

# Generative modeling of electricity imbalance prices for battery optimization

Victor Mylle

Promotors: prof. dr. ir. Chris Develder  
prof. Bert Claessens

Supervisor: Jonas Van Gompel

**Abstract**—In this study, different models are trained to model the imbalance prices of the Belgian electricity market and optimize battery usage for energy trading. The models are trained on published data from Elia, the Belgian Transmission System Operator (TSO), and evaluated using Mean Absolute Error (MAE), Mean Squared Error (MSE), and Continuous Ranked Probability Score (CRPS). The model types include linear models, non-linear models, recurrent neural networks and diffusion models. Optimizing battery usage is done in a two-step approach. First, the models are used to generate full-day Net Regulation Volume (NRV) samples, which are then used to reconstruct the imbalance prices. These prices are then incorporated into the decision-making to charge and discharge a battery. For each of the imbalance price samples, a charging and discharging threshold is determined by a simple grid search to optimize the maximum profit. For each day, the mean of these thresholds is used to charge and discharge the battery. The results show that the diffusion model outperforms the other models in terms of profit generation. The diffusion model achieves a profit increase of 9.74% over the baseline policy, which uses the previous day's NRV as a prediction. This demonstrates the potential benefits of advanced generative models for enhancing decision-making in energy trading.

**Index Terms**—Generative modeling, imbalance prices, diffusion model, battery optimization

## I. INTRODUCTION

The electricity market is a complex system influenced by various factors, with renewable energy sources adding significant volatility. Renewables, such as wind and solar power, are inherently variable and unpredictable, creating challenges in maintaining a balance between electricity supply and demand. This balance is managed by the Transmission System Operator (TSO), Elia in Belgium, using reserves to address potential imbalances, directly impacting electricity imbalance prices.

Market participants with flexible assets, like industrial batteries, can help stabilize the grid, reducing Elia's reliance on reserves and lowering system costs. These participants aim to maximize profits by buying electricity when prices are low and selling when prices are high.

Forecasting imbalance prices is crucial for market participants to make informed trading decisions. Current industry practices often use simplistic strategies, such as fixed price thresholds, which are suboptimal. This thesis aims to develop generative models to forecast imbalance prices in the Belgian electricity market, optimizing battery usage for profit maximization.

The thesis consists of two parts: modeling the Net Regulation Volume (NRV) for the next day and optimizing a policy using NRV forecasts to maximize profit through strategic battery usage. Various models will be trained and compared based on their profit optimization performance.

## II. BACKGROUND

### A. Electricity Market

The electricity market is a complex ecosystem comprising various stakeholders who work collectively to ensure the balance between electricity generation and consumption, while also aiming to achieve profitability. Key participants in this market include Producers, Consumers, Transmission System Operators (TSOs), Distribution System Operators (DSOs), Balance Responsible Parties (BRPs), and Balancing Service Providers (BSPs). Producers generate electricity using various methods such as coal, nuclear energy, and renewable sources like wind and solar. Consumers, which include households, businesses, and industries, use this electricity. [1]

The grid's stability is maintained by the TSO, responsible for the high-voltage transmission of electricity. In Belgium, this role is performed by Elia. The TSO ensures the grid's balance by activating reserves provided by BSPs, who submit bids for supplying these reserves when needed. This balance is crucial because any discrepancy between electricity generated and consumed can lead to instability, potentially causing blackouts and equipment damage.

A fundamental aspect of maintaining grid stability involves the role of BRPs, which are entities responsible for forecasting and balancing electricity consumption and generation. They submit daily balance schedules to the TSO, and any deviation from these schedules results in imbalance charges, which are calculated based on the System Imbalance (SI) and the Net Regulation Volume (NRV). The NRV represents the volume of reserves activated by the TSO to maintain balance, while the SI is derived from the discrepancy between scheduled and actual power exchanges minus the NRV.

Elia uses three types of reserves to manage imbalances: Frequency Containment Reserves (FCR), Automatic Frequency Restoration (aFRR), and Manual Frequency Restoration (mFRR) [2–4]. The activation of these reserves follows a specific order based on their response times to ensure grid stability. The imbalance price, a key component in managing grid stability, is determined by the highest marginal price of activated reserves for a given quarter-hour. This price calculation involves various factors including the SI, NRV, and specific pricing formulas set by the TSO [5].

### B. Generative Modeling

Given the complexity and variability in factors affecting the imbalance price, forecasting it is challenging. Traditional models struggle due to the dynamic nature of these variables and the changing formulas used by the TSO. An alternative approach involves forecasting the NRV and using this to calculate the imbalance price using the TSO's formulas. This approach requires accurate modeling of the NRV distribution, which can be achieved through

generative modeling techniques. The modeling is done using multiple conditional input features like load, wind, photovoltaic and nominal net position. Deterministic forecasting is often not accurate and does not capture the uncertainty in the NRV, which is crucial for managing risk and optimizing trading strategies.

Generative modeling, a branch of machine learning, aims to generate new data samples that resemble the training data. Techniques such as Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), Normalizing Flows, and Diffusion Models are commonly used in this field [6–9]. These models can learn the underlying distribution of the NRV and generate multiple scenarios, aiding in better decision-making.

Quantile regression is another technique that can be used to estimate the distribution of the NRV without prior knowledge of its form. By predicting multiple quantiles, the model can reconstruct the cumulative distribution function (CDF) of the NRV and generate samples accordingly. This method, introduced by Koenker and Bassett [10], is advantageous in capturing the full range of possible outcomes, especially in the tails of the distribution which are critical for risk management. The models are optimized using the Pinball Loss, a metric that penalizes underestimation and overestimation of quantiles, ensuring a balanced prediction across the distribution. The formula for the Pinball Loss is given by:

$$L_{\tau}(y, \hat{y}) = \begin{cases} (\tau - 1)(y - \hat{y}) & \text{if } y < \hat{y} \\ \tau(y - \hat{y}) & \text{if } y \geq \hat{y} \end{cases} \quad (1)$$

where  $y$  is the true value,  $\hat{y}$  is the predicted value, and  $\tau$  is the quantile level.

In this thesis, both autoregressive and non-autoregressive models are explored. Autoregressive models generate samples sequentially, incorporating dependencies between consecutive time steps, while non-autoregressive models generate all values simultaneously, allowing for faster but potentially less realistic samples. Evaluating the performance of these models involves metrics such as the Mean Absolute Error (MAE), Mean Squared Error (MSE), and the Continuous Ranked Probability Score (CRPS) [11], ensuring a comprehensive assessment of their accuracy and reliability. A more advanced diffusion model is also considered, which has shown promising results in generating realistic samples and capturing complex distributions [9]. Various implementations of diffusion models exist, such as the Denoising Diffusion Probabilistic Model (DDPM) [12].

### C. Battery Optimization Policies

Battery optimization in the electricity market involves strategic decisions on when to charge and discharge batteries to maximize profit and maintain battery health. Organizations aim to buy electricity when prices are low and sell when prices are high, based on predictions of future market conditions. In this thesis, the optimization policy leverages generative models to forecast imbalance prices by reconstructing these prices from Net Regulation Volume (NRV) samples for the next day. The policy aims to maximize profit by charging and discharging, while also considering battery health. The maximum number of charge cycles for a battery is around 400 cycles a year. Exceeding this limit can lead to battery degradation, reducing its capacity and efficiency. To prevent excessive charging and discharging, a penalty parameter is introduced in the optimization policy, ensuring a balance between profit maximization and battery health.

1) *Baseline Policies*: Two baseline policies are established for comparison. The first policy uses fixed thresholds for charging and discharging the battery, determined by a grid search on historical

imbalance price data to maximize profit. A penalty parameter is introduced to minimize excessive charging and discharging, preserving battery health.

The second baseline policy sets thresholds based on the NRV of the previous day, under the assumption that the next day’s NRV will be similar. These thresholds are also optimized using a grid search on reconstructed imbalance prices derived from the previous day’s NRV, including a penalty parameter for battery health.

2) *Policies Based on NRV Generations*: A more advanced policy utilizes multiple NRV predictions for the next day, generated by a trained generative model. Each NRV sample is used to reconstruct imbalance prices, and optimal charging and discharging thresholds are determined for each sample through a grid search incorporating a penalty parameter to reduce battery degradation. The final thresholds for the next day are obtained by averaging these optimal thresholds. This approach showcases the potential of using NRV generations to enhance decision-making in battery optimization.

## III. RESULTS & DISCUSSION

### A. NRV Modeling

Multiple model types are trained to generate new full-day samples of the NRV for a certain day. This is done by training the models on data published by Elia, the TSO in Belgium. The data consists of the NRV, load history and forecast, photovoltaic power history and forecast, wind power history and forecast, and the nominal net position. The models are evaluated using MAE, MSE, and CRPS.

The first set of models includes linear models, non-linear models, and recurrent neural networks like GRU. These models are trained using Quantile Regression to estimate the cumulative distribution of the NRV values. The models output multiple quantile values which can be interpolated to generate a cumulative distribution function. For each quarter of the day, such distribution is predicted and reconstructed. The distributions are then used to sample NRV values and generate full-day NRV predictions.

The quantiles used during this study are 1%, 5%, 10%, 15%, 30%, 40%, 50%, 60%, 70%, 85%, 90%, 95%, and 99%. These quantiles are chosen to get a good approximation of the cumulative distribution of the NRV. More quantiles are used in the tails of the distribution to ensure a good approximation of the extreme values which is important for risk management. An example of a reconstruction of the cumulative NRV distribution for a certain quarter is shown in Figure 1. This distribution can then be used to sample NRV values for that quarter.

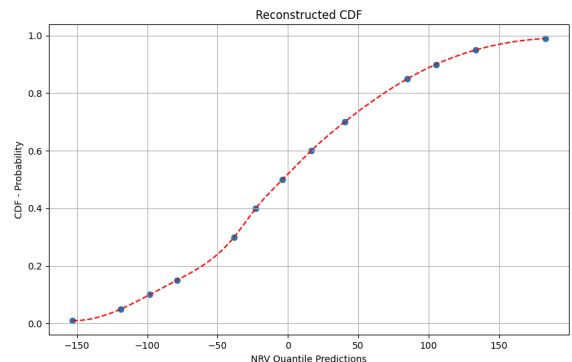


Fig. 1. Example of a reconstructed cumulative distribution function of the NRV for a certain quarter.

The models using Quantile Regression are trained in both autoregressive and non-autoregressive settings. The autoregressive models output the NRV quantile values for each quarter sequentially, while the non-autoregressive models output all quantile values simultaneously. Sampling from the autoregressive models is done by feeding the sampled NRV value back into the model to predict the quantiles of the NRV for the next quarter. This makes sure the dependencies between the quarters are captured in the samples and leads to smoother samples. The non-autoregressive models do not capture these dependencies because they output all quantile values simultaneously for a full day. For each quarter, the cumulative distribution functions are reconstructed and sampled. The sample for a certain quarter does not depend on which value was sampled for the previous quarter. This leads to less smooth samples but allows for faster generation of samples. An example of this behavior is shown in Figure 2. The autoregressive model generates smoother samples compared to the non-autoregressive model.



Fig. 2. Comparison between autoregressive and non-autoregressive models. The samples are generated using a linear model using all input features.

The different models are evaluated using MAE, MSE and CRPS to get an insight into the modeling performance of the models. The results show that non-autoregressive models achieve a better CRPS score than autoregressive models while having worse MAE and MSE scores. An explanation for this behavior is the error propagation in the sampling using autoregressive models. When a sampled value is fed back into the model to predict the quantiles of the next quarter, the error in the sampled value is propagated and the outputted quantiles are shifted. The CRPS evaluates these shifted distributions against the target values which can result in a worse score.

The results also show that using more input features leads to better modeling performance for the autoregressive models. This is not always the case for the non-autoregressive models. Because non-autoregressive models need to output the quantiles for every quarter of

the day simultaneously, they need to capture more complex patterns in the input data. The input data itself also contains more values for non-autoregressive models because the forecasts for the whole day are included while the autoregressive models only use the forecasts for the next quarter. This can lead to the non-autoregressive models not being able to capture the complex patterns in the input data and having worse modeling performance when more input features are used.

A more recent model type, diffusion models, is also explored. These models are trained using the Denoising Diffusion Probabilistic Model (DDPM) which has shown promising results in generating realistic samples and capturing complex distributions. The diffusion models are trained using the same input features as the other models and are evaluated using the same metrics. The architecture used for the diffusion models is a simple feedforward neural network that uses linear hidden layers with a ReLU activation function. More advanced architectures can be explored in future work to improve the modeling performance of the diffusion models. Each layer of the diffusion model is conditioned on the input features which guide the sampling process. This guidance is very simple and can be improved by using more complex conditioning mechanisms.

Samples are generated using the diffusion models by sampling noise from a standard normal distribution and feeding it through the model together with the input features. The model denoises the noise in multiple steps to generate a full-day sample of the NRV. This process is shown in Figure 3. First, the generations are very noisy and do not resemble the target distribution. After multiple denoising steps, the generations become more realistic and resemble the target distribution more closely.




Fig. 3. Intermediate steps of the diffusion model for example 864 from the test set. The confidence intervals shown in the plots are made using 100 samples.

When comparing the different models based on the evaluation metrics MAE, MSE, and CRPS, several trends emerge. Firstly, non-autoregressive models tend to outperform autoregressive models in terms of CRPS, suggesting they are better at capturing the overall distribution of the target variable. However, autoregressive models typically achieve lower MAE and MSE scores, indicating they provide more accurate point predictions.

In terms of complexity, non-linear models generally perform better than their linear counterparts across all metrics, indicating that capturing non-linear relationships in the data is crucial for improving model performance. Interestingly, GRU models, despite their complexity and higher number of parameters, do not always outperform simpler non-linear models, particularly in the autoregressive setting.

The diffusion models, while promising in their ability to generate realistic samples, show worse performance in terms of MAE, MSE, and CRPS. This discrepancy is likely due to the narrow confidence intervals produced by these models, which do not capture the variability present in the data.

It is difficult to draw definitive conclusions about the best model type, as their performance varies across different metrics and visual inspection of the generated samples. However, diffusion models show potential for capturing complex distributions, while non-linear models offer a good balance between accuracy and complexity.

### B. Policy Evaluation

The goal of this study is to use the imbalance price generations to optimize battery usage for profit maximization. The simple policy determines a charging and discharging threshold for each imbalance price sample of a certain day by performing a simple grid search to maximize the profit. The mean of these thresholds is then used to charge and discharge the battery for that day. The policy is evaluated on a test set that starts on 01-01-2023 and ends on 12-12-2023. Days with missing data are excluded from the evaluation for a fair comparison.

One of the key questions is if the metrics used to evaluate the models also correlate with the profit made by the policies. This is not the case and the models that perform best in terms of the evaluation metrics do not necessarily generate the most profit. Because of this, the profit needs to be evaluated during the training process of the models. If another metric is used to evaluate the models and do early stopping, the models might be overfitting too much. A validation set is used consisting of the last two months of 2022 to evaluate the profit of the policy during the training of the models. Evaluating the whole test set during the training process is not feasible because of the computational cost. The validation set gives a good indication of the performance of the profit during the training process. This metric is used to do early stopping and prevent overfitting of the models.

The penalty parameter was tuned for every model to prevent excessive charging and discharging of the battery. A total of 283 charge cycles can be used for the battery during the test set. This also makes comparing the profit of the different models and policies fair because more charge cycles would lead to more profit. For the evaluation of the policy using the different models, a battery of 2MWh with a charge/discharge power of 1MW is used.

The results show that the diffusion model outperforms the other models in terms of profit generation. The diffusion model achieves a profit increase of 9.74% over the baseline policy, which uses the previous day's NRV as a prediction. Only the diffusion model was able to outperform the baseline policy. The results are shown in Figure 4.

## IV. CONCLUSION

This thesis explored the use of generative models to model the Net Regulation Volume (NRV) and optimize battery charging and discharging policies in the electricity market. Various models were trained and assessed using metrics such as Mean Absolute Error (MAE), Mean Squared Error (MSE), and Continuous Ranked Probability Score (CRPS). The primary goal was to model imbalance prices and use these predictions to enhance battery optimization policies.

The results indicated that traditional evaluation metrics do not always align with the profitability of the policies. Consequently, models were evaluated based on the profit they generated, revealing that better modeling performance does not necessarily lead to higher profits.

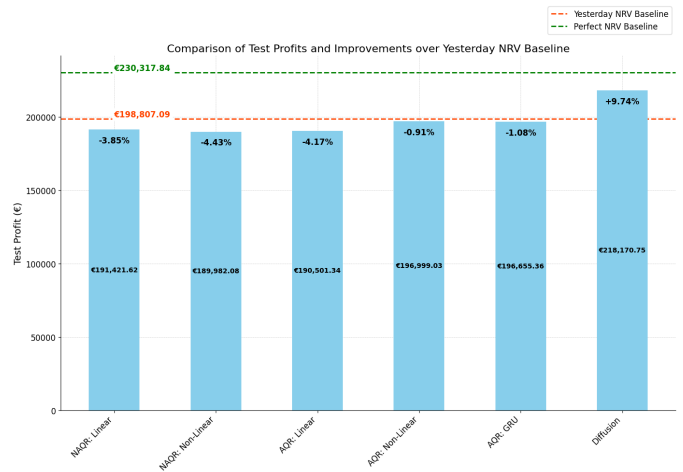


Fig. 4. Comparison of the profit made by the different policies and baselines using the test set.

Among the tested models, only the diffusion model outperformed the baseline policy, achieving a 9.74% increase in profit.

The findings highlight the potential of generative modeling in modeling imbalance prices and optimizing energy trading strategies. Future work could involve more sophisticated diffusion model implementations and advanced conditioning techniques to further improve battery utilization and profitability. This thesis demonstrates that focusing on profitability as a measure of success can lead to more practical applications in the energy market.

## REFERENCES

- [1] De geliberaliseerde elektriciteitsmarkt omvat vele partijen die allen samen moeten werken en tegelijkertijd proberen winst te maken. hieronder volgt een... [Online]. Available: <https://www.next-kraftwerke.be/nl/weten/spelers-energiemarkt>
- [2] FCR. [Online]. Available: <https://www.elia.be/en/electricity-market-and-system/system-services/keeping-the-balance/fcr>
- [3] aFRR. [Online]. Available: <https://www.elia.be/en/electricity-market-and-system/system-services/keeping-the-balance/afrr>
- [4] mFRR. [Online]. Available: <https://www.elia.be/en/electricity-market-and-system/system-services/keeping-the-balance/mfrr>
- [5] Elia, "Tariffs for maintaining and restoring the residual balance of individual access responsible parties." [Online]. Available: <https://www.elia.be/-/media/project/elia/elia-site/customers/tariffs-and-invoicing/tariffs-and-invoicing/en/grille-tarifaire-desequilibre-2022-en-v20220214s.pdf>
- [6] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks." [Online]. Available: <http://arxiv.org/abs/1406.2661>
- [7] D. P. Kingma and M. Welling, "Auto-encoding variational bayes." [Online]. Available: <http://arxiv.org/abs/1312.6114>
- [8] D. Rezende and S. Mohamed, "Variational inference with normalizing flows," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. PMLR, pp. 1530–1538. [Online]. Available: <https://proceedings.mlr.press/v37/rezende15.html>

- [9] J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics." [Online]. Available: <http://arxiv.org/abs/1503.03585>
- [10] R. Koenker and G. Bassett, "Regression quantiles," vol. 46, no. 1, pp. 33–50, publisher: [Wiley, Econometric Society]. [Online]. Available: <https://www.jstor.org/stable/1913643>
- [11] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," vol. 102, no. 477, pp. 359–378. [Online]. Available: <http://www.tandfonline.com/doi/abs/10.1198/016214506000001437>
- [12] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models." [Online]. Available: <http://arxiv.org/abs/2006.11239>