DIVISION SAFETY AND
TRANSPORT
ELECTRIC POWER
SYSTEM

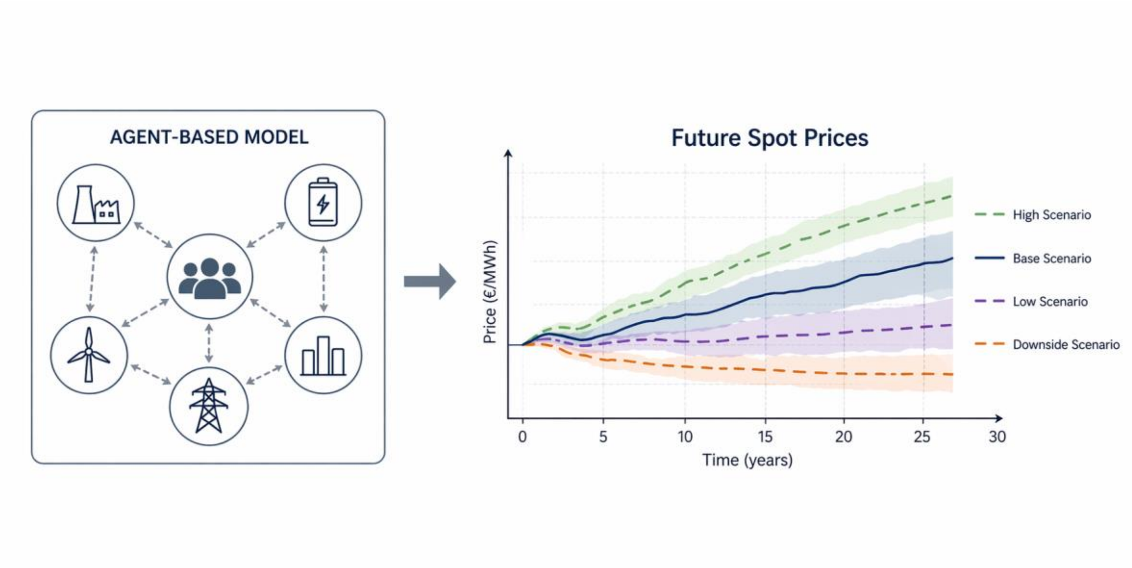

Agent-Based Model for Long-Term Price Scenarios of Frequency Regulation Services

Erik Alvarez, Markus Eriksson, Jan Hendrik Rolfes

RISE Report 2026:50

Agent-Based Model for Long-Term Price Scenarios of Frequency Regulation Services

Erik Alvarez, Markus Eriksson, Jan Hendrik Rolfes

Abstract

Agent-Based Model for Long-Term Price Scenarios of Frequency Regulation Services

This report presents the design, implementation, and validation of an Agent-Based Model (ABM) developed for long-term price scenario generation of Frequency Containment Reserve (FCR) services in the Nordic power system. The ABM simulates the Nordic day-ahead spot market through 11 meta-agents whose bidding behaviour is learned from 2019 historical data via feedforward neural networks. The model is coupled with a Temporal Fusion Transformer (TFT) that translates the simulated spot prices—together with LMA2021 scenario variables—into long-term price trajectories for FCR-N, FCR-D up, and FCR-D down through 2050. Results for the 2045 scenario horizon show that the ABM-supported TFT produces forecasts that closely track LMA2021 benchmarks while exhibiting substantially lower variance and fewer extreme price events, particularly for FCR-D up. The coupled pipeline is evaluated under the Elektrifiering förnybart (EF) and Elektrifiering planerbart (EP) scenarios, with FCR-N showing the tightest alignment between approaches. Key limitations include the use of pre-pandemic training data, hourly temporal resolution, and potential error propagation across the coupled models.

Keywords: FCR-N, FCR-D, ABM, TFT, ML, Long-term, Forecasts, Scenarios

RISE Research Institutes of Sweden AB

RISE Report 2026:50
ISBN: 978-91-90109-79-3
Linköping/Göteborg 2026

Content

Preface

This report details the work with regards to the development of an agent-based simulation model (ABM) in order to facilitate long-term forecasting of the Swedish balancing market, specifically for Frequency Containment Reserve services (FCR-N, FCR-D up, and FCR-D down). It is part of Project “Long-term price scenarios of frequency regulation services” (Project number: P2022-00752), funded by Energimyndigheten and supported by Svenska kraftnät, Vattenfall, Flower, Kungliga Tekniska Högskolan (KTH), Chalmers Tekniska högskolan, Skellefteå Kraft and Bixia elbolag.

The project was initiated to capture the rising complexities of a transforming energy landscape, particularly the integration of higher levels of variable renewable energy. This shift brings increased volatility and a need for greater flexibility, impacting how ancillary service markets are structured and managed. This report is intended to supplement the main project report with technical details on the agent-based model and serves as a complement to another technical report on the temporal fusion transformer (TFT) part of the project - referred to hereafter as the TFT Technical Report [1]. At last, it further describes how the two models support each other and how they can be coupled to achieve the final long-term forecast for frequency regulation services. To this end, the coupled approach combines the mainly simulation-based ABM with the machine learning driven TFT and established long-term electricity price scenarios provided by Svenska Kraftnät in their Långsiktig Marknadsanalys 2021 (LMA2021). Hence, this approach aligns with the official system development scenarios, but it also imposes specific limitations that are addressed later in this document. To note is that during the duration of the project, a newer Långsiktig Marknadsanalys 2024 (LMA2024) was released. However, since this was not available at the start of the project, the electricity price scenarios of the slightly older LMA2021 were used.

Summary

This report presents a coupled Agent-Based Model (ABM) and Temporal Fusion Transformer (TFT) pipeline for generating long-term price scenarios of Frequency Containment Reserve (FCR) services in the Nordic power system to the year 2050. The ABM simulates the Nordic day-ahead spot market using 11 meta-agents trained on 2019 bidding curve data. Resulting spot price forecasts for SE1–SE4 are fed into a TFT model that derives FCR-N, FCR-D up, and FCR-D down price trajectories, with scenario inputs drawn from Svenska Kraftnät's LMA2021 long-term market analysis.

Evaluated against LMA2021 benchmark forecasts for the 2045 horizon under the Elektrifiering förnybart (EF) and Elektrifiering planerbart (EP) scenarios, the ABM-supported TFT tracks weekly-average spot prices closely while producing notably lower forecast variance—particularly in SE3 and SE4. Downstream FCR price scenarios show strong agreement with the LMA2021-supported TFT: FCR-N mean squared error is below 8 $(\mathrm{EUR/MWh})^2$ in both scenarios, and FCR-D up variance is reduced by more than 60% under EF when using ABM inputs. The ABM's stabilising effect on the forecast is most pronounced for FCR-D up, where LMA2021-driven inputs generate multi-year price spike events that are absent in the ABM-supported output.

The primary limitations of the approach are: (i) dependence on pre-pandemic training data that does not reflect post-2020 market variations; (ii) an hourly temporal resolution that limits sub-hourly accuracy; and (iii) potential error propagation from the ABM into the TFT stage. Future work should focus on updating training data to include recent market conditions, training the ABM directly on ancillary service bidding curves to eliminate the need for the TFT bridging step, and integrating LMA2024 scenario inputs once available.

1 Background

1.1 Frequency Regulation Services in the Nordic Power System

The Nordic power system—spanning Sweden, Norway, Finland, and Denmark—relies on a set of ancillary services to maintain stable grid frequency at 50 Hz. Frequency Containment Reserves (FCR) form the first line of automatic frequency regulation and are procured through competitive day-ahead auctions operated jointly by the Nordic Transmission System Operators (TSOs). FCR-N (Normal) is activated during normal frequency deviations (49.9–50.1 Hz), while FCR-D (Disturbance) is activated in response to large frequency disturbances falling below 49.5 Hz—with separate products for upward (FCR-D up) and downward (FCR-D down) regulation.

Historically, hydropower has dominated FCR provision in the Nordic region due to its fast response characteristics and large installed capacity. However, rising shares of variable renewable energy (wind and solar), combined with the phased retirement of synchronous generators, are transforming the composition of FCR providers. Battery energy storage systems (BESS) have become increasingly active participants, and industrial demand response is growing. These structural changes make long-term price forecasting of FCR markets both more important and more challenging.

1.2 Outlook to 2030

The Nordic TSOs have collectively identified a trend of increasing FCR demand driven by higher variability in generation and rising net load fluctuations associated with electrification. Svenska Kraftnät's Balancing Market Outlook 2030 [3] projects a continued growth in FCR-D procurement volumes as system inertia declines with the displacement of conventional synchronous generation. This creates upward pressure on FCR prices, particularly during periods of tight reserve margins.

The entry of new technology types—particularly BESS and aggregated distributed resources—introduces competitive supply that may partly offset price increases, but also introduces new patterns of strategic bidding behaviour. Understanding and anticipating these variations is essential for market participants making long-term investment decisions in flexible assets.

1.3 Long-term Electricity Price Scenarios

Svenska Kraftnät (SvK) publishes Långsiktig Marknadsanalys (LMA) scenario frameworks that define long-term trajectories for key system variables including electricity demand, generation mix, reservoir levels, and wholesale spot prices to 2045 and beyond. The LMA is structured around distinct narrative scenarios that capture different plausible futures for the Swedish and Nordic energy system. The two primary scenarios used in this project are:

1. Elektrifiering planerbart (EP) – characterised by a more moderate penetration of variable renewable energy, with a greater role for controllable (plannable) generation including nuclear power.
2. Elektrifiering förnybart (EF) – characterised by a high penetration of variable renewable energy (wind), leading to greater price volatility and stronger seasonal patterns in spot prices.

Each LMA scenario encompasses multiple weather years, meaning electricity spot prices are represented as distributions rather than single deterministic trajectories. LMA2021 [2] was used as the primary scenario input for this project (see Preface for the rationale regarding LMA2021 vs. LMA2024).

1.4 Why a Long-Term Forecasting Model for FCR?

Accurate long-term price forecasts for FCR markets are critical for several categories of decision-makers. Asset owners—including hydro operators, battery investors, and industrial flexibility providers—require credible price scenarios over 10–30 year horizons to evaluate the business case for investments in reserve-capable assets. TSOs and regulators need scenario-based price insights to assess the adequacy of reserve procurement frameworks and to design market rules that remain fit for purpose as the power system transforms.

Existing approaches for long-term FCR price forecasting are limited. Fundamental equilibrium models capture structural drivers but struggle to represent the emergent market variations arising from strategic agent behaviour. Statistical extrapolation approaches are inherently backward-looking and unable to account for structural breaks. Agent-Based Modelling offers a complementary approach: by explicitly simulating the decision-making of market participants under changing conditions, ABMs can generate forward-looking price scenarios that reflect both the physical constraints of the power system and the strategic behaviour of agents—including new entrants such as BESS operators.

The project 'Long-term price scenarios of frequency regulation services' (project number P2022-00752) was initiated to address this gap, developing a coupled ABM–TFT modelling pipeline capable of generating FCR price scenarios consistent with SvK's official long-term market analysis scenarios to the year 2050.

2 Methodology

The methodology developed in this project combines two complementary modelling approaches: an Agent-Based Model (ABM) for simulating the Nordic day-ahead electricity spot market, and a Temporal Fusion Transformer (TFT) for translating spot market variations into long-term ancillary service price scenarios. This section provides an overview of the general Agent-Based modelling rationale and a brief summary of the workflow used in this project; detailed technical descriptions of each component are provided in Sections 3, 4, and the companion TFT Technical Report [1].

2.1 Review of Agent-Based Modelling

A systematic review of existing modelling tools was conducted to identify a suitable platform for simulating the reserve market and generating long-term FCR price time series. The review considered both optimisation-based equilibrium models and agent-based modelling (ABM) frameworks, evaluating them against the specific requirements of this project: the ability to represent individual market participants with heterogeneous bidding strategies, compatibility with evolving Nordic market rules and product definitions, support for scenario-driven simulation over long time horizons, and the flexibility to incorporate learning-based agent behaviour trained on historical data.

The review concluded that agent-based modelling is the most appropriate paradigm for this project's objectives. Unlike traditional optimisation or equilibrium approaches, which assume perfectly rational actors and produce a single cost-minimising dispatch, ABM represents individual decision-making by market participants and captures the emergent changes that arise from their interactions — including strategic bidding, imperfect information, and adaptive behaviour [4][5]. These properties are particularly relevant for reserve markets, where a small number of heterogeneous actors (hydropower producers, battery operators, aggregators, wind farms) interact under evolving rules, and where the relationship between spot prices and reserve bids involves strategic considerations that cannot be reduced to a static cost function. The general motivations for using agent computation in the social sciences, including electricity markets, have been articulated by Axtell[6], while the potential of ABMs for economic forecasting has been demonstrated by Poledna et al. [11].

2.1.1 ABM Selection

From a pool of 75 open-source energy models screened in the literature, three agent-based tools with explicit balancing market functionality were shortlisted: ASAM (Ancillary Services Acquisition Model), AMIRIS [7], and Elba-ABM (Agent-Based Model of ELectricity BAlancing market)[8]. In addition, a set of equilibrium-based tools was reviewed but found to be less suited to the project's needs due to their limited representation of agent heterogeneity and strategic behaviour.

ASAM is an open-source model built on the Mesa ABM framework [10]and PyPSA (Python for Power System Analysis). It simulates day-ahead and intraday markets with agent-level bidding, portfolio dispatch with inter-temporal constraints, and zonal pricing. However, the current version of ASAM does not include balancing capacity markets, does not cover the Swedish or Nordic market context, and would require substantial extension to represent FCR products and their specific technical requirements.

AMIRIS is a Java-based ABM (with a Python wrapper) running on the FAME framework [7]. It was developed to assess how policy frameworks affect the profitability and behaviour of energy market actors and has been applied to study battery storage revenues in day-ahead and aFRR markets under high-renewables scenarios. However, AMIRIS operates at hourly resolution, which limits its ability to capture the sub-hourly variations relevant to FCR. The current version focuses on the German market and day-ahead trading; FCR and intraday markets are not implemented and have been identified by the developers as areas for future work. Additional limitations include the absence of demand response and Power-to-X modelling, perfect-foresight assumptions in the optimisation, and a lack of competition modelling between flexibility providers.

Elba-ABM is a bottom-up agent-based model specifically designed for European balancing markets. Among the three tools, it is the closest to the project's scope: it models hourly FCR-N prices, total demand, offered volumes and system costs, and supports both marginal pricing and pay-as-bid rules. Notably, Elba-ABM has been applied to the Swedish FCR market, where it demonstrated that the entry of new technologies (battery storage and wind) reduces system costs and stabilises market outcomes, and that strategic bidding behaviour under different pricing rules can substantially affect total system costs. However, Elba-ABM is a commercial tool with restricted access, which limits its suitability for an open research project. Furthermore, adapting it to the full set of FCR products (FCR-N, FCR-D up, FCR-D down) with the updated Nordic technical requirements and integrating it with the TFT forecasting chain and the LMA2024 scenario framework would require close collaboration with the tool's developers that was not feasible within the project's timeline and resources.

Given that none of the existing tools could be adopted without major modifications, the project chose to develop a custom agent-based model built on the Mesa framework in Python. Mesa was selected as the underlying platform for several reasons. It is the most widely used open-source ABM framework in the scientific computing ecosystem, with over 500 publications and 800 contributing authors since its initial release in 2014 [9][10]. It provides built-in core components (agent schedulers, data collection, visualisation) while allowing full customisation of agent logic, market clearing rules and model structure. Its Python-native architecture ensures seamless integration with the scientific computing stack (NumPy, pandas, scikit-learn) used for data handling, neural network training and result analysis in this project. Furthermore, ASAM's use of Mesa demonstrated that the framework is suitable for electricity market applications and can be extended with market-specific modules.

2.1.2 Mesa-based ABM

The custom Mesa-based ABM developed in this project represents the Nordic day-ahead market through 11 meta-agents, each equipped with a dedicated neural-network bidding function trained on historical Nordic bidding curve data from 2019. This approach — learning agent behaviour from observed data rather than prescribing cost-based or rule-based strategies — was chosen because comprehensive data on individual market participants' costs and strategies is not available, and because a data-driven approach allows the model to capture the revealed competitive effects and strategic patterns embedded in historical bids. The concept of learning meta-agent behaviour from observed data has been similarly applied in other domains, for example by Tian et al. for household electricity consumption [12]. The model structure (main.py, agents.py, model.py, training_data.py), training methodology, input features (date, marginal price, reservoir levels, energy mix), and the rationale for using meta-agents are documented in detail in Section 4.

2.2 Coupled ABM–TFT Pipeline

The overall forecasting pipeline consists of two sequentially coupled models. In the first stage, the ABM simulates hourly Nordic spot prices for a specified future period, drawing on scenario inputs (demand, reservoir levels, energy mix) from the LMA2021 scenario framework [2]. In the second stage, the forecasted spot prices—together with other scenario variables—are fed as inputs to the TFT model, which generates price scenarios for FCR-N, FCR-D up, and FCR-D down.

This coupled approach leverages the strengths of both methods: the ABM's capacity to generate internally consistent spot price trajectories reflecting agent behaviour under changing market conditions, and the TFT's data-driven ability to learn the empirical relationships between spot market variations and ancillary service prices. By anchoring both models to SvK's LMA2021 scenario inputs, the resulting FCR price scenarios are consistent with the official long-term system development outlook.

3 Input data

3.1 Training Data for the ABM

The ABM's meta-agents learn their bidding behaviour from historical Nordic spot market data. The training dataset is structured around observations from the year 2019, chosen to represent typical pre-pandemic market behaviour in the Nordics—a period free from the COVID-19 demand disruption that characterises 2020 and subsequent years.

3.1.1 Input Features

Each meta-agent's neural network is trained on the following input features:

1. Exact date and time of the hour to predict (hour, day, month)
2. Marginal day-ahead spot price observed in the previous hour
3. Reservoir levels in SE1 and SE2 observed in the previous hour (sourced from LMA2021 for scenario periods)
4. Energy production disaggregated by source: nuclear, hydro, renewable, and thermal

The response (output) dataset consists of observed bidding curves for the Nordic spot market during 2019, specifically the observed bid volumes and prices at each hour. The bidding curve data was obtained from Nord Pool.

3.1.2 Scenario Input Data (LMA2021)

For the long-term forecast periods, scenario-driven inputs replace historical observations. Future demand trajectories, reservoir levels in SE1 and SE2, and energy production mix are sourced from the LMA2021 EP and EF scenarios published by Svenska Kraftnät. These inputs are used both to drive the demand side of the ABM market clearing (via model.py) and to provide the reservoir level inputs to the agent bidding functions.

3.2 Data Handling

3.2.1 Data Cleaning

The dataset underwent standard cleaning procedures to remove inconsistencies and missing values. Where short gaps in time series data existed, linear interpolation was applied. Outlier bids in the training data were reviewed; the minimum observed bid in the training data was −500 EUR/MWh, which was retained as the lower bound for Agent 0 representing must-run capacity forced into the market.

3.2.2 Data Adjustments and Normalisation

Feature engineering was limited to scaling: all input variables were normalised to a uniform range prior to neural network training, which is a standard requirement for stable MLP regressor performance. No additional feature transformations were applied.

The ABM operates at Nordic-system level for spot price generation but adjusts outputs to individual Swedish pricing zones (SE1–SE4) using price zone factors derived from historical data. These factors represent the historic spread between each pricing zone and the Nordic average, computed as hourly averages across the years 2016–2023.

4 Model

4.1 Structure of the ABM

The ABM is implemented in Python using the Mesa agent-based modelling framework and consists of three main modules: main.py (steering and orchestration), agents.py (agent behaviour and bidding), and model.py (market clearing). These modules interact in a simulation loop that iterates through each hour of the forecast period.

4.1.1 main.py – Steering Module

main.py serves as the entry point for the simulation. It accepts user-defined parameters (forecast time frame, number of agents, scenario inputs) and orchestrates the simulation loop by alternately calling functions in model.py and agents.py. After simulation, main.py handles output processing—generating visualisations via pyplot or writing results to CSV/Parquet files. It also applies the price zone factors to translate the Nordic average spot price output from the ABM into zone-specific prices for SE1–SE4.

4.1.2 agents.py – Agent Behaviour

agents.py implements the bidding behaviour of N=11 meta-agents that collectively simulate the Nordic spot market bidding curve. Agent 0 represents the bulk of must-run capacity, bidding at a fixed price of −500 EUR/MWh. Agents 1–10 represent aggregated bids across different segments of the bidding curve. Each agent's bid (price and volume) is determined by a feedforward neural network (MLP regressor) trained on historical bidding curve data. Agent bids are capped at a minimum of −500 EUR/MWh and a maximum volume of 20,000 MWh (~80% of minimum observed demand) to prevent extreme bidding behaviour and convergence issues.

4.1.3 model.py – Market Clearing

model.py simulates market clearing based on agent bids and an external demand forecast drawn from the LMA2021 EP or EF scenarios [2]. The clearing mechanism determines the marginal spot price for each simulated hour. The resulting marginal price is fed back as an input to agents.py for the subsequent hour, creating the autoregressive loop that characterises ABM simulation. The following Figure 1 illustrates both, the training and the usage of the ABM.

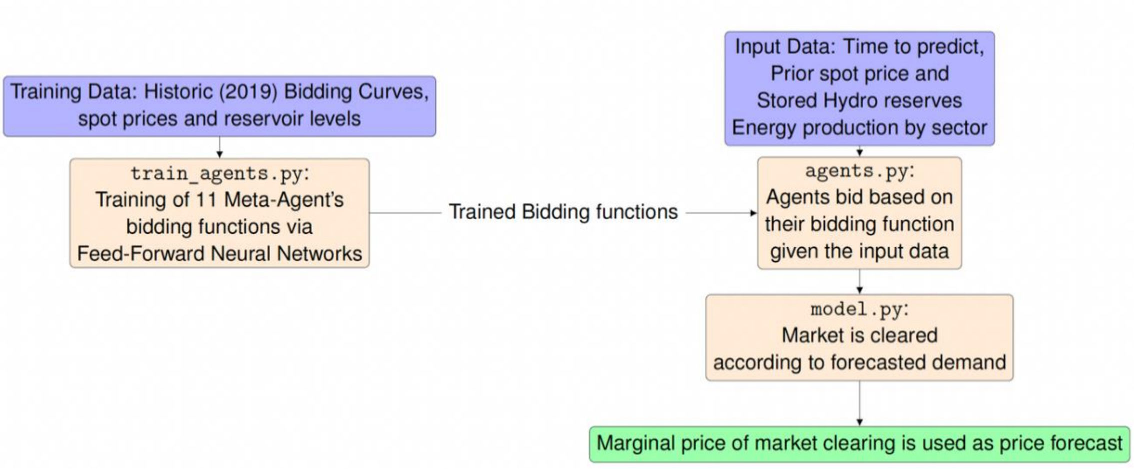

Figure 1: On the left-hand side the training procedure of the meta-agents is illustrated, whereas on the right-hand side the ABM loop is described to forecast the Nordic spot prices.

It is worth noting that the individual meta-agent's bids consist of a bid volume and price, which is then cleared by model.py. The resulting marginal price then again serves as an input for the ABM to forecast the spot price for the next time period, i.e., the ABM loops through the steps illustrated on the right-hand side of Figure 1.

4.2 Learning Agent Behaviour

Accurately modelling agent behaviour is critical for reliable ABM-based forecasting. In this implementation, each agent's bidding function is represented by a separate feedforward artificial neural network (MLP regressor, scikit-learn), trained on observed market behaviour from 2019. This meta-agent concept—where a small number of learned agents collectively represent the aggregate behaviour of many individual market participants—has precedent in the literature for applications where individual agent-level data is unavailable [12].

Each of the 11 neural networks is characterised by:

1. Network architecture: number of hidden layers and neurons per layer, tuned per agent.
2. Training data: historical marginal prices, reservoir levels, and energy mix from 2019, with corresponding bidding curve observations as targets.
3. Hyperparameters: regularisation penalty and optimiser (Adam solver used throughout).

4.3 Training Process

The model was trained on a standard notebook computer, with each agent's neural network converging in under 10 minutes. The primary evaluation metric used was the prediction score (R^2 coefficient of determination from Python's MLPRegressor), which measures the proportion of variance in the target variable explained by the model. After hyperparameter tuning, final R^2 scores for each agent are reported in the Appendix.

4.4 Coupling of ABM and TFT Models

After the ABM generates hourly Nordic spot price forecasts for SE1–SE4, these are used as inputs to the Temporal Fusion Transformer (TFT) model. The TFT exploits correlations between the forecasted spot prices, energy production mix, and hydropower reserves to derive price scenarios for FCR-N, FCR-D up, and FCR-D down. The following Figure 2 illustrates the coupling of the ABM and the TFT model.

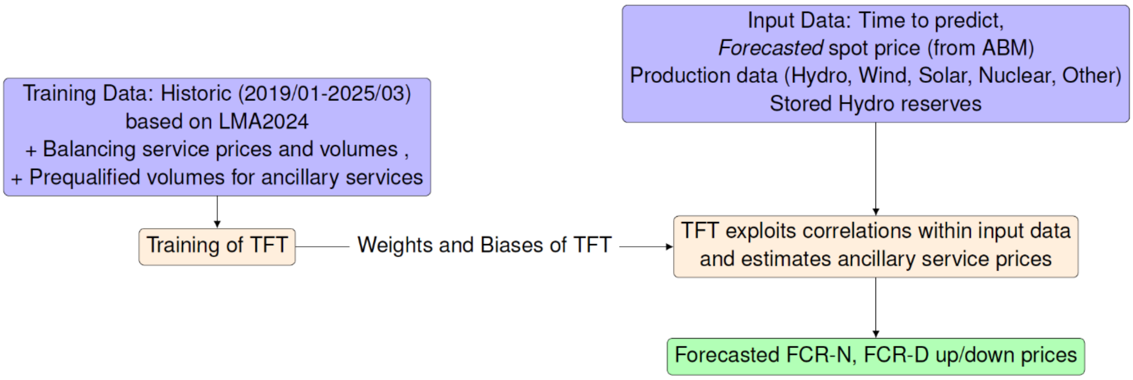

Figure 2: Illustration of TFT model. The spot price forecasts by the ABM are inserted as input data to the TFT model (see blue box on the top right).

The ABM forecasts spot prices at time t+1 based on inputs available at t; the TFT then uses the forecasted inputs at t+1 to derive ancillary service prices at t+1. For further details on the TFT model architecture, training, and evaluation, refer to the TFT Technical Report [1].

4.5 Limitations

The ABM has several structural limitations that bound the scope of its application:

1. Temporal resolution: The ABM is limited to hourly resolution, reflecting the historical market clearing rate. This limits accuracy for modelling sub-hourly ancillary service effects.
2. Technology coverage: Demand response, Power-to-X, and large-scale EV/BESS penetration are not explicitly modelled as separate agents. Their market impact is implicitly captured in the training data to the extent they were present in 2019 market behaviour.
3. External long-term impacts: Political impacts—taxes, tariffs, levies—are excluded. Degradation effects are implicitly present through scenario inputs but not explicitly modelled.
4. Price zone spreads: The ABM models SE1–SE4 price differences using historical seasonal spread factors, which is an acknowledged simplification of the complex physical and commercial drivers of zonal price divergence.

5 Results

In our analysis, we draw on scenario data for 2045, focusing on the two scenarios “Elektrifiering planerbart” (EP) and “Elektrifiering förnybart” (EF) as outlined in [2]. Each scenario encompasses multiple weather years, which means electricity spot prices are represented as a distribution rather than a fixed hourly value. As a benchmark we use the mean spot price over those weather years.

The main result of the analysis is that the forecasts of spot prices simulated by the ABM closely track the benchmark forecasts from [2] on a weekly average basis. The ABM however seems to produce significantly more stable forecasts over most of the 4 Swedish pricing zones in the EF scenario and similarly stable results in the EP scenario. The future demand in the market as well as the levels of the main water reservoirs in SE1 and SE2 in both approaches are based on [2].

For the remaining analysis of ancillary service prices, we compare the outputs given by the TFT, where on the one hand the electricity spot price trajectories from LMA2021 serve as known inputs and on the other hand, the ABM outputs serve as the inputs for the TFT. By aligning our models with these scenario-based price paths, we ensure that forecasts for ancillary services—such as FCR-N, FCR-D up, and FCR-D down—are consistent with the underlying assumptions of SvK’s long-term market analysis.

5.1 Spot Price Forecasts: ABM vs. LMA2021

The ABM closely tracks the weekly-average spot price trajectories given in LMA2021 in both scenarios. A key observation is that the ABM produces significantly more stable forecasts than LMA2021 in most pricing zones, reflected in substantially lower forecast variance. The ABM displays a clear seasonal pattern with higher spot prices in winter and lower prices in summer. Unlike LMA2021, the ABM does not generate severe outlier prices above ~100 EUR/MWh, while LMA2021 forecasts contain rare events exceeding 200 EUR/MWh (and even 500 EUR/MWh in extreme cases).

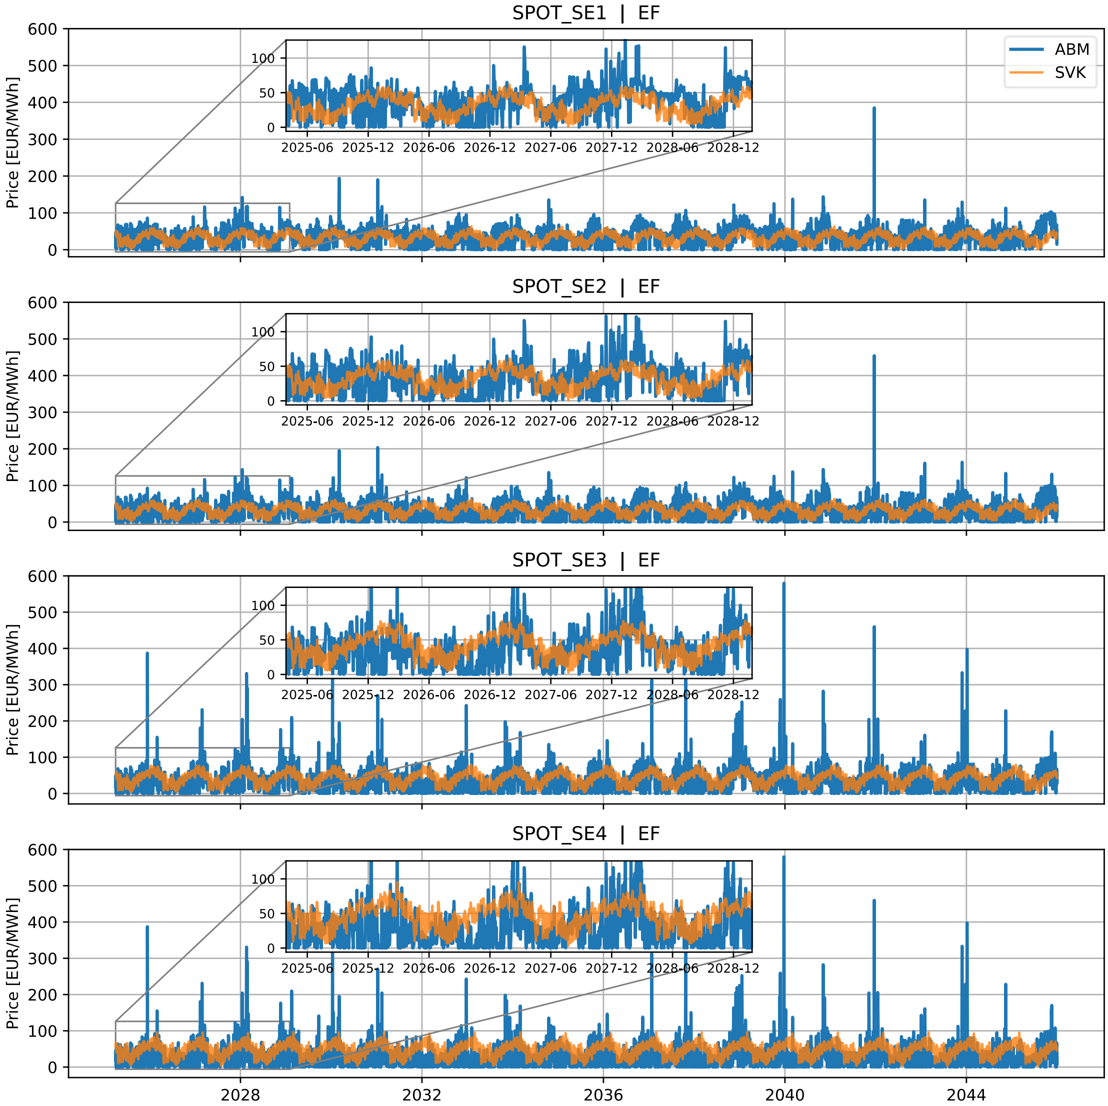

Figure 3: Spot price forecasts for SE1-SE4 in Scenario EF by SVK in [2] in blue and the ABM in orange shown on a daily resolution. Both forecasts are based on forecasted demands, reservoir levels in SE1 and SE2 from [2].

In the EF scenario, the statistical comparison between ABM and LMA2021 forecasts for each pricing zone is summarised in Table 1. The average prices level of the ABM is slightly lower in SE1 compared to LMA2021, about equal in SE2 and sufficiently higher in SE3 and SE4. The latter might be attributed to the fact the SVK expects that investment in transfer capacity between north and south will result in a changing power balance in the different price zones, which will cause the price differences first to equalise and then to reverse. Such overall trends are not captured by the ABM, but in this particular case could be implemented by changing the price spreads for the respective pricing zones in main.py.

EF Scenario	SE1 ABM	SE1 SVK	SE2 ABM	SE2 SVK	SE3 ABM	SE3 SVK	SE4 ABM	SE4 SVK
Mean [EUR/MWh]	32.72	38.37	32.91	31.89	43.69	36.58	49.55	28.87

EF Scenario	SE1 ABM	SE1 SVK	SE2 ABM	SE2 SVK	SE3 ABM	SE3 SVK	SE4 ABM	SE4 SVK
Variance [(EUR/MWh)²]	226.9	700.9	224.5	786.3	426.5	1459	499.1	1512
MSE [(EUR/MWh)²]	913.6		935.2		1721		2304	

Table 1: Statistical parameters for EF scenario spot price forecasts (ABM vs. SVK/LMA2021). MSE indicates tracking error between the two forecasts.

In the more conservative EP scenario, a lower penetration of variable renewable energy is assumed. The ABM forecast is very similar to the EF scenario, whereas the LMA2021 forecast is significantly more stable in EP, leading in some zones to lower LMA2021 variance than the ABM. The difference in mean spot price between the two approaches widens in SE2 and SE3 under EP.

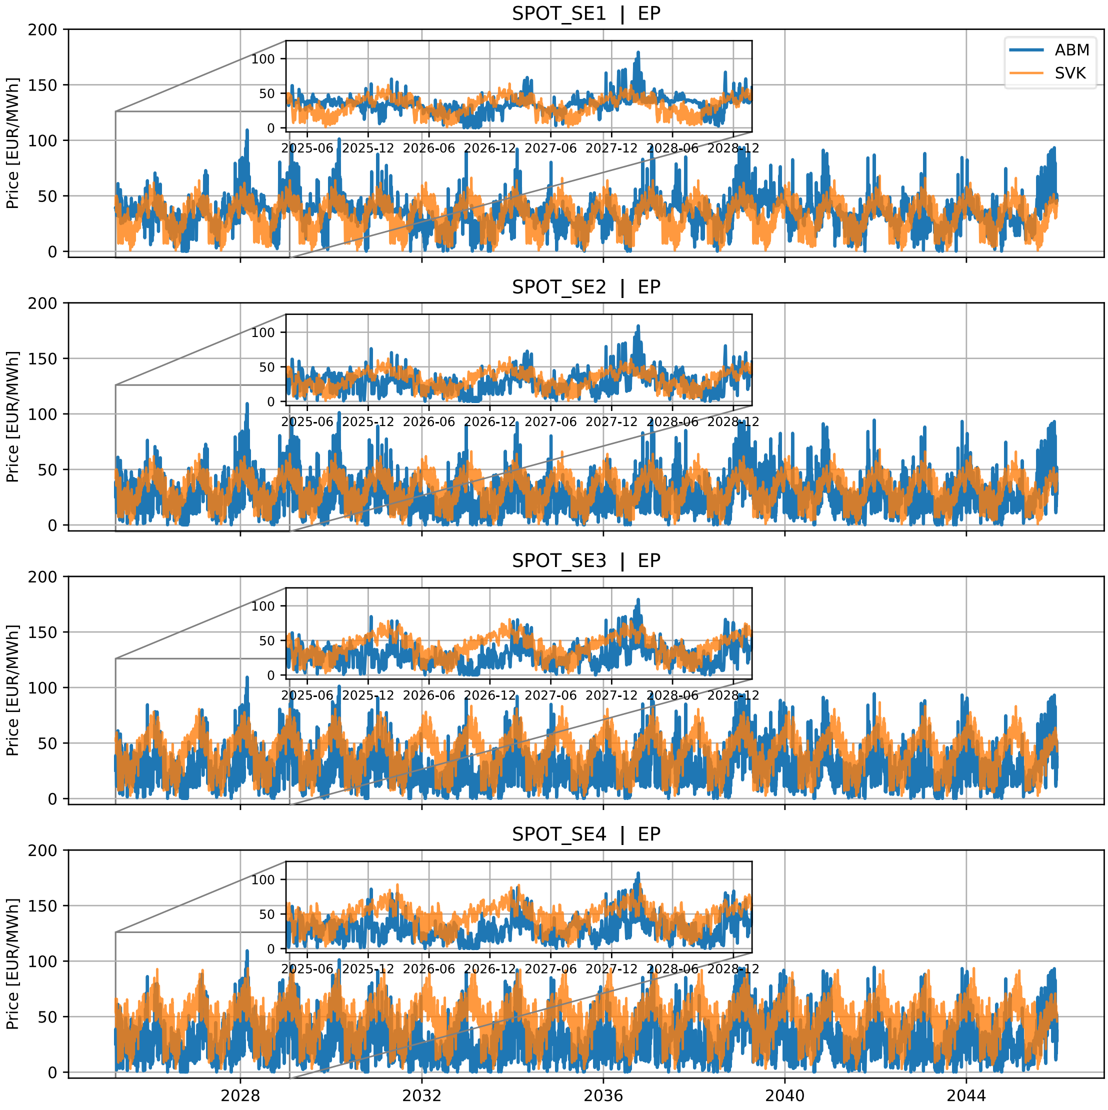

Figure 4: Spot price forecasts for SE1-SE4 in Scenario EP by SVK in [2] in blue and the ABM in orange shown on a daily resolution. Both forecasts are based on forecasted demands, reservoir levels in SE1 and SE2 and energy mix from [2].

As the mean spot prices for the pricing zones in the EP scenario are considered to be significantly lower in LMA2021, the difference in spot prices compared to the ABM forecast widens, particularly in SE2 and SE3, see Table 2 below. The tracking error, measured via MSE, between the two forecasts however is reduced as well, which appears to be caused by the drastic reduction in variance of the LMA2021 forecast.

EP Scenario	SE1 ABM	SE1 SVK	SE2 ABM	SE2 SVK	SE3 ABM	SE3 SVK	SE4 ABM	SE4 SVK
Mean [EUR/MWh]	32.61	34.02	32.80	27.91	43.51	28.47	49.36	28.72
Variance [(EUR/MWh)^2]	221.2	185.0	219.0	283.6	414.4	311.8	487.0	341.9
MSE [(EUR/MWh)^2]	371.5		463.4		846.5		1149	

Table 2: Statistical parameters for EP scenario spot price forecasts (ABM vs. SVK/LMA2021).

Lastly, it is worth pointing out that the consistency in the ABM forecasts may very well be caused by the cutting behaviour described in agents.py. However, it may also be caused by the fact that the energy mix seems only to have a minimal influence on the ABM forecast overall as we observe a very similar behaviour of the ABM also in the more low-variance EP scenario.

5.2 FCR-D down Forecasts

FCR-D down forecasts were generated using TFT (Model 1), which achieved the highest composite performance score for FCR-D. Both ABM-supported and LMA2021-supported TFT inputs produce very closely aligned FCR-D down forecasts in the EF scenario, with a Mean Squared Error of 2.53 EUR/MW between the two approaches. This suggests that the choice of spot price input source has limited impact on FCR-D down forecasts in the EF scenario, likely because direct TFT inputs (energy mix, hydro reserves) dominate the forecast signal.

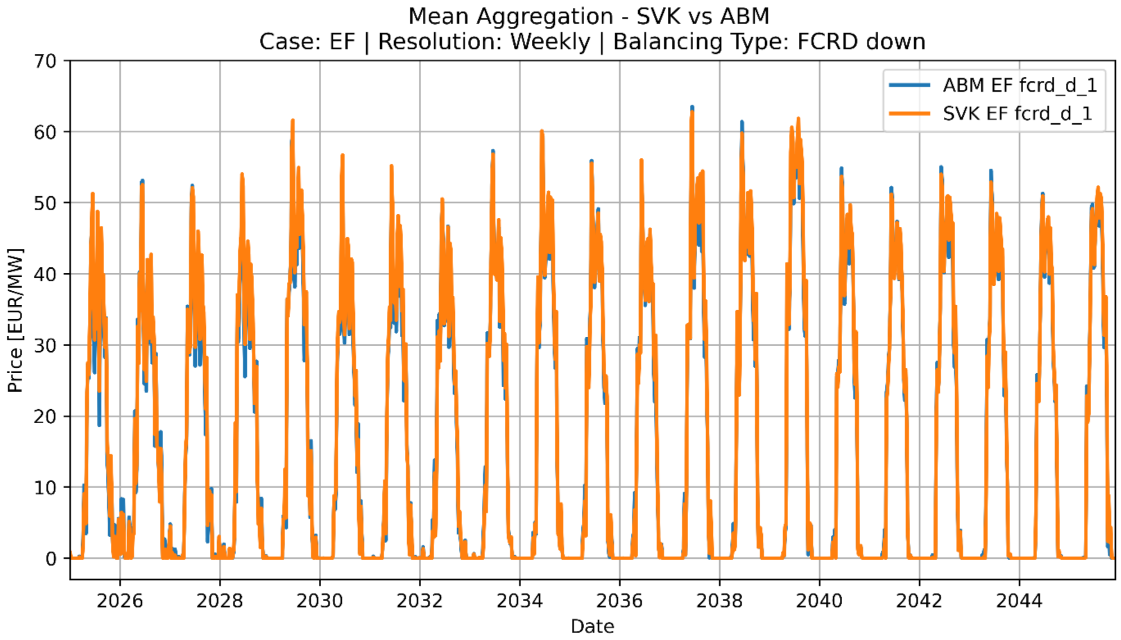

Figure 5: EF Scenario comparison ABM vs LMA2021 input data translated to FCR-D down forecasts via TFT Model 1. Both forecasts are calculated on raw daily output data of TFT Model 1.

One observes that although the fit in the spot price comparison between the two spot price forecasts was not significant, the alignment between the resulting FCR-D-down forecasts is very close. Consequently, it appears that the choice of spot prices overall has no major influence in the EF scenario. This effect might be explained by a reduced influence of the spot prices compared to other input features to the TFT, such as the energy distribution and hydropower reserves, that vary significantly between the scenarios. These direct influences appear to dominate particularly in the EF scenario, which captures a change from the usual patterns observed until 2024. The key metrics support the overall close alignment between the two spot price forecasting techniques with a notably low Mean-Square Error of 2.53.

FCR-D down EF (Model 1)	ABM	SVK/LMA2021
Mean [EUR/MW]	16.7	16.5
Variance [(EUR/MW)²]	408.4	396.7
MSE [(EUR/MW)²]	2.53	

Table 3: Statistical parameters for FCR-D down forecasts in EF scenario.

In the EP scenario, the two approaches diverge more noticeably, with the ABM-supported TFT forecasting a higher mean FCR-D down price (12.3 vs. 9.9 EUR/MWh). This supports the interpretation that spot prices play a more prominent role in the EP scenario, where direct input features (energy mix, hydro reserves) show less extreme variation.

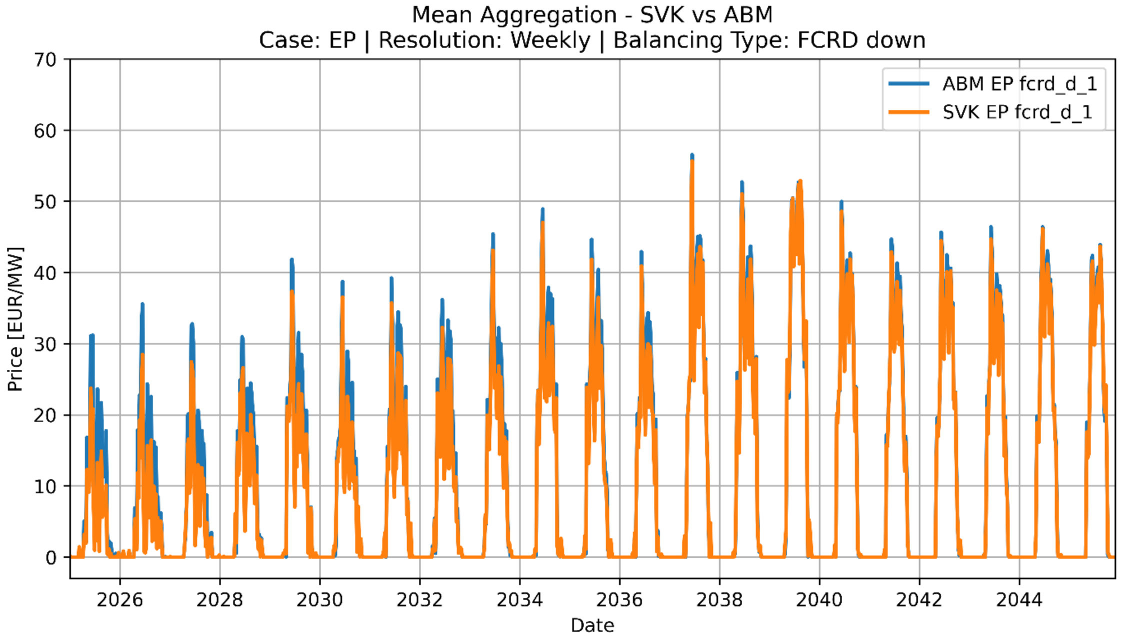

Figure 6: EP Scenario comparison ABM vs LMA2021 input data translated to FCR-D down forecasts via TFT Model 1. Both forecasts are calculated on raw daily output data of TFT Model 1.

The key parameters are again summarized in the following Table:

FCR-D down EP (Model 1)	ABM	SVK/LMA2021
Mean [EUR/MWh]	12.3	9.9
Variance [(EUR/MWh)²]	278.7	206.6
MSE [(EUR/MWh)²]	6.87	

Table 4: Statistical parameters for FCR-D down forecasts in EP scenario.

5.3 FCR-D up Forecasts

FCR-D up forecasts were generated using TFT Model 11. Under the EF scenario, the ABM-supported TFT produces a significantly more reliable and consistent forecast (variance 4.49 vs. 12.10 EUR/MW²), avoiding the outlier events forecasted by the LMA2021-supported TFT in late 2037, 2039, 2041, and 2043—periods where LMA2021 inputs lead the TFT to forecast weekly prices exceeding 20–30 EUR/MW.

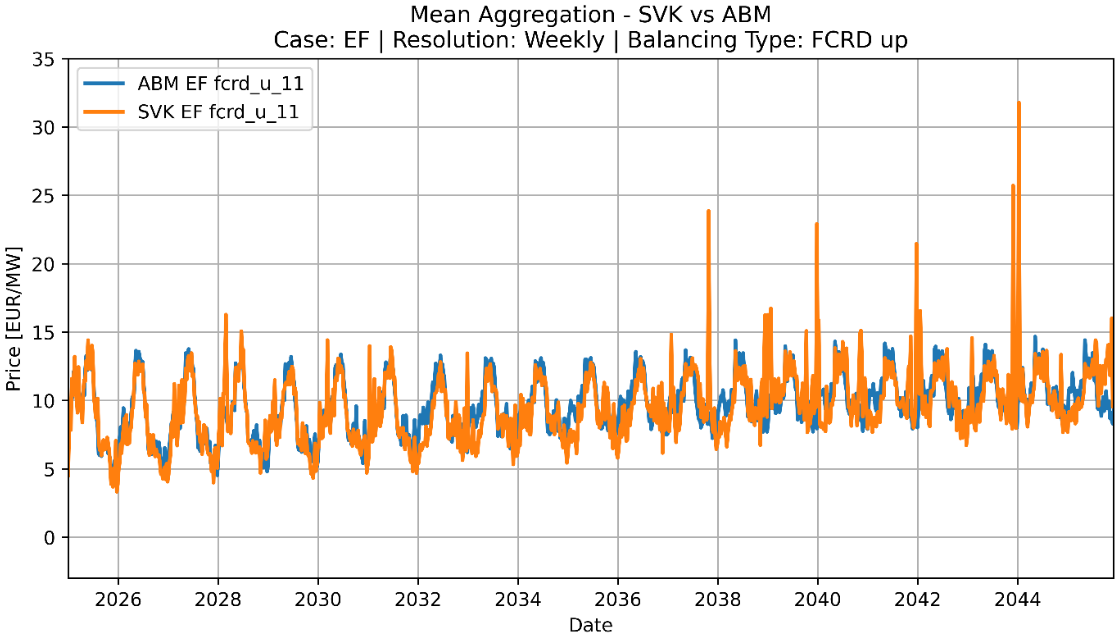

Figure 7: EF Scenario comparison ABM vs LMA2021 input data translated to FCR-D up forecasts via TFT Model 11. Both forecasts are calculated on raw daily output data of TFT Model 11.

Observe that the overall trends align well up to certain outliers by LMA2021 in late 2037, 2039, 2041 and 2043, where weekly FCR-D up prices are forecasted to surpass 20 or even 30 EUR per MW and even a smaller, but early outlier in early 2028 of more than 15 EUR/MW. These results indicate that the ABM price building mechanism is slightly more in line with the expected input data for the TFT, which is expected as the ABM does not pivot too far from previously observed training data. These outliers and a more reliable prediction for FCR-D up are also supported by a significantly smaller variance of the ABM-supported TFT of 4.49 vs 12.10 (EUR/MWh)^2 for the LMA2021-supported TFT.

FCR-D up EF (Model 11)	ABM	SVK/LMA2021
Mean [EUR/MW]	10.0	9.64
Variance [(EUR/MW)^2]	4.49	12.10
MSE [(EUR/MW)^2]	7.76	

Table 5: Statistical parameters for FCR-D up forecasts in EF scenario.

In the EP scenario, the difference in forecast variability is less pronounced (variances of 6.78 vs. 9.03), and the tracking error between the two approaches drops to an MSE below 2. The ABM-supported TFT still forecasts a slightly higher mean price, particularly in the years after 2040.

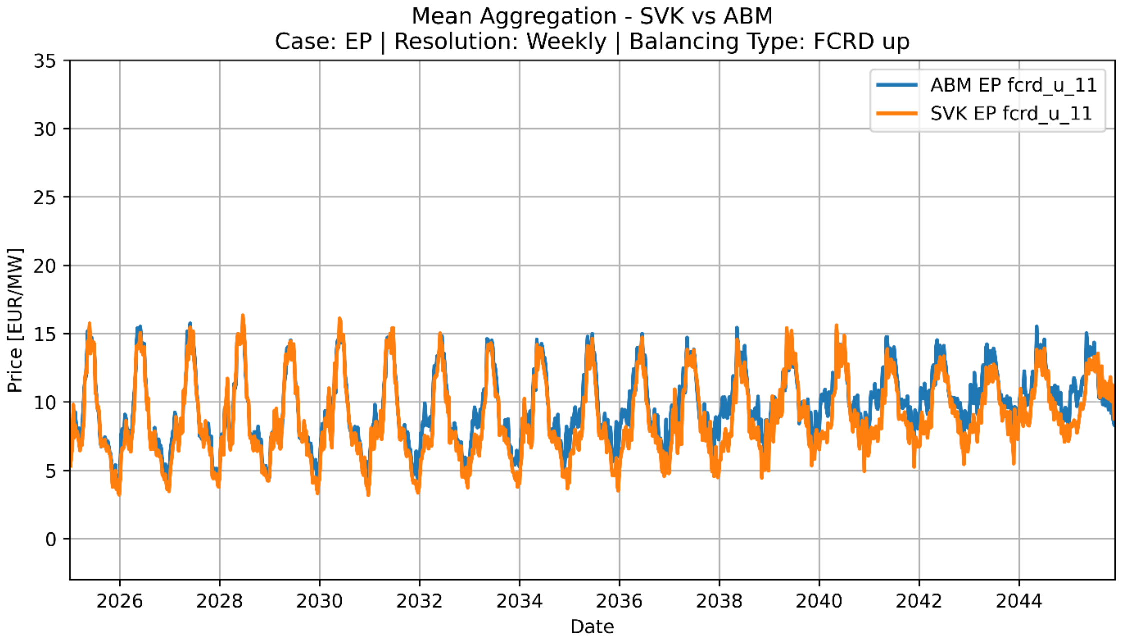

Figure 8: EP Scenario comparison ABM vs LMA2021 input data translated to FCR-D up forecasts via TFT Model 11. Both forecasts are calculated on raw daily output data of TFT Model 11.

However, the overall tracking behaviour between the two forecasts is the strongest observed in our scenarios with an MSE below 2, which however may be caused due to the significant lower price range of 5-15 EUR of FCR-D up compared to the other two ancillary services and the lack of outliers in the standard TFT. It is further worth noting that in contrast to the EF scenario, the mean FCR-D up price forecasts diverge slightly with the ABM supported TFT predicting a slightly higher price, particularly in the years after 2040.

FCR-D up EP (Model 11)	ABM	SVK/LMA2021
Mean [EUR/MW]	10.22	8.87
Variance [(EUR/MW)²]	6.78	9.03
MSE [(EUR/MW)²]	1.94	

Table 6: Statistical parameters of FCR-D up forecasts in EP scenario for ABM and LMA2021-based input spot prices.

5.4 FCR-N Forecasts

FCR-N forecasts were generated using TFT Model 6. FCR-N shows the closest alignment between the ABM-supported and LMA2021-supported TFT approaches across both scenarios. The mean price levels and variances are well-matched in both EF and EP scenarios. Figure 9 below presents the FCR-N forecast results for both models under the EF scenario.

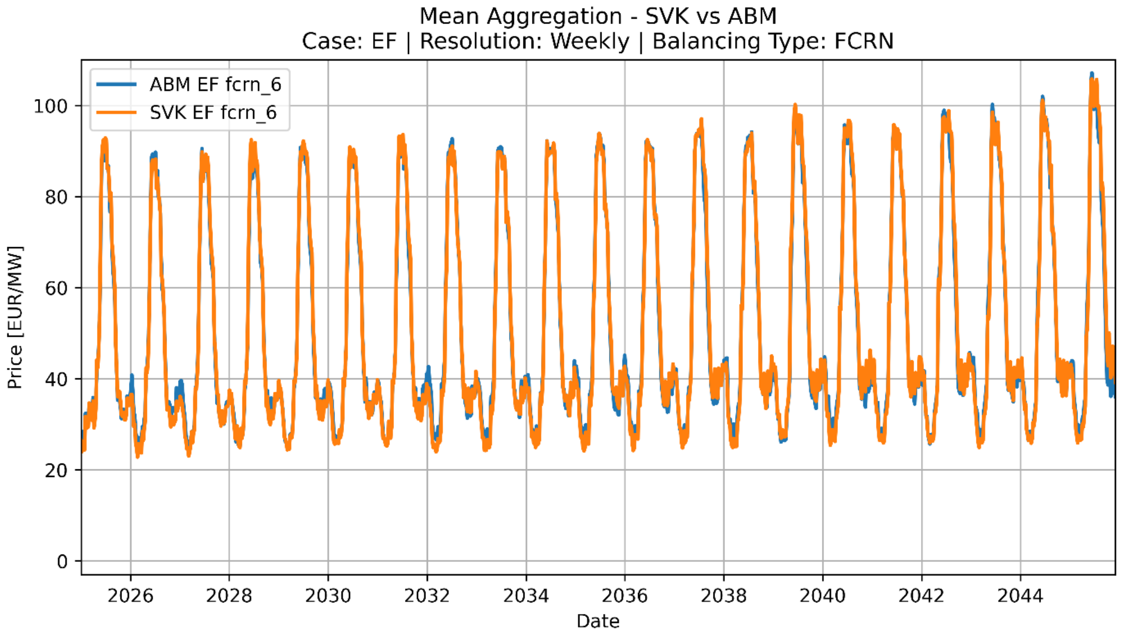

Figure 9: EF Scenario comparison ABM vs LMA2021 input data translated to FCR-N forecasts via TFT Model 6. Both forecasts are calculated on raw daily output data of TFT Model 6.

We note that the EP scenario yields similar results as is illustrated in Figure 10 below.

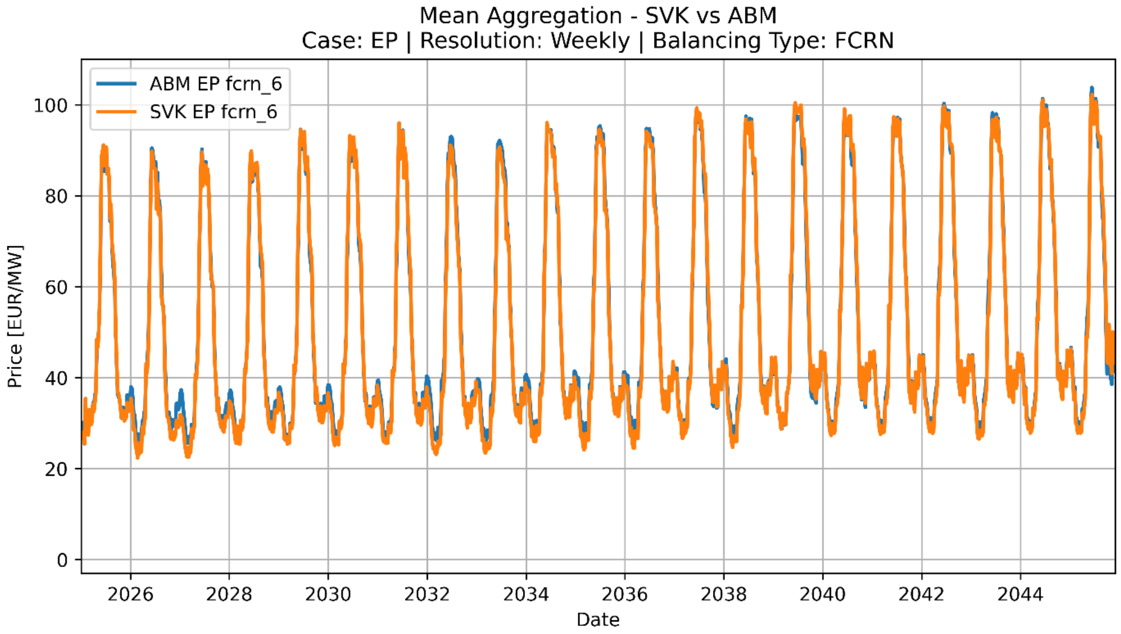

Figure 10: EP Scenario comparison ABM vs LMA2021 input data translated to FCR-N forecasts via TFT Model 6. Both forecasts are calculated on raw daily output data of TFT Model 6.

In general, FCR-N seems to be the most aligned of all considered ancillary services. This may be due to the fact that FCR-N maintains frequency stability during normal conditions, while FCR-D is a safeguard against extreme frequency deviations caused by major disturbances. Hence, FCR-N is expected to be more strongly correlated to spot prices and the overall system circumstances, i.e., energy mix and reserves, and thus the considered features should be able to capture the main behavior of FCR-N prices better than for FCR-D prices. The overall slightly higher MSE of approximately 6.9-97.8

(EUR/MW)² compared to the previous FCR-D forecasts might be explained by the fact that FCR-N does not share the natural lower baseline of zero characteristic of FCR-D down, and spans a significantly wider price range than FCR-D up. The other key metrics considered, i.e. mean price level and variance are very much aligned for FCR-N forecasts as can be seen in Table 7.

FCR-N Scenario (Model 6)	ABM Mean	SVK Mean	ABM Var	SVK Var	MSE
EF	53.4	51.0	653.4	578.0	7.82
EP	54.1	50.9	648.5	591.4	6.91

Table 7: Statistical parameters for FCR-N forecasts in EF and EP scenarios (mean and variance in EUR/MW and (EUR/MW)² respectively).

6 Discussion

6.1 Input Data and Methodological Considerations

The training of both the ABM meta-agents and the TFT model relies on historical market data from the Nordic electricity and ancillary service markets. This creates an inherent dependency on the representativeness of the training period relative to future market conditions. The use of 2019 as the primary ABM training year was deliberate—it represents a year of relatively normal pre-pandemic Nordic market behaviour—but it also means that the model does not learn from the significant price variations of 2021–2023, which were characterised by energy crisis conditions and extreme price spikes.

The reliance on LMA2021 scenario inputs for reservoir levels, demand trajectories, and energy mix introduces a structural alignment between the ABM and the broader LMA2021 scenario framework. While this ensures consistency with SvK's official outlook, it also means that the ABM's spot price forecasts are bounded by the LMA2021 scenario assumptions and cannot depart significantly from them.

6.2 Behaviour of Forecasts

A consistent finding across all results is that the ABM-supported TFT produces more reliable forecasts with lower variance and fewer extreme outliers compared to the LMA2021-supported TFT. This behaviour is particularly pronounced for FCR-D up, where the LMA2021-driven TFT generates several multi-year price spike episodes that are absent in the ABM-supported forecast.

This price-dampening behavior likely reflects two aspects of the ABM design: first, the hard cap on agent bid volumes (20000 MWh) prevents the market from clearing at extreme prices; second, the ABM's reliance on 2019 training data tends to anchor forecasts near historically observed behaviour, limiting extrapolation to scenarios far outside the training distribution.

The ABM also does not capture the expected convergence of SE3 and SE4 spot prices toward the Nordic average that SvK anticipates following investment in north-south transmission capacity. As a result, the ABM consistently forecasts higher spot prices in SE3 and SE4 compared to LMA2021, which propagates through to slightly higher ancillary service prices in the ABM-supported TFT outputs.

6.3 Interpretation

The coupling of the ABM for spot prices with the TFT-based translation provides a versatile forecasting tool that allows for straightforward scenario adaptation while tracking SvK's expert projections reasonably well under shared input assumptions. Across all tested scenarios and ancillary service products, the ABM-supported pipeline consistently forecasts slightly higher price levels than the LMA2021-supported TFT. This is largely attributable to the ABM's treatment of SE3 and SE4 pricing: because the ABM applies static historical zone-spread factors and does not incorporate SvK's expectation of increasing north–south transmission capacity, it does not replicate the projected long-run convergence of SE3/SE4 prices toward the Nordic average. The result is a persistently higher spot price in these southern zones, which propagates into higher ancillary service prices through the TFT.

A notable strength of the coupled approach is its robustness against extreme forecast events. In the EF scenario, the LMA2021-driven TFT generates several multi-year FCR-D up price spike episodes—periods where weekly average prices exceed 20–30 EUR/MW—that do not appear in the ABM-supported forecast. This smoothing effect arises because the ABM's bidding agents were trained on 2019 data and are unlikely to generate spot price trajectories far outside the range of the training distribution; the hard cap on agent bid volumes (20000 MWh) further prevents extreme clearing prices. Whether this conservatism is a strength or a weakness depends on the application: for risk-assessment purposes, the LMA2021-supported TFT may be more appropriate, while for central-estimate long-run planning, the ABM-supported pipeline offers a more tractable baseline.

Among the three FCR products, FCR-N shows the tightest alignment between the two approaches. This is consistent with the expectation that FCR-N—activated during normal frequency deviations under ongoing system conditions—is more strongly correlated with spot market variables (energy mix, reserves) that are shared as direct inputs to the TFT in both pipeline variants. FCR-D products, being triggered by disturbance events, are more sensitive to the specific trajectory of spot prices and exhibit greater divergence between the two approaches, especially under the EP scenario where the stabilising direct inputs (extreme energy mix shifts, low reservoir levels) are less pronounced.

6.4 Limitations and Directions for Future Work

Several limitations are identified for the current version of the ABM–TFT pipeline:

1. The most significant source of uncertainty is error propagation: errors introduced in the ABM spot price forecast propagate into the TFT ancillary service price forecast. While the tested scenarios showed limited impact, this risk grows for scenarios that diverge significantly from the training data.
2. Climate change effects on reservoir levels—through altered precipitation patterns and more frequent extreme weather—are not explicitly modelled. The reservoir inputs are anchored to LMA2021 projections, which themselves embody a set of climate assumptions.
3. The model does not account for future technologies such as large-scale hydrogen production, nuclear fusion, or widespread Power-to-X deployment.
4. No severe computational constraints exist: both models run on standard hardware without high-performance computing requirements.

Directions for future work include:

a) Retraining the ABM on updated data incorporating post-2020 market behaviour, and integrating LMA2024 scenario inputs.
b) Training the ABM directly on ancillary service market bidding curves, which would eliminate the need for the TFT translation step and reduce error propagation.
c) Adding dedicated meta-agents for BESS operators and aggregated demand response, enabling explicit simulation of these growing market participants.
d) Implementing additional post-processing procedures to reflect expected long-run trends in SE1–SE4 price spreads as grid integration increases.

7 References

[1] G. Bengtsson and E. Weihs, Long-Term Forecasts of Frequency Containment Reserve Prices – TFT Technical Report (Supplementary Report 2). RISE Research Institutes of Sweden AB, 2026.

[2] Svenska kraftnät, Långsiktig Marknadsanalys 2021 (LMA2021). Svenska kraftnät, Sundbyberg, 2021. Available at: https://www.svk.se/om-oss/nyheter/aktuellt/2021/lansering-av-svks-langsiktiga-marknadsanalys-lma-2021/ [A newer LMA2024 was released during the project; its data file is available at: https://www.svk.se/siteassets/om-oss/rapporter/2024/lma_2024.pdf]

[3] Svenska kraftnät, Balancing Market Outlook 2030. Svenska kraftnät, Sundbyberg, 2021. Available at: https://www.svk.se/48d672/siteassets/om-oss/rapporter/2024/svenska-kraftnat-balancing-market-outlook-2030.pdf

[4] J. M. Epstein and R. L. Axtell, Growing Artificial Societies: Social Science from the Bottom Up. MIT Press, Cambridge, MA, 1996.

[5] J. M. Epstein, "Agent-based computational models and generative social science," Complexity, vol. 4, no. 5, pp. 41–60, 1999. doi: 10.1002/(SICI)1099-0526(199905/06)4:5<41::AID-CPLX9>3.0.CO;2-F

[6] R. L. Axtell, "Why Agents? On the Varied Motivations for Agent Computing in the Social Sciences," CSED Working Paper No. 17, Center on Social and Economic Dynamics, The Brookings Institution, Washington, D.C., 2000.

[7] C. Schimeczek, K. Nienhaus, U. J. Frey, E. Sperber, S. Sarfarazi, F. Nitsch, J. Kochems, and A. A. El Ghazi, "AMIRIS: Agent-based Market model for the Investigation of Renewable and Integrated energy Systems," Journal of Open Source Software, vol. 8, no. 84, art. 5041, 2023. doi: 10.21105/joss.05041

[8] K. Poplavskaya, J. Lago, and L. de Vries, "Effect of market design on strategic bidding behavior: Model-based analysis of European electricity balancing markets," Applied Energy, vol. 270, art. 115130, 2020. doi: 10.1016/j.apenergy.2020.115130

[9] E. ter Hoeven, J. H. Kwakkel, V. Hess, T. Pike, B. Wang, and J. Kazil, "Mesa 3: Agent-based modeling with Python in 2025," Journal of Open Source Software, vol. 10, no. 107, art. 7668, 2025. doi: 10.21105/joss.07668

[10] D. Masad and J. Kazil, "Mesa: An Agent-Based Modeling Framework," in Proceedings of the 14th Python in Science Conference (SciPy 2015), pp. 51–58, 2015. doi: 10.25080/Majora-7b98e3ed-009

[11]S. Poledna, M. G. Miess, C. H. Hommes, and K. Rabitsch, "Economic Forecasting with an Agent-Based Model," Macroeconomics: Aggregative Models eJournal, 2020.

[12] S. Tian, Y. Lu, X. Ge, and Y. Zheng, "An agent-based modeling approach combined with deep learning method in simulating household energy consumption," Journal of Building Engineering, 2021.

Appendix

Here, we focus on the training process of the ABM as the training of the TFT model is described in the corresponding technical report.

Data Preparation for ABM:

The dataset underwent standard cleaning to remove inconsistencies and missing values. Feature engineering was limited to scaling, ensuring all input variables were normalized to a uniform range, which is crucial for neural network performance. The data was then split into training and validation sets to facilitate model evaluation.

Training Time:

The model was trained on a standard notebook, with each agent's neural network converging in under 10 minutes, making the approach both efficient and accessible.

Hyperparameter Tuning:

The neural network architecture for each agent was iteratively adjusted to maximize predictive performance. The primary evaluation metric used was the prediction score, which corresponds to the R^2 score (coefficient of determination) provided by Python's MLPRegressor. The R^2 score indicates the proportion of variance in the dependent variable that is predictable from the independent variables, with values closer to 1 indicating better predictive accuracy. After tuning, the final R^2 scores for each agent were as follows:

Agent Prediction Scores (R^2)

Agent	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Score	0.45	0.21	0.36	0.31	0.33	0.20	0.34	0.39	0.50	0.43	0.52

The results highlight variability in agent performance, with some achieving scores above 0.5, indicating strong predictive capability, while others suggest room for further optimization.

Through our international collaboration programmes with academia, industry, and the public sector, we ensure the competitiveness of the Swedish business community on an international level and contribute to a sustainable society. Our 2,800 employees support and promote all manner of innovative processes, and our roughly 100 testbeds and demonstration facilities are instrumental in developing the future-proofing of products, technologies, and services. RISE Research Institutes of Sweden is fully owned by the Swedish state.

I internationell samverkan med akademi, näringsliv och offentlig sektor bidrar vi till ett konkurrenskraftigt näringsliv och ett hållbart samhälle. RISE 2 800 medarbetare driver och stöder alla typer av innovationsprocesser. Vi erbjuder ett 100-tal test- och demonstrationsmiljöer för framtidssäkra produkter, tekniker och tjänster. RISE Research Institutes of Sweden ägs av svenska staten.

RISE Research Institutes of Sweden AB
Box 857, 501 15 BORÅS, SWEDEN
Telephone: +46 10-516 50 00
E-mail: info@ri.se, Internet: www.ri.se

Electric power system
RISE Report 2026:50
ISBN:978-91-90109-79-3