following the FBMC launch in October 2024.

The empirical analysis confirmed that the transition to FBMC has changed the dynamics between FI and SE3 price areas. Rather than driving increased price convergence, the new mechanism has increased the number of hours with counterintuitive flows between the two areas, and decreased the number of hours with converged prices. Crucially, the research identified a single internal transmission line, the Vuolijoki CNEC, as the primary bottleneck constraining the electricity flow within Finland. Because the variance in the RAM of this specific element was shown to be driven almost entirely by internal Finnish grid fundamentals, it emerged as an optimal target for forecasting.

Consequently, a machine learning framework was developed to forecast the week-ahead hourly RAM of the Vuolijoki CNEC. The results demonstrate that both XGBoost and NeuralProphet are able to capture the RAM fluctuations, though XGBoost, specifically when used recursively achieves a higher overall predictive accuracy. However, model performance is highly sensitive to exogenous feature quality, particularly regarding the alignment between realized and forecasted wind generation data. Based on model interpretability frameworks, specifically SHAP and additive component analysis, Finnish wind generation was identified as the primary feature driving RAM variance, as increased generation in the north directly increases the north-south flow and inversely impacts the RAM.

For electricity market participants, accurately forecasting the RAM of constraining CNECs has practical implications for electricity trading, since it could improve SPOT forecasting accuracy. By scaling the machine learning framework introduced in this thesis to other CNECs, and integrating the forecasts into market simulation tools, analysis and trading teams can better anticipate transmission bottlenecks and forecast the subsequent price spreads. This foresight ultimately supports decision-making and the optimization of generation dispatch, day-ahead bidding strategies, and hedging operations under FBMC in the Nordic market.

References

[1] Petr Spodniak, Kimmo Ollikka, and Samuli Honkapuro. *The Relevance of Wholesale Electricity Market Places: The Nordic Case*. Valtion Taloudellinen Tutkimuskeskus VATT, 2019.

[2] European Commission. Commission Staff Working Document SWD(2017) 383 final. Staff Working Document SWD(2017) 383 final, European Union, 2017.

[3] Heidar Jafarizadeh, Eliyad Yamini, Seyed Mohammad Zolfaghari, Farbod Esmaeilion, M. El Haj Assad, and M. Soltani. Navigating Challenges in Large-Scale Renewable Energy Storage: Barriers, Solutions, and Innovations. *Energy Reports*, 12:2179–2192, 2024.

[4] Nordic RCC. About the Nordic Regional Coordination Centre. https://nordic-rcc.net/about/, 2025. Accessed: 9.4.2025.

[5] Darryl R Biggar and Mohammad Reza Hesamzadeh. *The Economics of Electricity Markets*. John Wiley & Sons, 2014.

[6] Stephen Littlechild. Competition and Contracts in the Nordic Residential Electricity Markets. *Utilities Policy*, 14(3):135–147, 2006.

[7] Nordic Energy Regulators. An Overview of the Nordic Electricity Market. Nordic Energy Regulators (NordREG), 2025. Accessed: 9.5.2025.

[8] Volue. Energy Market Data & Forecasts, 2026. Accessed: 16.4.2026.

[9] Ember. Electricity Data Explorer. https://ember-energy.org/data/electricity-data-explorer/, 2026. Accessed: 4.5.2026.

[10] Oleksandr Prokhorov and Dina Dreisbach. The Impact of Renewables on the Incidents of Negative Prices in the Energy Spot Markets. *Energy Policy*, 167:113073, 2022.

[11] Nord Pool. Day-Ahead Market Prices: Hourly Data for Finland. https://data.nordpoolgroup.com/auction/day-ahead/prices?deliveryDate=latest&currency=EUR&aggregation=Hourly&deliveryAreas=FI, 2026. Accessed: 16.3.2026.

[12] Norwegian Energy Regulatory Authority. Wholesale Market Timeframes. Norwegian Energy Regulatory Authority (NVE), 2025. Accessed: 28.4.2025.

[13] Austan Goolsbee, Steven D. Levitt, and Chad Syverson. *Microeconomics*. Worth Publishers, New York, NY, USA, Second edition, 2016.

[14] Peter D. Lund, John Byrne, Reinhard Haas, and Detlef Flynn, editors. *Advances in Energy Systems: The Large-scale Renewable Energy Challenge*. Wiley, 2019.

[15] Europex. EPEX SPOT: Member Profile. https://www.europex.org/members/epex-spot/, 2025. Accessed: 9.5.2025.

[16] Day-Ahead Trading - Nord Pool. https://www.nordpoolgroup.com/en/trading/Day-ahead-trading/. Accessed 4.4.2025.

[17] All NEMO Committee. Euphemia Public Description: Single Price Coupling Algorithm. Technical report, All NEMO Committee, August 2024. Accessed: 16.4.2025.

[18] Camille Hamon and Akshaya Tammanur Ravi. Impact of Flow-Based on the Intraday Market. Technical Report 2023:962, Energiforsk, Stockholm, Sweden, November 2023. Accessed: 28.4.2025.

[19] Fingrid Oyj. Balancing Energy and Balancing Capacity Markets (mFRR). https://www.fingrid.fi/en/electricity-market/reserves_and_balancing/balancing-energy-and-balancing-capacity-markets/, 2025. Accessed: 28.4.2025.

[20] Nasdaq. European Commodities Market Prices. Nasdaq Commodities, 2025. Accessed: 28.4.2025.

[21] Hans Nylund. *Integration of the Nordic Electricity Grids: Incentives, Cost Sharing and Regional Perspective in Transmission Investments*. Licentiate thesis, Luleå University of Technology, Luleå, Sweden, 2010. Accessed: 15.5.2025.

[22] Peter Sores, Daniel Divenyi, and David Raisz. Flow-Based Capacity Calculation Method Used in Electricity Market Coupling. In *Proceedings of the 10th International Conference on the European Energy Market (EEM)*, 2013.

[23] Kenneth Van den Bergh, Jonas Boury, and Erik Delarue. The Flow-Based Market Coupling in Central Western Europe: Concepts and Definitions. *The Electricity Journal*, 29(1):24–29, 2016.

[24] Rafael Finck. Impact of Flow Based Market Coupling on the European Electricity Markets. *Sustainability Management Forum*, 29(4):173–186, November 2021.

[25] Tarjei Kristiansen. The Flow Based Market Coupling Arrangement in Europe: Implications for Traders. *Energy Strategy Reviews*, 27:100444, January 2020.

[26] Nordic TSOs. Supporting Document for the Nordic Capacity Calculation Region's Proposal for Capacity Calculation Methodology in Accordance with Article 20(2) of Commission Regulation (EU) 2015/1222 of 24 July 2015 Establishing a Guideline on Capacity Allocation and Congestion Management. Technical report, Nordic Capacity Calculation Region, May 2018. Amended Proposal Dated 16.5.2018.

[27] Joint Allocation Office (JAO). Nordic Publication Tool Handbook, 2024. Accessed: 19.2.2026.

[28] Ferenc Nagy, Melinda Nagy, Luca Toth, Ágnes Esze, Ákos Arnold, Dániel Diveny, and Gábor Szathmari. *Day-Ahead Flow-Based Capacity Calculation and Market Coupling in Core CCR*, pages 219–258. 2025.

[29] Pieter Schavemaker. *Day-Ahead Flow-Based Capacity Calculation in CCR Nordic*, pages 259–277. Springer Nature Switzerland, Cham, 2025.

[30] Fingrid Oyj. Aurora Line Power Transmission Connection to Be Commissioned on November 13. https://www.fingrid.fi/en/news/news/2025/aurora-line-power-transmission-connection-to-be-commissioned-on-november-13/, 2025. Accessed: 18.2.2026.

[31] Marten Ovaere, Michiel Kenis, Kenneth Van den Bergh, Kenneth Bruninx, and Erik Delarue. The Effect of Flow-Based Market Coupling on Cross-Border Exchange Volumes and Price Convergence in Central Western European Electricity Markets. *Energy Economics*, 118:106519, 2023.

[32] Joint Allocation Office (JAO). JAO Publication Tool – Nordic Flow-Based Market Coupling Data. https://publicationtool.jao.eu/, 2024. Accessed: 12.2.2026.

[33] Fingrid Oyj. Fingrid Grid Map Service. https://karttapalaute.fingrid.fi/?setlanguage=en#, 2025. Accessed: 19.2.2026.

[34] Mika Andersson. Forecasting of Zonal Power Transfer Distribution Factors in Flow-Based Market Coupling. Master's Thesis, Aalto University, School of Science, Helsinki, Finland, November 2023. Supervisor: Prof. Ahti Salo, Advisor: M.Sc. (Tech.) Jouni Mäenpää.

[35] Tianqi Chen and Carlos Guestrin. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, page 785–794, New York, NY, USA, 2016. Association for Computing Machinery.

[36] Oskar Triebe, Hansika Hewamalage, Polina Pilyugina, Nikolay Laptev, Christoph Bergmeir, and Ram Rajagopal. NeuralProphet: Explainable Forecasting at Scale. arXiv:2111.15397, 2021.

[37] Scott M. Lundberg and Su-In Lee. A Unified Approach to Interpreting Model Predictions. *CoRR*, abs/1705.07874, 2017.

[38] Kehe Wu, Yanyu Chai, Xiaoliang Zhang, and Xun Zhao. Research on Power Price Forecasting Based on PSO-XGBoost. *Electronics*, 11(22), 2022.

[39] Ciaran O'Connor, Joseph Collins, Steven Prestwich, and Andrea Visentin. Electricity Price Forecasting in the Irish Balancing Market. *Energy Strategy Reviews*, 54:101436, 2024.

[40] Kedong Zhu, Jian Geng, and Ke Wang. A Hybrid Prediction Model Based on Pattern Sequence-Based Matching Method and Extreme Gradient Boosting for Holiday Load Forecasting. *Electric Power Systems Research*, 190:106841, 2021.

[41] Xin Zhao, Qiushuang Li, Wanlei Xue, Yihang Zhao, Huiru Zhao, and Sen Guo. Research on Ultra-Short-Term Load Forecasting Based on Real-Time Electricity Price and Window-Based XGBoost Model. *Energies*, 15(19), 2022.

[42] Md Jamal Ahmed Shohan, Md Omar Faruque, and Simon Y. Foo. Forecasting of Electric Load Using a Hybrid LSTM-Neural Prophet Model. *Energies*, 15(6), 2022.

[43] Shuai Lu and Taotao Bao. Short-Term Electricity Load Forecasting Based on NeuralProphet and CNN-LSTM. *IEEE Access*, 12:76870–76879, 2024.

[44] Kamran Ali, Essa Q. Shahra, Shadi Basurra, and Moad Idrissi. Forecasting Electricity Demand for UK Consumers: A Time Series Machine Learning Approach. In *Advances on Intelligent Computing and Data Science II*, pages 209–220, Cham, 2025. Springer Nature Switzerland.

[45] Oskar Triebe, Nikolay Laptev, and Ram Rajagopal. AR-Net: A Simple Auto-Regressive Neural Network for Time-Series. *CoRR*, abs/1911.12436, 2019.

[46] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A Next-Generation Hyperparameter Optimization Framework, 2019.

[47] Fingrid. Fingrid Open Data, 2026. Accessed: 2026-04-14.

Appendix

A Optimized Hyperparameters

Table A.1: XGBoost (Direct) configurations and hyperparameter search space

Parameter	Search Space	Optimized Value	Description
Fixed Configuration			
Forecast Horizon (h)	–	168	Week-ahead hourly forecasting
Loss Function	–	MSE	Symmetric mean squared error loss
Max Estimators	–	2000	Upper cap for boosting rounds
Instance Weighting	–	$w = 5.0$	5x weight for top/bottom 15% wind
Optuna Search Space			
Learning Rate	$[10^{-2}, 10^{-1}]$ (Log)	0.0167	Step size shrinkage to prevent overfitting
Max Depth	[4, 15]	4	Controls tree complexity and interactions
Min Child Weight	[1, 5]	5	Minimum sum of instance weight in a leaf
Subsample	[0.5, 1.0]	0.84	Fraction of rows sampled per tree
Colsample by Tree	[0.5, 1.0]	0.80	Fraction of features sampled per tree
Reg Alpha (α)	$[10^{-8}, 1.0]$ (Log)	7.34×10^{-5}	L_1 regularization for sparsity
Reg Lambda (λ)	$[10^{-8}, 1.0]$ (Log)	0.0025	L_2 regularization for smoothness

Table A.2: XGBoost (Recursive) configurations and hyperparameter search space

Parameter	Search Space	Optimized Value	Description
Fixed Configuration			
Forecast Horizon (h)	–	168	Week-ahead hourly forecasting
Loss Function	–	MSE	Symmetric mean squared error loss
Max Estimators	–	2000	Upper cap for boosting rounds
Instance Weighting	–	$w = 5.0$	5x weight for top/bottom 15% wind
Optuna Search Space			
Learning Rate	$[10^{-2}, 10^{-1}]$ (Log)	0.0672	Step size shrinkage to prevent overfitting
Max Depth	[4, 15]	13	Controls tree complexity and interactions
Min Child Weight	[1, 100]	74	Minimum sum of instance weight in a leaf
Subsample	[0.5, 1.0]	0.55	Fraction of rows sampled per tree
Colsample by Tree	[0.5, 1.0]	0.89	Fraction of features sampled per tree
Reg Alpha (α)	$[10^{-8}, 1.0]$ (Log)	0.0034	L_1 regularization for sparsity
Reg Lambda (λ)	$[10^{-8}, 1.0]$ (Log)	0.0019	L_2 regularization for smoothness

Table A.3: NeuralProphet configurations and hyperparameter search space

Parameter	Search Space	Optimized Value	Description
Fixed Configuration			
Forecast Horizon (h)	–	168	Week-ahead hourly forecasting
Loss Function	–	Huber	Fixed internally in NeuralProphet
Seasonality	–	Yearly, Weekly, Daily	Fourier-based periodic components
Training Epochs	–	100	Number of gradient descent iterations
Normalization	–	Auto	Internal Min-Max scaling of inputs
Optuna Search Space			
Learning Rate	$[10^{-4}, 10^{-1}]$ (Log)	0.00456	Step size for stochastic gradient descent
n_lags (AR-Net)	[48, 336]	107	Observations for auto-regression (2–14 days)
Trend Reg (λ_{trend})	[0.01, 10.0] (Log)	3.49	Regularization for piecewise linear trend

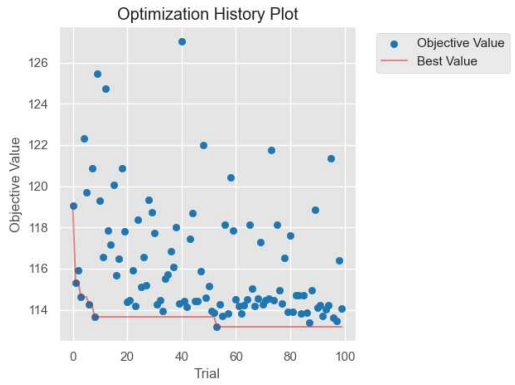

Figure A.1: Hyperparameter optimization history for direct XGBoost.

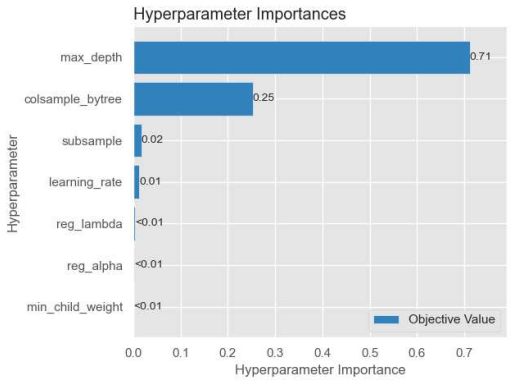

Figure A.2: Hyperparameter importances for direct XGBoost.

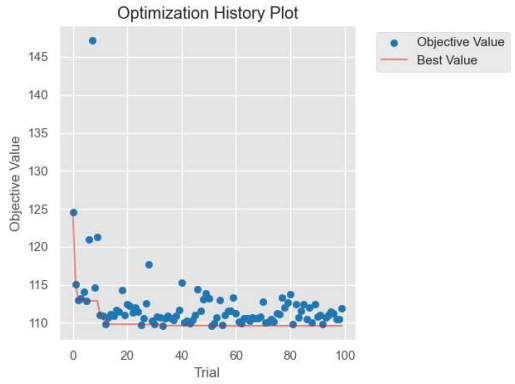

Figure A.3: Hyperparameter optimization history for recursive XGBoost.

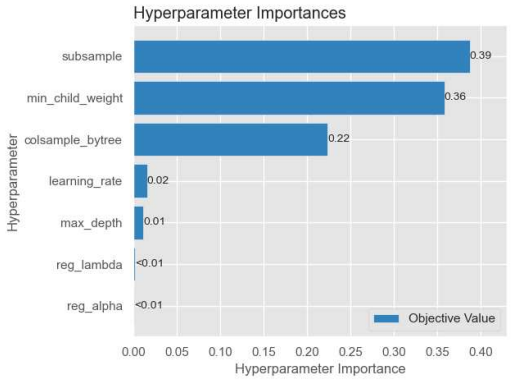

Figure A.4: Hyperparameter importances for recursive XGBoost.

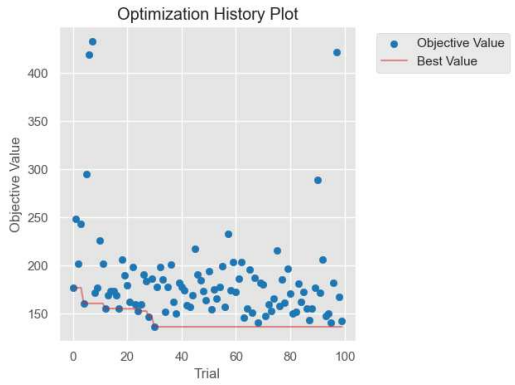

Figure A.5: Hyperparameter optimization history for NeuralProphet.

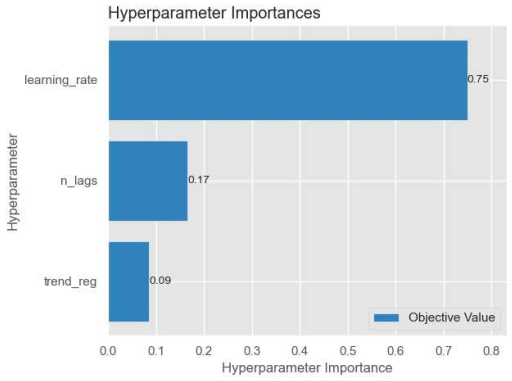

Figure A.6: Hyperparameter importances for NeuralProphet.