Cross-Border Information in Electricity Price Forecasting: Benefits of a Panel Data Approach

Tobias Schnabel
Department of Statistics
University of Oxford
Oxford, United Kingdom
ctobiasschnabel@gmail.com

***Abstract*—European short-term electricity prices are determined in day-ahead auctions matching hourly power supply and demand. Forecasting these prices is typically seen as a country-specific task, but increasing intermittent renewable generation and market integration introduce spillover effects. We therefore evaluate if treating forecasting as a panel data problem improves model accuracy and stability by leveraging cross-country dependencies alongside longitudinal country data. Using day-ahead prices from 13 European countries in the Core Capacity Calculation Region from 2019 to 2024, we show that panel methods can incorporate cross-sectional information from interconnected markets. An open-source forecasting evaluation engine systematically compares panel and traditional univariate approaches under various market conditions. Our results indicate that panel methods produce more stable predictions, especially during high price volatility and longer forecasting horizons. This work highlights the prediction stability benefits and conceptual simplicity of panel approaches for market monitoring, cross-border trading, and regulatory assessments of market coupling effects.**

***Index Terms*—Multivariate Regression, Open Source Software, Power markets, Prediction methods, Time series analysis**

I. Introduction

Electricity markets across Europe are increasingly interconnected—not only through market integration but also via physical transmission lines that facilitate cross-border energy flows. With the rising share of intermittent renewable energy sources, electricity prices have become more volatile and subject to pronounced spillover effects between bidding zones. While nearly all published forecasting work treats each country's market in isolation, the reality is that electricity flows freely across borders, creating inherent interdependencies that can be harnessed for improved prediction accuracy [1].

By adopting a panel data framework, our approach explicitly accounts for these transmission-induced spillovers. Rather than modeling each market separately, we integrate cross-sectional information from multiple European countries. This method provides a more robust representation of the system dynamics, capturing both the temporal evolution within individual markets and the simultaneous interactions via the physical grid. In doing so, it reflects the true nature of the European electricity system where prices are not solely determined by local conditions but are significantly influenced by cross-border flows.

Our submission presents a benchmarking study that compares traditional country-specific forecasting techniques with the proposed panel-based approach. The results demonstrate that by leveraging the well-established advantages of panel data analysis, forecasts become not only more stable but also more competitive in terms of accuracy. This practical framework offers valuable insights for market participants—ranging from policymakers and grid operators to traders—providing a clearer picture of the underlying market dynamics and better anticipating price movements under different market conditions.

The remainder of the submission is organized as follows. First, we describe the panel data framework and the details of data preprocessing needed to integrate cross-border electricity price information. Next, we outline our forecasting methodology and present a performance comparison with traditional univariate approaches. Finally, we discuss the implications of our findings for market monitoring and regulatory assessments and suggest directions for future research.

II. Data and Problem Description

We focus our analysis on the Core Capacity Calculation Region (CCR). This region—comprising the bidding zones of 13 EU Member States (Austria, Belgium, Croatia, Czech Republic, France, Germany, Hungary, Luxembourg, Netherlands, Poland, Romania, Slovakia, and Slovenia) and managed by 16 TSOs—represents the heart of European power market integration, with coordinated capacity calculation and flow-based market coupling that optimize cross-border exchanges, and provides a robust setting for capturing the known physical spillovers inherent in electricity transmission.

The dataset used in this application comprises hourly day-ahead electricity prices for the 12 European bidding zones (with Luxembourg and Germany forming a joint bidding zone) listed above from January 1, 2019, 00:00 to June 30, 2024, 23:00, resulting in 578,028 observations.

This dataset has both a cross-sectional (between bidding zones) and a time dimension, which makes it a panel. The objective is the prediction on day D-1 of the hourly day-ahead price for all bidding zones (day D). Additionally, we use the following exogenous predictors for each bidding zone:

- Coal futures settlement prices for API2 European (thermal) coal futures of the evening of day D-2 (forward-filled for each hour of day D-1). Measured in US Dollars per 1,000 metric tons.

- Gas futures settlement prices for TTF natural gas futures of the evening of day D-2 (forward-filled for each hour of day D-1). Measured in Euros per MWh.
- Calendar information: hour of the day (0-23), day of the week (0-6, where 0 is Monday), day of the year (1-366), month (1-12), quarter (1-4), week of the year (1-53), and a ”weekend” dummy are all included as numerical features, following standard forecasting practice.
- Day-ahead generation forecasts of the total power generation per bidding zone released on day D-1. Measured in MW, hourly resolution.
- Generation forecasts for wind and solar power per bidding zone released on day D-1. Measured in MW, hourly resolution.
- Day-ahead load forecasts of the total power demand per bidding zone released on day D-1. Measured in MW, hourly resolution.

All data is publicly available and sourced from the ENTSO-E transparency platform and Yahoo Finance. Preprocessing was performed with minimal intervention. Missing exogenous values were imputed by cross-sectional weighting using corresponding values from neighboring countries. No variable had more than 0.2% of missing entries. No standardization or other transformation was applied given that the market conditions spanned by the dataset vary dramatically, which makes for an interesting robustness exercise for each of the forecasting methods used.

III. Forecasting Methodology

At its core, our forecasting challenge can be framed as a problem of data imputation. In the panel of electricity prices—where rows represent bidding zones and columns represent hourly observations—future prices are effectively ”missing.” By leveraging the information contained in the observed portions of this matrix, our goal is to reliably impute these missing future values. This perspective allows us to jointly exploit both the temporal dynamics within each market and the cross-sectional interdependencies across markets. In the sections that follow, we detail the various methodologies we employ to accomplish this imputation task, ranging from matrix completion models to classic linear regression-based approaches.

A. Model Selection and Explainability

Our methodological approach prioritizes model interpretability alongside forecasting accuracy. While deep learning models have shown promising results in electricity price forecasting, as demonstrated by [2] and [3], their black-box nature obscures the cause-effect relationships driving their predictions. As [4] argue, this lack of transparency limits their practical applicability in critical infrastructure. We therefore focus on linear models, which offer clear interpretability of how different variables influence price forecasts.

We implement five main approaches: Matrix Completion with Nuclear Norm Minimization (MC-NNM), Elastic Net (EN) regression, LASSO regression, and two variants of LEAR (Lasso Estimated AutoRegressive) [5], [6], [7]. MC-NNM, developed by [8], leverages the panel structure of electricity markets to capture both temporal and cross-sectional dependencies. We extend their approach with a temporal autocorrelation structure to better handle forecasting applications (MC-NNM-TSR) [9]. The technical details of MC-NNM are thoroughly documented in [8].

EN combines LASSO and Ridge regression penalties to handle correlated predictors while maintaining some sparsity. LASSO performs variable selection by shrinking some coefficients exactly to zero, making it particularly suitable for high-dimensional settings with many potential predictors. LEAR, specifically designed for electricity price forecasting, estimates separate LASSO models for each hour of the day, capturing intraday patterns. While LEAR has shown strong performance in single-market applications, it treats each bidding zone independently.

To address this limitation, we formulate Panel LEAR, which extends the LEAR framework to explicitly model cross-border dependencies. Rather than estimating separate models for each bidding zone, Panel LEAR pools data across zones while maintaining hour-specific coefficients. This reduces the total number of models from 288 (24 hours × 12 zones) to 24, while allowing the capture of cross-sectional relationships. The approach represents a middle ground between fully independent models like LEAR and fully integrated approaches like MC-NNM.

B. Open-Source Forecast Engine

To ensure reproducibility and standardized comparison across methods, we developed a custom forecasting engine in Python. The engine standardizes data preparation, model fitting, prediction, and performance evaluation across all estimators, offering insights into the relative strengths of different approaches in this complex panel data setting. This standardization is particularly important given the high dimensionality of our data, which includes 12 price series and multiple exogenous variables per country.

The engine performs consistent data preprocessing across all models, creating rolling windows of configurable sizes (56, 84, and 112 days, similar to [6]) and handling missing values through cross-sectional weighting. For model training, it employs automated hyperparameter optimization using Optuna [10], with performance-based recalibration triggers ensuring models remain responsive to changing market conditions. Models are recalibrated daily using the previous 90 days of data, with optimization triggered when performance degrades by more than 10% compared to the previous three days. Our implementation emphasizes practical usability alongside theoretical rigor. The engine executes all operations in parallel where possible, significantly reducing computation time compared to sequential processing. This efficiency is particularly important for real-world applications where forecasts must

TABLE I
COMPARISON OF ESTIMATORS WITH 56-DAY SLIDING WINDOW

Estimator	RMSE				MAE				rMAE				Mean Execution Time (s)		
	min	mean	median	max	min	mean	median	max	min	mean	median	max	`fit()`	`predict()`	`optimize()`
MC-NNM	8.411	49.331	34.793	320.343	5.752	38.460	23.114	301.834	0.278	1.887	1.534	10.365	0	5.796	2.529
MC-NNM-TSR	8.032	49.140	34.614	319.978	5.328	38.273	23.110	301.834	0.278	1.876	1.560	9.807	0	10.442	5.734
ElasticNet	3.764	25.972	17.371	711.283	2.785	17.768	11.687	273.206	0.138	0.930	0.771	31.205	0.871	0.000	1.092
LASSO	15.714	81.371	69.208	446.215	12.245	57.267	44.296	300.592	0.375	3.507	2.797	19.674	5.079	0.000	1.667
LEAR	6.186	393.940	47.033	103 176.328	3.776	80.316	23.454	13 831.132	0.264	5.076	1.255	1151.798	1.202	0.025	0.065
LEAR-Panel	6.434	503.392	46.835	180 758.706	3.791	96.716	23.568	25 112.787	0.271	6.120	1.252	2031.926	1.185	0.025	0.064

be generated within tight time constraints. All components are fully documented and available on GitHub, facilitating transparency and reproducibility of our results.

The evaluation framework implements multiple error metrics including Mean Absolute Error (MAE), and relative Mean Absolute Error (rMAE), following the recommendations of [6]. These metrics are computed consistently across all models and time periods, enabling robust comparison of forecasting performance. The modular design of the engine allows for easy integration of new models and evaluation metrics.

IV. Results and Discussion

Our empirical analysis provides compelling evidence that treating electricity price forecasting as a panel data problem yields tangible improvements in forecast stability and overall performance. In particular, the panel-based methods—most notably the matrix completion approaches (MC–NNM and its temporal extension, MC–NNM-TSR)—consistently deliver forecasts that are less volatile and more reliable over varying time horizons. Compared to traditional regression models such as Elastic Net and LASSO, these integrated approaches not only maintain competitive accuracy but also reduce the incidence of extreme prediction errors, especially during periods of heightened market stress. In all tables in this section, the rows of each table correspond to estimators, while the columns present summary statistics for each metric as well as execution times. The best (lowest) value for each summary statistic (for example mean RMSE) is highlighted in green, the second-best (second lowest) value per column in cyan. All values are for the entire test period spanning 2006 days. The visualizations in this section display rMAE due to its intuitive interpretation as a relative performance metric [6].

A. Results Using a 56-day Sliding Window

Table I and Fig. 1 display the performance of the estimators using a 56-day sliding window. The EN estimator demonstrates the best overall performance across all metrics, achieving the lowest minimum, mean, and median RMSE, MAE, and rMAE. This suggests that for this short window size, the EN's ability to balance between the ℓ_1 and ℓ_2 penalties provides an effective approach to capturing the underlying patterns in the data. MC-NNM-TSR and MC-NNM show very similar performance, ranking second and third in terms of mean and median RMSE and MAE.

The LASSO estimator performs notably worse than the EN and MC-NNM approaches. This may indicate that the strict sparsity induced by LASSO is too restrictive for capturing

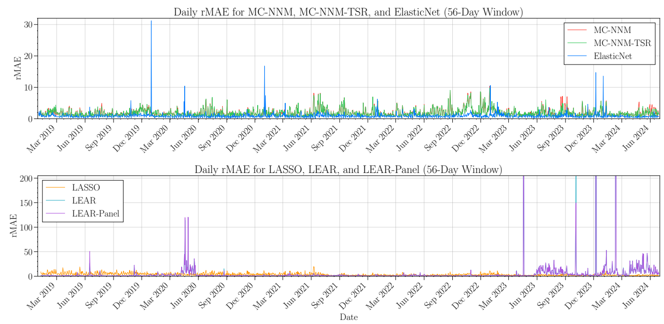

Fig. 1. Daily rMAE of Estimators With 56-Day Sliding Window

the complex dependencies in electricity prices over this time horizon. Fig. 1 is split into two panels due to the significant difference in rMAE scales between the three top-performing estimators and the three LASSO-based methods. The y-axes have different scales, which is also reflected in the maximum rMAE values reported in Table I. Both LEAR and LEAR Panel exhibit extremely high maximum errors, showing that they are prone to occasional severe mispredictions. However, their median performance is competitive, particularly in terms of rMAE. Comparing LEAR and LEAR Panel, we can observe that their performances are quite similar, with LEAR Panel showing slightly higher mean and maximum rMAE. This is somewhat surprising, as one might expect the panel version to perform better by leveraging cross-country information.

B. Results Using an 84-day Sliding Window

Table II and Fig. 2 display the performance comparison of the estimators using an 84-day sliding window. Many of the trends observed in the 56-day window persist, but there are some notable differences. The gap between EN and the MC-NNM variants narrows slightly, particularly in terms of rMAE. MC-NNM-TSR now performs second-best across RMSE, MAE, and rMAE, indicating that the incorporation of temporal structure becomes more beneficial with longer windows.

The stability advantage of both MC-NNM variants becomes more pronounced with the 84-day window. While EN maintains the lowest mean errors, it shows increased volatility in its rMAE, while both MC-NNM estimators maintain more consistent performance. The LASSO estimator's performance improves relative to the 56-day window but still lags behind EN and MC-NNM approaches.

TABLE II
COMPARISON OF ESTIMATORS WITH 84-DAY SLIDING WINDOW

Estimator	RMSE				MAE				rMAE				Mean Execution Time (s)		
	min	mean	median	max	min	mean	median	max	min	mean	median	max	`fit()`	`predict()`	`optimize()`
MC-NNM	8.064	51.841	36.789	383.996	4.897	40.630	24.910	369.120	0.286	2.001	1.631	13.186	0	16.858	12.856
MC-NNM-TSR	7.908	51.861	36.486	375.507	4.950	40.662	24.777	360.373	0.296	2.004	1.655	13.117	0	27.125	19.785
ElasticNet	3.523	26.208	18.110	661.197	2.406	18.173	11.983	370.660	0.127	0.956	0.788	28.481	1.111	0.000	1.409
LASSO	17.368	84.523	74.656	361.459	12.490	59.708	46.076	320.579	0.368	3.689	2.927	20.924	7.361	0.000	2.521
LEAR	6.164	297.727	45.463	71 530.336	3.929	64.211	24.031	9458.683	0.266	4.035	1.201	793.786	1.847	0.027	0.218
LEAR-Panel	6.049	294.556	45.389	62 382.795	3.855	61.266	23.942	5960.928	0.269	3.685	1.203	468.175	1.819	0.025	0.216

TABLE III
COMPARISON OF ESTIMATORS WITH 112-DAY SLIDING WINDOW

Estimator	RMSE				MAE				rMAE				Mean Execution Time (s)		
	min	mean	median	max	min	mean	median	max	min	mean	median	max	`fit()`	`predict()`	`optimize()`
MC-NNM	7.825	53.550	37.992	401.004	5.678	42.105	25.980	386.470	0.300	2.100	1.732	12.992	0	50.471	31.906
MC-NNM-TSR	7.277	53.391	37.727	401.004	5.131	41.940	25.874	386.470	0.295	2.094	1.729	13.639	0	51.182	33.835
ElasticNet	3.566	26.458	19.024	368.116	2.252	18.496	12.468	294.846	0.117	0.978	0.802	14.195	1.360	0.000	1.742
LASSO	19.991	86.776	79.405	412.320	16.242	61.541	48.840	365.495	0.356	3.780	3.012	22.281	9.251	0.000	3.083
LEAR	6.809	174.872	42.637	40 532.072	4.300	48.550	23.487	5246.787	0.245	2.823	1.179	159.449	2.420	0.027	0.276
LEAR-Panel	6.870	186.739	42.721	33 253.808	4.308	48.753	23.480	3293.416	0.255	2.855	1.186	191.374	2.439	0.027	0.278

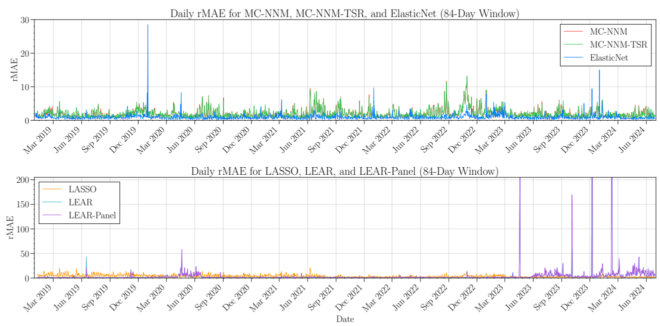

Fig. 2. Daily rMAE of Estimators With 84-Day Sliding Window

C. Results Using a 112-day Sliding Window

With a 112-day window (Table III, Fig. 3), the relative performance of the models shifts further. MC-NNM-TSR achieves comparable accuracy to EN while maintaining its stability advantage. This suggests that longer windows allow the model to better capture both cross-sectional dependencies and temporal patterns in the data.

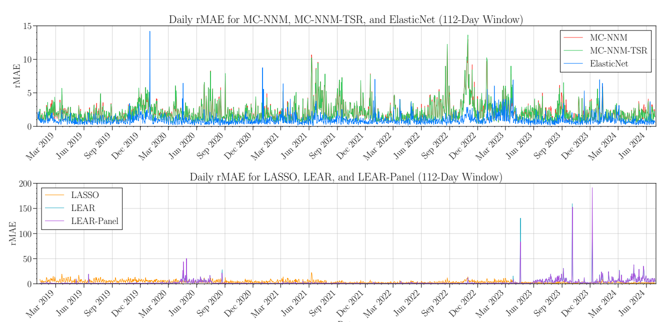

Fig. 3. Daily rMAE of Estimators With 112-Day Sliding Window

Notably, both LEAR variants show improved performance with the longer window, but their maximum errors remain significantly higher than those of MC-NNM and EN. This shows that while longitudinal modeling becomes more effective with more historical data, it still cannot match the stability of panel methods on the same window size during extreme events.

V. CONCLUSION

This paper demonstrates the benefits of treating electricity price forecasting as a panel data problem. Our results show that Elastic Net regression performs best on most metrics across all three window sizes. We demonstrate that panel methods, particularly MC-NNM variants, offer superior stability compared to traditional univariate approaches while maintaining competitive accuracy. This stability advantage becomes more pronounced with longer forecasting windows and during periods of market stress.

A key insight from our analysis is that panel methods achieve better results with simpler models. While traditional approaches require separate models for each bidding zone—288 models in the case of LEAR—panel methods like MC-NNM effectively capture market dynamics with a single integrated model. This reduction in model complexity is not merely a computational convenience; it reflects the fundamental reality of Europe's increasingly integrated electricity markets. As the CORE CCR implements flow-based market coupling and coordinates capacity calculation across its 13 member states, prices in different bidding zones become more tightly linked. Panel methods naturally mirror this market structure by modeling cross-border dependencies directly rather than treating them as external shocks to individual markets.

This alignment between model structure and market reality becomes particularly relevant as European power markets move toward closer integration. The CORE CCR's role as the central hub of European market coupling makes it an ideal testing ground for panel approaches. As renewable penetration increases and cross-border flows become more volatile, the ability to capture simultaneous interactions across multiple bidding zones will only grow in importance. Our results suggest that panel methods are well-positioned to handle these challenges while maintaining model interpretability and computational efficiency.

The open-source implementation and standardized evaluation framework provide practitioners with tools to leverage these insights in real-world applications. Future research could explore extending the temporal structure in MC-NNM-TSR to capture more complex dependencies and incorporating additional cross-border information such as transmission capacity data.

Our work demonstrates that leveraging cross-border information significantly enhances electricity price forecasting in European markets. As physical power flows freely across national boundaries, forecasting approaches should similarly transcend siloed market analysis. The panel data methods we evaluated capture these interconnections, offering superior stability during market volatility—a crucial advantage for risk management and operational planning. This stability becomes increasingly valuable with longer historical windows, suggesting these approaches better reflect the evolving integration of European electricity markets. The practical implications are substantial: grid operators can anticipate price movements across multiple markets simultaneously, traders can reduce extreme prediction errors during market stress, and regulators can better assess market coupling effectiveness. As European energy markets continue to integrate under initiatives like the Clean Energy Package, methodologies that explicitly account for these interdependencies will become essential tools for all market participants navigating the energy transition.

Acknowledgment

I gratefully acknowledge the thoughtful supervision of Prof. Frank Windmeijer during this dissertation work.

References

[1] J. Lago, F. De Ridder, P. Vrancx, and B. De Schutter, "Forecasting day-ahead electricity prices in Europe: The importance of considering market integration," Appl. Energy, vol. 211, pp. 890–903, Feb. 2018. [Online]. Available: https://doi.org/10.1016/j.apenergy.2017.11.098

[2] P. Joshi, S. Størdal, G. Lien, D. Mishra, and E. Haugom, "Dynamic Electricity Price Forecasting via Deep Learning," in 2023 IEEE International Conference on Energy Technologies for Future Grids (ETFG), Dec. 2023, pp. 1–6. [Online]. Available: https://doi.org/10.1109/ETFG55873.2023.10407554

[3] K. G. Olivares, C. Challu, G. Marcjasz, R. Weron, and A. Dubrawski, "Neural basis expansion analysis with exogenous variables: Forecasting electricity prices with NBEATSx," Int. J. Forecast., vol. 39, no. 2, pp. 884–900, Apr. 2023. [Online]. Available: https://doi.org/10.1016/j.ijforecast.2022.03.001

[4] M. M. Mascarenhas, M. Amelin, and H. Kazmi, "Bridging Accuracy and Explainability in Electricity Price Forecasting," in 2024 20th International Conference on the European Energy Market (EEM), Jun. 2024, pp. 1–6. [Online]. Available: https://doi.org/10.1109/EEM60825.2024.10608857

[5] T. Matsumoto and F. Ziel, "Electricity Price Forecasting with Principal Component-Guided Sparse Regression," in 2024 20th International Conference on the European Energy Market (EEM), Jun. 2024, pp. 1–6. [Online]. Available: https://doi.org/10.1109/EEM60825.2024.10609019

[6] 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," Appl. Energy, vol. 293, p. 116983, Jul. 2021. [Online]. Available: https://doi.org/10.1016/j.apenergy.2021.116983

[7] S. B. Amor, P. A. Devi, and F. Müsgens, "Meta-Forecasting for Solar Power Generation: Algorithm-Based Swarm Intelligence," in 2024 20th International Conference on the European Energy Market (EEM), Jun. 2024, pp. 1–5. [Online]. Available: https://doi.org/10.1109/EEM60825.2024.10608959

[8] S. Athey, M. Bayati, N. Doudchenko, G. Imbens, and K. Khosravi, "Matrix Completion Methods for Causal Panel Data Models," J. Am. Stat. Assoc., vol. 116, no. 536, pp. 1716–1730, Oct. 2021. [Online]. Available: https://doi.org/10.1080/01621459.2021.1891924

[9] T. Schnabel, "lightweight-mcnnm: A Python package for Matrix Completion with Nuclear Norm Minimization," 2024. [Online]. Available: https://pypi.org/project/lightweight-mcnnm/

[10] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A Next-generation Hyperparameter Optimization Framework," in Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2019, pp. 2623–2631. [Online]. Available: https://doi.org/10.1145/3292500.3330701