Multi-Day Day-Ahead Price Scenario Generation using Conditioned Multivariate Elliptical Copulas

Rick Verhage*†, Gijs Verhoeven*, Bart van der Holst*, Koen Kok*
*Electrical Energy Systems (ESS) Group, Department of Electrical Engineering
Eindhoven University of Technology, Eindhoven, the Netherlands
{g.verhoeven.1, b.v.d.holst, j.k.kok}@tue.nl
†Lamb Weston EMEA, Kruiningen, the Netherlands
rickverhage@lambweston.eu

***Abstract*—This paper develops a multi-day probabilistic day-ahead electricity price forecasting model based on the Conditioned Multivariate Elliptical Copula. By integrating exogenous variables and inter-day dependencies, the model generates robust price scenarios for industrial demand-side flexibility (DSF) applications. Coupled with a stochastic optimization framework for flexible battery operations, stability tests identify optimal parameters and a useful forecast horizon. The results demonstrate the potential application of the model for DSF to optimize grid usage and reduce electricity costs.**

***Index Terms*—Demand-side Flexibility, Day-ahead market, Multi-day Forecasting, Scenario generation, Stochastic optimization**

I. Introduction

The European energy transition, driven by electrification and renewable energy integration, is straining grid capacity. In the Netherlands, this has led to congestion in the grid, limiting industrial sites from expanding their existing grid connections and delaying sustainable growth [1]. As network reinforcements may take years, industries seeking to electrify their processes face significant challenges. Optimizing existing grid connections has therefore become essential. Demand-side flexibility (DSF) offers a solution by enabling consumers to optimize electricity usage within existing grid constraints while adapting to market and network conditions. With increasing short-term electricity trading and price volatility, integrating DSF into daily operations creates new opportunities, for example in the Day-Ahead (DA) market.

However, most existing studies limit DA price forecasts to single days, whereas many industrial processes operate on multi-day cycles. Extending DSF scheduling over multiple days could improve cost efficiency by leveraging price fluctuations over longer time horizons. Evaluating DSF's feasibility in the DA market requires scenario-based modeling to simulate operational decision-making. Electricity market mechanics play a crucial role in DSF's cost-effectiveness, particularly regarding price uncertainty. Literature regarding price dynamics of electricity markets has been explored widely. Existing literature on DA market price forecasting, extensively described in [2], includes both deterministic and probabilistic models. Deterministic DA price forecasting models provide single-point estimates but fail to capture price fluctuations [3], [4]. Probabilistic forecasts, while gaining attention, remain less explored. A relevant study by Van Wijngaarden [5] demonstrated strong results using Conditioned Multivariate Elliptical Copulas (CMEC) for probabilistic DA forecasting, balancing effectiveness with accessibility. However, existing models remain limited to daily horizons, with limited research on multi-day forecasting. A notable exception is [6], which introduces a hybrid multi-day forecast model incorporating price spike analysis.

This paper introduces a multi-day probabilistic forecasting model integrated into a stochastic optimization framework to assess DSF's potential for industrial applications.

The main contributions of this work are:

- Development and validation of a multi-day probabilistic forecast model to generate DA price scenarios, built on and extending the CMEC framework adapted from [5].
- Development of a stochastic optimization model for DSF load shifting techniques applied to a flexible battery system in a continuous industry context.

The remainder of this paper is structured as follows: First, Section II gives a detailed explanation of the proposed methodology. Then, the results are presented in Section III. Finally, the paper concludes with key findings and insights in Section IV.

II. Methodology

This section outlines the methodology for probabilistic DA price forecasting using the CMEC. First, the CMEC-based single-day DA forecasting model is introduced. The approach is then extended to a multi-day framework by incorporating inter-day dependencies. Finally, the stability and performance of the forecasts are evaluated through a real-world use-case involving a stochastic optimization approach.

II.1 Single-Day DA Price Forecasting with CMEC

The study of [5] presents a robust probabilistic DA price forecasting method, overcoming the limitations of traditional models like linear regression. These models assume normally distributed errors, which fail to capture the extreme fluctuations and non-linear patterns typical of electricity prices. The

Conditioned Multivariate Elliptical Copula overcomes these limitations by offering a more flexible and accurate framework for modeling the dependencies between exogenous variables and electricity prices. Unlike some traditional approaches, the CMEC does not assume normality in data and is able to effectively capture the asymmetries and irregularities.

The CMEC, based on the work of Duque et al. [7], is based on Sklar's theorem [8] and models a multivariate joint distribution of the 24 DA prices $\Lambda_1, \ldots, \Lambda_d$ by combining the marginal distributions $F_i(\Lambda_i)$ for $i = 1, \ldots, d$ of random variables, and exogenous variables Z with a copula function $C : [0,1]^T \rightarrow [0,1]$ that captures the dependencies between them such that the (cumulative) probability distribution function F can expressed as [8]:

$$F(\Lambda_1, \ldots, \Lambda_d | Z_1, \ldots) = C(F_1(\Lambda_1), \ldots, F_d(\Lambda_d) | Z_1, \ldots) \tag{1}$$

This method extends unconditioned copulas to include external factors such as weather forecasts, load predictions, or past price trends. By conditioning on these exogenous variables, the model refines the sample space, generating forecasts that align with specific external conditions.

Sampling in this context involves generating subsets of data that reflect the relationship between the random variables and the exogenous variables. Instead of sampling from the entire distribution, the model utilizes the conditional copula to focus only on regions that are consistent with the provided exogenous inputs. This process effectively narrows the sample space, ensuring that the generated samples, which contain resulting random variables, are relevant to the given exogenous inputs. After generating samples, scenario reduction is required to create a manageable set of representative scenarios. This process involves selecting a subset of the generated samples and assigning probabilities to each, ensuring that the reduced set maintains the characteristics of the original distribution [9]. By doing so, the computational complexity of actually using the scenario forecasts in optimization schemes is reduced while preserving the accuracy.

To generate single-day DA price scenarios, the CMEC is constructed using 24 hourly DA prices per day as random variables over a specified training period. By incorporating exogenous variables, the CMEC gains enhanced adaptability, allowing it to capture diverse market conditions and adjust its application to different cases. This adaptability is particularly valuable for this research to generate multi-day DA price forecasts. The performance of the CMEC in generating probabilistic DA price forecasts has been evaluated in [5]. This study demonstrated improved accuracy when exogenous variables included the total load forecasts for the DA market, solar irradiation forecasts—particularly at hours 6, 12, and 18— and the pre-mean, defined as the average of the previous day's 24 hourly DA prices. From now on, this model will be referred to as the 'base CMEC' which will serve as the basis for developing probabilistic multi-day DA price forecasts.

II.2 Multi-day Day-Ahead price forecasting

The base CMEC is designed to capture the statistical dependencies between exogenous variables and hourly DA prices as random variables within a single 24-hour period. Extending the model's dimensionality to incorporate multiple days of DA prices as random variables risks introducing artificial dependencies between price hours that do not reflect actual market dynamics. Additionally, expanding the model with more variables would significantly increase the CMEC's dimensionality. This could lead to overfitting and make it more difficult to obtain reliable results. Consequently, restricting training and sampling to a 24-hour basis is both practical and methodologically sound, as it focuses on the strongest relationships within daily DA prices and exogenous variables.

II.2.1 Incorporating Inter-Day Dependencies

To enable multi-day forecasts, the sample space must be enriched with an exogenous variable that captures the price dependency between consecutive days. Ideally, this exogenous variable is simultaneously a byproduct of the sampling process. This would allow the sampling result of each day to be used for the next day, creating a link between days within the sampling period T_s. In this case, the logical focus is therefore on capturing the relationship between hourly prices of hour 23 on day $D-1$ ($\lambda_{23\text{(D-1)}}$) and hour 0 on day D ($\lambda_{0\text{(D)}}$). Pearson's Correlation (PC, ρ_{23-0}) is used as a metric that measures the strength and direction of a linear relationship between variables. While PC does not capture non-linear dependencies, it is able to provide meaningful insights for smaller datasets. In this case, a ρ_{23-0} value close to one or minus one indicates a significant linear dependency between $\lambda_{23\text{(D-1)}}$ and $\lambda_{0\text{(D)}}$, whereas a value near zero suggests little linear relationship between these two variables. In the historical data, a strong linear relationship between $\lambda_{23\text{(D-1)}}$ and $\lambda_{0\text{(D)}}$ has been identified across a majority of eight selected periods, as indicated in Table I.

Table I: Pearson's Correlation ρ_{23-0} over each period

	Winter 1	**Winter 2**	**Spring 1**	**Spring 2**
ρ_{23-0}	0.895	0.999	0.987	-0.113
	Summer 1	**Summer 2**	**Fall 1**	**Fall 2**
ρ_{23-0}	0.973	0.972	0.921	0.792

One period, 'Spring 2', deviates from the usual strong trend observed across other periods. Given the volatility of electricity prices, such outliers are expected and do not undermine the proposed method. Due to the overall strong positive correlation, $\lambda_{23\text{(D-1)}}$, is added to the base CMEC as an exogenous variable. This expands the sample space to capture dependencies across days, enabling a multi-day forecast. To contribute effectively to the model's performance, all exogenous variables must conform to the characteristic rhombus shape of the t-distributed CMEC. Since the previous day's average price met this criterion in prior research [5], it is assumed that $\lambda_{23(\text{D-1})}$ will meet this criterion as well, enabling its inclusion.

II.2.2 Training and Scenario Reduction

The training set of the CMEC consists of data from the 365 days prior to the starting date of a sampling period T_{s}, with a separately fitted CMEC for week- and weekend days. All assessments within this paper are conducted out-of-sample, as in-sample stability has been proven for the base CMEC. During sampling, real forecasted load and solar irradiation values ($FL_{6,12,18}$- $FS_{6,12,18}$) are provided for each day. $\lambda_{23(\text{D-1})}$ is known only for the first day and estimated for subsequent days by averaging across sampled scenarios, discarding the 10th and 90th percentiles to reduce outlier influence. This computed value is then inserted as exogenous variable into the sampling for the following day, for all days in sampling period T_{s}.

The constructed CMEC enables independent sampling of multiple days in a period, with daily samples randomly concatenated to form a multi-day forecast. This is feasible since the samples M are drawn without inherent probabilities. Within this study, these samples are independently sampled for four days in each period. To obtain a manageble DA scenario set N, scenario reduction is required. The Fast-Forward Selection (FFS) algorithm is used to select representative scenarios while preserving probabilistic structure [9]. Equal probability is initially assumed, but sample sets with a high correlation between $\lambda_{23(\text{D-1})}$ and $\lambda_{0(\text{D})}$ are prioritized by assigning a weighting factor ρ_{23-0} in the FFS algorithm. This ensures that the resulting scenarios maintain strong inter-day dependencies while preserving the probabilistic structure of the original distribution.

II.3 Forecast Stability and Performance

Naturally, the probabilistic forecast stability and overall performance must be assessed. Additionally, to be able to use the forecasts in stochastic optimization, the length of the useful forecast horizon must be determined. It is necessary to evaluate the required number of samples M and the stability of the scenario reduction. However, due to the stochastic nature of probabilistic forecasting and the fact that multiple scenarios are generated for a given date range, it is difficult to directly evaluate the model's stability against historical, real data using traditional metrics like Root Mean Square Error or Mean Absolute Percentage Error [10]. For this reason, the stability of the model is evaluated through an application.

II.3.1 Use-Case and Optimization Problem

To assess the stability of the probabilistic forecasting model, a real-world application is used. The stability of the model is evaluated through its performance in a stochastic optimization framework for demand-side flexibility in a continuous industrial setting. The use-case in this paper features a flexible battery system enabling load shifting techniques, leveraging real load data from an industrial facility where the battery acts as a flexible electricity storage unit. By strategically storing available electricity and discharging it to support the industrial process load, the system effectively shifts energy consumption to reduce costs. This helps to minimize electricity costs by optimizing power imports from the grid. For the DA scheduling of the battery, the total electricity costs, based on the power imported from the grid $P_{\text{t}}^{\text{grid,s}}$ for each forecast period T_{f} are minimized over all DA scenarios $\omega_n \in \Omega$ with a corresponding probability π_{ω_n}:

$$C_{\text{total}}^{\text{DA,bat,e}} = \min\Big(\sum_{\omega_n \in \Omega} \pi_{\omega_n} \cdot \sum_{t \in T_{\text{f}}} [P_{\text{t}}^{\text{grid,s}} \cdot \lambda_{\text{t},\omega_n}]\Big) \tag{2}$$

subject to

$$P_{\text{t}}^{\text{ch,s}} \leq z_{\text{t}}^{\text{bat,s}} \cdot P^{\text{ch}_{\text{c}}} \tag{3}$$

$$P_{\text{t}}^{\text{dch,s}} \leq (1 - z_{\text{t}}^{\text{bat,s}}) \cdot P^{\text{dch}_{\text{c}}} \tag{4}$$

$$0 \leq \text{SOC}_{\text{t}}^{\text{s}} \leq \text{SOC}^{\text{max}} \tag{5}$$

$$\text{SOC}_{\text{t}}^{\text{s}} = \text{SOC}_{\text{t-1}}^{\text{s}} + \frac{\delta_{\text{t}}}{3600} \cdot \Big(P_{\text{t}}^{\text{dch,s}} \cdot \eta^{\text{bat,dch}} - \frac{P_{\text{t}}^{\text{ch,s}}}{\eta^{\text{bat,ch}}}\Big) \tag{6}$$

$$P_{\text{t}}^{\text{grid,s}} = P_{\text{t}}^{\text{TEPL}} - P_{\text{t}}^{\text{dch,s}} + P_{\text{t}}^{\text{ch,s}} \tag{7}$$

$$P_{\text{t}}^{\text{grid,s}} \leq P^{\text{grid}_{\text{c}}} \tag{8}$$

Constraints (3)-(6) ensure proper battery operations by defining limits and updating the battery's energy content. Specifically, constraints (5)-(6) set the energy content of the battery during scheduling in relation to the real capacity of the battery SOC^{max}. Furthermore, (7) defines the overall power balance of the scheduled grid demand $P_{\text{t}}^{\text{grid,s}}$, the total electric processing load $P_{\text{t}}^{\text{TEPL}}$ and the (dis)charging of the battery. Finally, (8) ensures that the scheduled grid demand $P_{\text{t}}^{\text{grid,s}}$ never exceeds the grid capacity $P^{\text{grid}_{\text{c}}}$. To illustrate the industrial load demand profile, Fig. 1 depicts the Total Electricity Processing Load (TEPL) for the facility. The grid capacity $P^{\text{grid}_{\text{c}}}$, set at 9.53 MW_e, is represented by the dashed line. Each sampling period is marked with a distinct color, which is maintained consistently throughout the analysis and will therefore not be repeated in other legends.

II.3.2 Model Stability and Performance Evaluation

To assess the stability and performance of DA scenarios, two key metrics are used. The first is the absolute expected-realized cost difference ($C_{|\Delta_{\text{exp-real}}|}^{\text{bat,e}}$), which measures the absolute difference between the expected cost from the stochastic optimization and the realized cost when applying the optimized DA schedule to real prices:

$$C_{|\Delta_{\text{exp-real}}|}^{\text{bat,e}} = \left| C_{\text{stoch. exp}}^{\text{bat,e}} - C_{\text{stoch. real}}^{\text{bat,e}} \right| \tag{9}$$

This quantifies how well DA scenarios predict real prices. The second metric, the Expected Value of Perfect Information (EVPI), evaluates the theoretical benefit of perfect foresight, highlighting the value of a stochastic approach. A high EVPI indicates significant variation in DA price scenarios, making stochastic optimization preferable over deterministic methods. It is calculated as:

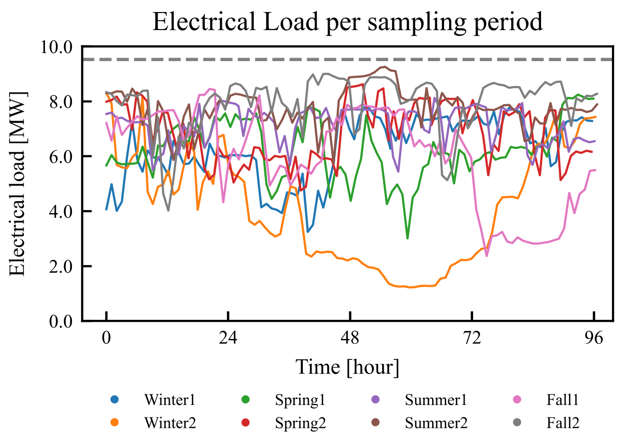

Fig. 1: Electrical load of the use-case over eight periods.

$$EVPI = \sum_{\omega \in \Omega} \pi_\omega \cdot C^{\text{bat,e}}_{\text{det. exp},\omega} - C^{\text{bat,e}}_{\text{stoch. exp}} \tag{10}$$

By analyzing these metrics across different sample sizes and scenario reductions, the optimal number of samples M and scenarios N is determined for stable performance. These values are then used as input for assessment of the optimal forecast horizon.

Since accurate DA scheduling depends on identifying the maximum useful forecast horizon, it is important to account for forecast performance degradation over time due to accumulating errors in exogenous inputs. The Continuous Ranked Probability Score (CRPS) measures probabilistic forecast accuracy by comparing the forecasted and observed Cumulative Distribution Functions (CDFs). The CRPS is computed for each DA scenario and aggregated across all scenarios, weighted by their respective probabilities. The equation for CRPS is adapted to include weights from [11] as follows:

$$CRPS_{\text{t}}(CDF_{\text{t}}, \lambda_{\text{t},\omega_n}) = \sum_{\omega_n \in \Omega} \pi_{\omega_n} \cdot$$
$$\int_{-\infty}^{\infty} \left[CDF_{\text{t}}(\lambda_{\text{t,real}}) - H_{\text{t}}(\lambda_{\text{t,real}} - \lambda_{\text{t},\omega_n}) \right]^2 \, d\lambda_{\text{t,real}} \tag{11}$$
$$\forall t \in T_{\text{s}}$$

where λ_{real} and λ_w represent historical and forecasted DA prices, respectively, and H_{t} is a Heaviside function. The forecast horizon limit is identified where CRPS shows a steep increase, indicating reduced reliability. The optimal forecast horizon is further refined using the absolute expected-realized cost difference metric from (9) by varying the forecast horizon from 6 hours to the maximum found via CRPS. The horizon is selected at the point where the forecast accuracy begins to degrade, signaling a decline in the forecast accuracy.

III. Results

This section presents the results of the stability tests and the impact of scenario reduction on the model's performance. The analysis evaluates the influence of sample size M and scenario count N on the model's stability and accuracy. After the appropriate values for M and N are selected, the optimal forecasting horizon is determined using DA forecasting and cost deviation analysis within the flexible battery use case.

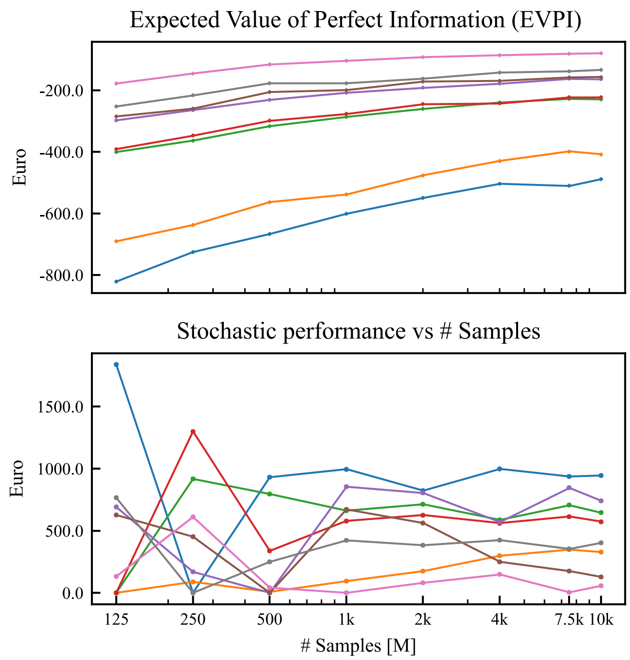

Fig. 2: DA Scenario stability test metrics across a varying number of samples.

III.1 Stability Analysis of Sampling and Scenario Reduction

For the stability tests, the sample stability and the performance of scenario reduction were analyzed independently. To isolate the effect of scenario size on sample stability, metrics were calculated for each sample size M and the weighted average was taken across all tested number of scenarios N. Conversely, for scenario reduction, stability metrics were computed for each N averaging over all tested M. Fig. 2 and 3 show the stability tests for each selected period. The EVPI is plotted on a common y-axis scale in euros, comparing the stochastic and deterministic solutions per period. However, $C^{\text{bat,e}}_{|\Delta_{\text{exp-real}}|}$ is normalized within each period, facilitating a comparison between the 8 periods. Both metrics are plotted against a logarithmically scaled x-axis for clarity.

Fig. 2 represents the sample stability resuts, showing that EVPI increases with a larger M. The negative values indicate that the stochastic solution consistently overestimates total costs within a period compared to the deterministic solution. However, as M increases, the overestimation decreases, suggesting that the variance within the scenarios is reduced. Stability is observed around 4000 samples, with winter periods acting as outliers. Similarly, stochastic performance declines with increasing M, showing reduced DA price forecast errors. Although 7500 samples yield marginal improvement, 4000 is computationally preferred. In Fig. 3, EVPI shows an opposite trend, as increasing N increases the variance between scenarios. However, beyond $N = 35$, additional scenarios offer little benefit. The same applies to stochastic performance,

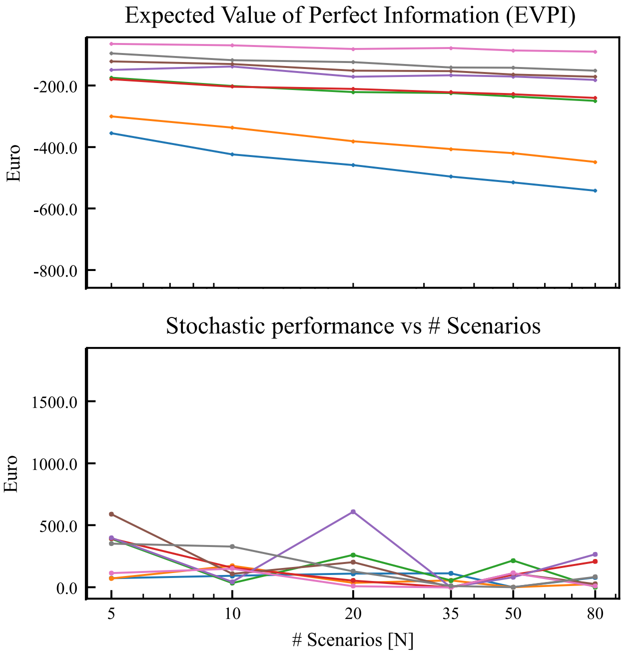

Fig. 3: DA Scenario stability test metrics across a varying number of scenarios.

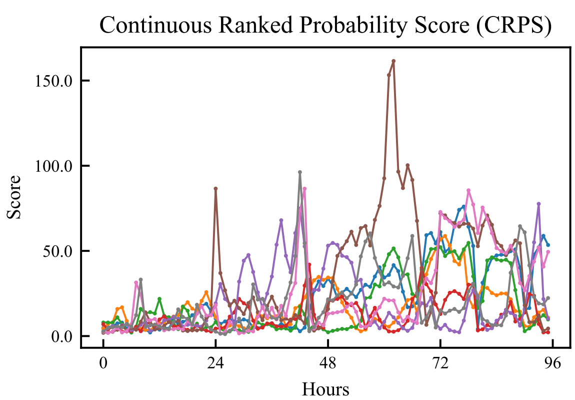

Fig. 4: Continuous Ranked Probability Score (CRPS) over time across eight sampling periods.

which weakens in trend but still declines with increasing N. Stability is reached at $N = 35$. For larger values noise may be reintroduced. Based on these findings, $M = 4000$ and $N = 35$ are used for further analysis.

III.2 Forecast Horizon Analysis

The DA forecasts were generated using the selected number of samples and scenarios for each period. Fig. 4 shows the computed CRPS, for the full sampling period T_s (96 hours). As expected, the first 24 hours exhibit low errors due to the stable behavior of the base CMEC.

Beyond the first 24 hours, the forecasts accuracy declines, with instability becoming evident after 48 hours across all periods. Therefore, the maximum forecast horizon T_f is set at 48 hours, beyond which forecasts become unreliable. However, to evaluate practical performance in the use-case, T_f is varied from 6 to 48 hours (in 3-hour steps), and $C^{\text{bat,e}}_{|\Delta_{\text{exp-real}}|}$ is computed.

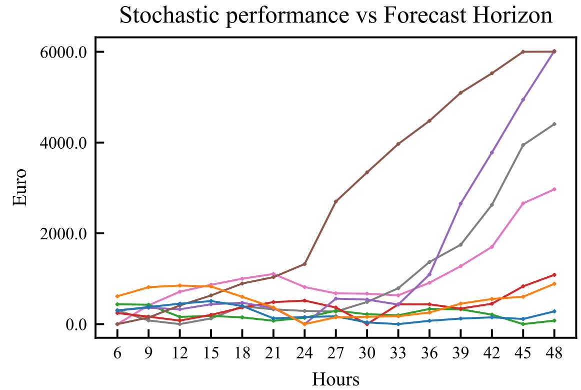

Fig. 5: Relative absolute delta between expected and real costs over varying DA forecast horizons within the simplified battery use case.

Fig. 5 shows the relative absolute delta between expected and real costs over varying DA forecast horizons. Excluding the outlier presented by the brown line, corresponding with the summer weekend period, forecasts remain stable up to 33–36 hours. Beyond this, most periods show steep increases in $C^{\text{bat,e}}_{|\Delta_{\text{exp-real}}|}$, indicating declining forecast utility. Therefore, the useful forecast horizon is concluded to be 33 hours.

IV. Conclusion

This study extends the CMEC framework to generate robust multi-day probabilistic day-ahead price scenarios integrated into a stochastic DSF scheduling model for flexible batteries. Stability analyses yield optimal parameters (M = 4000, N = 35) and identify a 33-hour forecast horizon, ensuring accurate cost reduction and grid optimization. These findings highlight the potential of DSF and CMEC forecast scenarios to enhance energy system performance, cost reduction, and support sustainable growth. Future work could focus on leveraging CMEC for inter-market arbitrage and imbalance market scenarios, further improving economic efficiency and evaluating DSF feasibility within different electricity markets.

References

[1] TenneT, “Afname congestiemanagement-onderzoek provincie Zeeland Analyse naar beschikbare transportcapaciteit voor afname van,” 2024. [Online]. Available: https://www.tennet.eu/nl/de-elektriciteitsmarkt/congestiemanagement/onderzoeken-congestiemanagement

[2] J. Lago, G. Marcjasz, B. De Schutter, and R. Weron, “Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark,” *Applied Energy*, vol. 293, p. 116983, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261921004529

[3] J. Nowotarski and R. Weron, “Recent advances in electricity price forecasting: A review of probabilistic forecasting,” *Renewable and Sustainable Energy Reviews*, vol. 81, pp. 1548–1568, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1364032117308808

[4] F. Ziel and R. Steinert, “Probabilistic mid- and long-term electricity price forecasting,” *Renewable and Sustainable Energy Reviews*, vol. 94, pp. 251–266, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1364032118303885

[5] E. Van Wijngaarden, B. Van der Holst, and N. G. Paterakis, “Day-ahead price scenario generation using conditioned multivariate elliptical copulas,” in *2024 International Conference on Smart Energy Systems and Technologies (SEST)*, 2024, pp. 1–6.

[6] D. Manfre Jaimes, M. Zamudio López, H. Zareipour, and M. Quashie, “A Hybrid Model for Multi-Day-Ahead Electricity Price Forecasting considering Price Spikes,” *Forecasting*, vol. 5, no. 3, pp. 499–521, 9 2023.

[7] E. M. S. Duque, P. P. Vergara, P. H. Nguyen, A. Van Der Molen, and J. G. Slootweg, “Conditional Multivariate Elliptical Copulas to Model Residential Load Profiles from Smart Meter Data,” *IEEE Transactions on Smart Grid*, vol. 12, no. 5, pp. 4280–4294, 9 2021.

[8] M. Sklar, “Fonctions de répartition à N dimensions et leurs marges,” *Annales de l’ISUP*, vol. VIII, no. 3, pp. 229–231, 1959.

[9] H. Heitsch and W. Römisch, “Scenario reduction algorithms in stochastic programming,” *Computational Optimization and Applications*, vol. 24, pp. 187–206, 2003. [Online]. Available: https://doi.org/10.1023/A:1021805924152

[10] M. Kaut and S. W. Wallace, “Evaluation of scenario-generation methods for stochastic programming,” *Pacific Journal of Optimization*, vol. 3, no. 2, pp. 257–271, 2007.

[11] H. H. Chen, A. Mohammad, T. Totsuka, A. Laliashvili, and M. Navarro, “Evaluation of Probabilistic Electricity Demand Forecasts,” in *PMAPS 2024 - 18th International Conference on Probabilistic Methods Applied to Power Systems*. Institute of Electrical and Electronics Engineers Inc., 2024.