



OPEN ACCESS

Operations Research and Decisions

www.ord.pwr.edu.pl

OPERATIONS
RESEARCH
AND DECISIONS
QUARTERLY



Regularization for electricity price forecasting

Bartosz Uniejewski^{1*} 

¹Department of Operations Research and Business Intelligence, Wrocław University of Science and Technology, Wrocław, Poland

*Email address: bartosz.uniejewski@pwr.edu.pl

Abstract

The most commonly used form of regularization typically involves defining the penalty function as a ℓ^1 or ℓ^2 norm. However, numerous alternative approaches remain untested in practical applications. In this study, we apply ten different penalty functions to predict electricity prices and evaluate their performance under two different model structures and in two distinct electricity markets. The study reveals that LQ and elastic net consistently produce more accurate forecasts compared to other regularization types. In particular, they were the only types of penalty functions that consistently produced more accurate forecasts than the most commonly used LASSO. Furthermore, the results suggest that cross-validation outperforms Bayesian information criteria for parameter optimization, and performs as well as models with ex-post parameter selection.

Keywords: *electricity price forecasting, regularization, power market, convex regularization, LQ regularization, elastic net*

1. Introduction

Electricity price forecasting (EPF) is a critical task for many participants in the power market, including power generators, traders and consumers. Accurate predictions of electricity prices enable one to make informed decisions regarding production, purchasing, and sales [23, 25]. In particular, forecasting of day-ahead market prices is of great importance for decision-making in electricity markets [16]. This auction-based market plays a crucial role in ensuring the reliability and efficiency of the power system [17]. The day-ahead market not only allows adjusting the long-term position to actual exposure [12, 13], but is also a reference point for over-the-counter (OTC) trading and settlement.

One of the main challenges in developing reliable electricity price forecasting models is the high volatility of the electricity market [33]. Many traditional forecasting models fail to capture these complex dynamics, resulting in inaccurate forecasts [14]. To overcome this problem, researchers have turned to regularization techniques to improve the accuracy of their models. Regularization is a method used in machine learning and statistical modeling to prevent the overfitting of a model by adding a penalty term to the model's objective function.

Although regularization has been widely applied in machine learning and related fields, its use in electricity price forecasting is still limited. To our knowledge, up to now, only three out of ten considered in this paper penalty functions have been tested in the EPF context. The first research that incorporated regularization in this context was conducted by Barnes and Balda [3], who used ridge regression to assess the financial viability of battery storage in electricity markets. More recently, Ludwig et al. [15] and Ziel et al. [37] applied the least absolute shrinkage and selection operator (LASSO) to improve prediction accuracy. Lastly, Uniejewski et al. [31] utilized elastic net to automate variable selection in the electricity market. In this paper, we investigate the performance of ten different types of regularization to forecast electricity prices. The techniques we consider are adaptive LASSO, clipped LASSO, concave potential function, elastic net, forward-LASSO adaptive shrinkage, LASSO, LQ, minimax concave PLUS, ridge regression and smoothly clipped absolute deviation.

We assess the performance of those regularization types using two datasets from the German EPEX SPOT and Iberian OMIE markets and two model structures well-established in the literature. We employ also cross-validation (CV) or Bayesian information criteria (BIC) to select the value of the tuning parameter (λ) and compare their performance. Finally, we propose a fully automated approach to select all the regularization parameters based on CV. The accuracy of price forecasts obtained with regularized models is compared with the benchmark models estimated with ordinary least square (OLS).

The paper is structured as follows. In Section 2, we provide a brief overview of the datasets. Next, in Section 3 we outline the regularization techniques considered and summarize related research on the EPF. In Section 4, we compare the accuracy of the forecasts obtained with the proposed methods, and finally, in Section 5, we conclude the results.

2. Datasets

This empirical study utilized datasets from two distinct European markets, each covering a period of seven years. The first market is the German EPEX SPOT. It has recently experienced a significant increase in the share of renewable generation and underwent major structural changes, making it an attractive target for analysis. The data were obtained from the transparency platform (<https://transparency.entsoe.eu>) and consist of five hourly time series (see top panel in Figure 1):

- day-ahead electricity prices for the DE-AT-LU bid zone until 30.9.2018 and DE-LU afterwards
- day-ahead total load forecasts for Germany
- day-ahead solar generation forecasts for Germany
- day-ahead wind (on-shore) generation forecasts for Germany
- day-ahead wind (off-shore) generation forecasts for Germany

The second market examined in the study is the Spanish OMIE market. Unlike the German market, the Iberian one does not allow prices to drop below zero. This feature, along with the rapid development of the market, makes it an interesting subject for analysis. Data were also collected from the transparency platform and comprise four hourly time series (bottom panel in Figure 1).

- day-ahead marginal electricity prices for Spain
- day-ahead total load forecasts for Spain

- day-ahead solar generation forecasts for Spain
- day-ahead wind (on-shore) generation forecasts for Spain




Figure 1. Day-ahead prices and day-ahead load forecast time series for German EPEX (top panel) and Spanish OMIE (bottom panel) from 1.1.2015 to 31.12.2021.

The vertical dashed line marks the beginning of the out-of-sample test period (29.12.2016)

It should be noted that until 30.06.2021, the Spanish electricity market had a price cap between 0 and 180 €/MWh. In the second half of 2021, this limit was changed to a wider price range between -500 and 3 000 €/MWh. However, as the change does not affect the price series, but rather was caused by a change in price dynamics, it has not been included in the model.

Both time series span 2557 days, ranging from 1.1.2015 to 31.12.2021. Missing or doubled values (corresponding to the time change) were replaced by the average of the closest observations and the arithmetic mean of the values from doubled hours, respectively.

The structure of the price series in both European markets underwent notable changes during the seven-year period under study, as shown in Figure 1. The onset of 2020, coupled with the emergence of the COVID-19 pandemic, followed by the Russian invasion of Ukraine in 2022, induced substantial shifts in electricity price dynamics in 2020–2021.

3. Methodology

The volatility and irregularity of electricity prices pose a challenge in producing accurate forecasts. Outliers in the data, caused by sudden spikes, can distort model coefficients and lead to higher in-sample errors for non-spiky periods. To address this issue, Uniejewski et al. [32] proposed various functions of variance stabilizing transformations (VST) to reduce the variation in the data and improve forecast accuracy. In this paper, the normal probability integral transform (N-PIT) is applied to each (price and exogenous variables) time series:

$$x_{d,h} = N^{-1}(\hat{F}_{X_{d,h}}(X_{d,h}))$$

where $\hat{F}_{X_{d,h}}$ is an empirical cumulative distribution function of sample $X_{d,h}$ and N^{-1} is the inverse function of the standard normal distribution.

In the next step, the transformed time series are used as the inputs of the forecasting models which are estimated via OLS or one of ten different regularization types. After generating forecasts of the transformed price ($\hat{p}_{d,h}$), the inverse transformation is applied to obtain the final forecasts of the electricity price $\hat{P}_{d,h} = \hat{F}_{P_{d,h}}^{-1}(N(\hat{p}_{d,h}))$.

3.1. Models

The price forecasting task is implemented separately for each hour. For each day of the out-of-sample period, this results in 24 different parameter sets. This so-called multivariate framework is a common approach in the EPF literature (see [38] for a discussion of the differences between the uni- and multivariate frameworks). The 24 individual models are independent of each other, but their estimation is based on the same set of information.

The day-ahead forecasts of the hourly electricity price for both markets are determined within a rolling window scheme using a 728-day calibration window. First, all considered models are calibrated to data from the initial calibration period (1.1.2015–29.12.2016), and forecasts for all 24 hours of the next day (30.12.2016) are determined. Then, the window is rolled forward by one day; the models are re-estimated, and forecasts for all 24 hours of 31.12.2016 are computed. This procedure is repeated until the predictions for the 24 hours of the last day in the sample (31.12.2021) are made.

The optimal choice of calibration window length is the subject of lively discussion [11, 29], in this paper the 728-day calibration window is used. The choice was based on limited empirical tests comparing the prediction accuracy of models calibrated to the considered window and three shorter calibration sizes.

3.1.1. Parsimonious ARX model

To assess the impact of different types of regularization on prediction accuracy, two different model structures are used in this study. The undoubted advantage of using regularization is the almost unlimited number of initially considered explanatory variables in the model. As a result, expert knowledge becomes less important. For this reason, studies incorporating regularization usually focus on parameter-rich model structures, and the concept is very rarely considered for estimating models with only a few inputs.

To address this literature gap, as the first underlying model, we propose a parsimonious autoregressive structure, which is a well-established model in the EPF literature [38]. The transformed price on day d

and hour h (i.e., $p_{d,h}$) is modeled by the following equation:

$$p_{d,h} = \beta_1 p_{d-1,h} + \beta_2 p_{d-2,h} + \beta_3 p_{d-7,h} + \beta_4 p_{d-1}^{\max} + \beta_5 p_{d-1}^{\min} + \beta_6 p_{d-1,24} + \\ + \beta_7 l_{d,h} + \beta_8 s_{d,h} + \beta_9 w_{d,h}^{\text{on}} + \beta_{10} w_{d,h}^{\text{off}} + \sum_{i=1}^7 \beta_{10+i} D_d^i + \varepsilon_{d,h}$$

where $p_{d-1,h}$, $p_{d-2,h}$ and $p_{d-7,h}$ account for the autoregressive terms and correspond to prices from the same hour of the previous day, two days before, and a week before, respectively. The midnight price of electricity for the previous day, which is the last known price at the time of prediction, is represented by $p_{d-1,24}$. p_{d-1}^{\max} and p_{d-1}^{\min} are the maximum and minimum prices of the previous day. The transformed day-ahead load forecast for a given hour of a day is represented by $l_{d,h}$, while transformed day-ahead forecasts of solar, on-shore wind, and off-shore wind generation are represented by $s_{d,h}$, $w_{d,h}^{\text{on}}$, and $w_{d,h}^{\text{off}}$, respectively. Finally, D_d^1, \dots, D_d^7 are weekday dummies and $\varepsilon_{d,h}$ is the noise term. Note that off-shore wind generation is excluded from the model for the OMIE market, as there is no installed capacity in Spain. The model consisting of 17 (16 for OMIE) regressors is referred to as ARX throughout the paper.

3.1.2. Parameter rich model

The second model structure, denoted by the fARX model, is a more complex autoregressive model that incorporates 277 (229 for OMIE) regressors. The first 72 regressors (p_{d-i,h^*}) refer to the autoregressive terms and consist of information on prices from all hours from one, two, and seven days before forecasted day d . The next six variables account for the non-linear terms and are reflected by the minimum (p_{d-i}^{\min}) and maximum (p_{d-i}^{\max}) of all prices from one, two, and three days before day d . Another 192 regressors (144 for OMIE) refer to exogenous variables. The model uses information on prediction of load ($l_{d-i,h}$), wind ($w_{d-i,h}^{\text{on}}$ and $w_{d-i,h}^{\text{off}}$) and solar generation ($s_{d-i,h}$) for all hours of the day d and the previous day. Lastly, similarly to the ARX model, to capture the weekly seasonality, the model consists of seven dummy variables D_d^1, \dots, D_d^7 . The formula for the fARX model is as follows:

$$p_{d,h} = \sum_{i \in \{1,2,7\}} \sum_{h^*=1}^{24} \beta_{i,h^*}^p p_{d-i,h^*} + \sum_{i=1}^3 (\beta_i p_{d-i}^{\max} + \beta_{i+3} p_{d-i}^{\min}) \\ + \sum_{i=0}^1 \sum_{h^*=1}^{24} (\beta_{i,h^*}^L l_{d-i,h^*} + \beta_{i,h^*}^S s_{d-i,h^*} + \beta_{i,h^*}^{\text{on}} w_{d-i,h^*}^{\text{on}} + \beta_{i,h^*}^{\text{off}} w_{d-i,h^*}^{\text{off}}) \\ + \sum_{i=1}^7 \beta_{6+i} D_d^i + \varepsilon_{d,h}$$

The ARX and fARX models are considered baseline models that are estimated with regularization procedures discussed in the following Section 3.2.

3.2. Regularization

The concept of regularization involves minimizing the objective function of the original model and adding a penalty to the model parameters. In this study, the objective function is the residual sum of squares (RSS), and the regularization is defined by the following equation:

$$\boldsymbol{\beta} = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} \underbrace{\sum_{d,h} \left(p_{d,h} - \sum_{i=1}^N \beta_i x_{d,h}^i \right)^2}_{\text{RSS}} + \underbrace{\sum_{i=1}^N g(\boldsymbol{\beta}; \lambda, \cdot)}_{\text{penalty function}} \quad (1)$$

where $p_{d,h}$ represents the dependent variable (it is the price series) and $x_{d,h}$ is a matrix of independent variables consisting of model inputs. Finally, the $g(\boldsymbol{\beta})$ is the penalty function, also known as the regularization term, applied to the parameter vector $\boldsymbol{\beta} = \{\beta_1, \beta_2, \dots, \beta_N\}$.

The most commonly used form of regularization involves defining the penalty function as an order norm ℓ^1 or ℓ^2 , scaled by a tuning parameter λ . However, numerous alternative approaches exist, which have yet to be tested in real-world applications. To our best knowledge, only three out of the ten considered regularization types have been previously tested in the EPF context and only LASSO is commonly used in the literature. In this study, ten penalty functions are used to generate the electricity price forecasts, and their performance is compared against the unregularized benchmarks.

Regularization can be utilized to identify the most important variables in the model by selecting an appropriate penalty function. This procedure involves fitting the full model with all predictors using an algorithm that reduces the coefficients of less significant explanatory variables toward zero. One of the benefits of using automatic variable selection methods is their ability to handle a large number of explanatory variables, thus reducing the dependence on expert knowledge, which is often unverified [31].

To conduct the research, we used the Matlab toolbox of McIlhagga [19]. The toolbox utilizes Fisher scoring over an active set with orthant projection [26, 28]. It offers an automated selection of the tuning parameter λ using information criteria or cross-validation. In this study we decided to compare Bayesian information criteria and cross-validation with 7 folds. Moreover, for penalty functions with an extra regularization parameter, the toolbox allows one to utilize CV to select the optimal value of both parameters simultaneously, making the model fully automated (this feature is unavailable for BIC).

3.2.1. Adaptive LASSO

The adaptive least absolute shrinkage and selection operator (aLASSO) penalty function was first introduced by Zou [39]. The adaptive weights in the model are used for penalizing different coefficients in the ℓ^q penalty. According to the author, the adaptive LASSO performs as well as if the true underlying model was given in advance. Although originally the adaptive LASSO model was defined only for the ℓ^1 penalty, it was further generalized to allow for the ℓ^q norm penalty. The adaptive penalty function is defined as follows:

$$g(\boldsymbol{\beta}; \lambda, q) = \lambda \frac{|\beta_i|^q}{\beta_i^*} \quad (2)$$

where β_i^* is a vector of weights estimated with an unregularized model, λ is the tuning parameter and q is the additional parameter that refers to the ℓ^q norm. In this paper β_i^* is defined as a vector of weights estimated with OLS, λ is selected automatically via BIC or CV and q takes one of the values from the grid $Q = \{1, 1.5, 2\}$.

3.2.2. Clipped LASSO

The clipped least absolute shrinkage and selection operator (cLASSO) [2] is another penalty function that aims to generalize LASSO. The idea of clipped LASSO is to set a threshold for the maximum value of the penalty for estimated weights. It is defined as follows:

$$g(\boldsymbol{\beta}; \lambda, \alpha) = \lambda \min \{|\beta_i|, \alpha\} \quad (3)$$

where λ is the tuning parameter and α is responsible for the threshold of the penalty. Based on limited pretests we consider α from a grid $A = \{0.5, 1, 1.5\}$, whereas λ is selected automatically via BIC or cross-validation.

3.2.3. Concave potential function

The regularization term of concave potential function (CPF) [24] is the sum of the values obtained by applying the potential function (PF) to each regressor. The penalty function is defined as follows:

$$g(\boldsymbol{\beta}; \lambda, k) = \lambda \frac{k |\beta_i|}{k + |\beta_i|} \quad (4)$$

where λ is the tuning parameter and k is an additional parameter that refers to the shape of the penalty. Note that if $k \rightarrow \infty$ the penalty becomes equivalent to a standard LASSO penalty. In this paper, the additional parameter k is selected from the initial grid $K = \{5, 15, 25\}$.

3.2.4. Elastic net

The elastic net (EN) [40] can be viewed as an extension of ridge regression and LASSO. This penalty function incorporates a combination of linear and quadratic penalty terms:

$$g(\boldsymbol{\beta}; \lambda, \alpha) = \lambda(\alpha |\beta_i| + (1 - \alpha) |\beta_i|^2) \quad (5)$$

where $\alpha \in [0, 1]$. When $\alpha = 1$, the elastic net reduces to LASSO, and with $\alpha = 0$, it becomes a ridge regression. Note also that every elastic net problem can be rewritten as a LASSO problem on augmented data. Therefore, for fixed λ and α , the computational difficulty of the elastic net solution is similar to the LASSO problem [9]. Just like in the other cases the value of λ is automatically selected via BIC or cross-validation. The additional parameter α takes one of the values from the grid $A = \{0.25, 0.5, 0.75\}$.

The elastic net was already applied in the context of forecasting electricity prices. In particular, it was first used by Uniejewski et al. [31], where it produced the most accurate forecast across all considered automated variable selection methods. Recently, the method is gaining in popularity and more and more authors decide to estimate their forecasting models with elastic net [for example 1, 4, 21, 22]

3.2.5. FLASH

The forward-LASSO adaptive shrinkage (FLASH) was introduced by Radchenko and James [27]. To be precise, the FLASH algorithm does not fall specifically under penalized optimization techniques, but it has an implicit penalty. It follows a hierarchical approach similar to least-angle regression [5] (LARS) and forward selection by incrementally adding one variable to the model, but with a unique feature of

adjusting the level of shrinkage at each step to optimize the selection of the next variable. The penalty for coefficients in the active set is specified as:

$$g(\boldsymbol{\beta}; \lambda, \gamma) = \lambda(1 - \gamma) |\beta_i| \quad (6)$$

In this paper, the additional parameter γ takes one of the values from the grid $G = \{0.25, 0.5, 0.75\}$.

3.2.6. LASSO

Least absolute shrinkage and selection operator (LASSO) was formally introduced by Tibshirani [30]. It has achieved great success in statistics and is widely used in various applications. The penalty function is defined as a ℓ^1 norm scaled by the tuning parameter λ :

$$g(\boldsymbol{\beta}; \lambda) = \lambda |\beta_i| \quad (7)$$

where λ is the tuning parameter. It indicates how significant the variables have to be to remain in the final model. While for $\lambda = 0$ the method reduces to ordinary least squares (OLS), as the parameter increases, more and more variables are considered irrelevant and eliminated from the final model. In this research LASSO is treated as fully-automated model as the value of λ parameter is selected automatically via BIC or cross-validation. LASSO is a regularization method that has been widely used in EPF but only in the past decade. Some of the earliest examples of research that employed LASSO in the context of EPF include [7, 15, 31, 35–37].

3.2.7. LQ

LQ regularization [8] is another type of penalty function, after elastic net, that tries to fit between LASSO and ridge regression. This time instead of mixing linear and quadratic terms, the LQ regularization proposes to use a ℓ^q norm, for $q \in (1, 2)$:

$$g(\boldsymbol{\beta}; \lambda, q) = \lambda |\beta_i|^q \quad (8)$$

In this paper, we consider three values of q from the grid $Q = \{1.25, 1.5, 1.75\}$. Note that due to the toolbox requirements [19], we had to impose the upper limit of the set of possible λ values. Thus, in the LQ regularization, the tuning parameter cannot exceed a threshold equal to 2.

3.2.8. MC+

MC+ is a penalized variable selection method for high-dimensional linear regression [18, 34]. Compared to LASSO, which is fast and continuous but biased, MC+ is nearly unbiased and accurate. The method consists of two elements: a minimax concave penalty (MCP) and a penalized linear unbiased selection (PLUS) algorithm. The MCP ensures the convexity of the penalized loss in sparse regions to the greatest extent possible, given certain thresholds for variable selection. The PLUS algorithm computes multiple exact local minimizers of a possibly nonconvex penalized loss function in a certain main branch of the graph of critical points of the penalized loss. The penalty is defined as follows:

$$g(\boldsymbol{\beta}; \lambda, \gamma) = \begin{cases} \lambda |\beta_i| - \frac{\beta_i^2}{2\gamma} & \text{for } |\beta_i| \leq \gamma\lambda \\ \frac{\lambda^2\gamma}{2} & \text{for } \gamma\lambda < |\beta_i| \end{cases} \quad (9)$$

The additional parameter $\gamma \geq 1$ takes one of the values from the grid $G = \{1, 3, 5\}$. As for all other methods, the tuning parameter λ is selected automatically via BIC or cross-validation.

3.2.9. Ridge

Ridge regression is one of the first regularization method introduced in statistics [10]. Ridge regression is very similar to LASSO but the linear penalty function is substituted with a quadratic norm:

$$g(\boldsymbol{\beta}; \lambda) = \lambda |\beta_i|^2 \quad (10)$$

where $\lambda \geq 0$. Note that for $\lambda = 0$, we get the standard OLS estimator; for $\lambda \rightarrow \infty$, all β_i 's tend to zero; while for intermediate values of λ , we are balancing two ideas: minimizing the RSS and shrinking the coefficients toward zero (and each other). The toolbox [19] requires us to specify the maximum value of λ . In this study, the maximum value has been set to 5.

Although ridge regression is one of the most classical forms of regularization, it very rarely appears in the EPF literature. Apart from Barnes and Balda [3] – the first paper that utilizes ridge regression in the context of evaluating the profitability of battery storage – the method has only been mentioned a couple of times [for example 20, 31].

3.2.10. SCAD

The smoothly clipped absolute deviation (SCAD) [6] corresponds to a quadratic spline function with knots at λ and $\alpha\lambda$. This penalty function leaves large values of β not excessively penalized and makes the solution continuous. The resulting solution is given by:

$$g(\boldsymbol{\beta}; \lambda, \alpha) = \begin{cases} \lambda |\beta_i| & \text{for } |\beta_i| \leq \lambda \\ \frac{-|\beta_i|^2 + 2\alpha\lambda |\beta_i| - \lambda^2}{2(\alpha - 1)} & \text{for } \lambda < |\beta_i| \leq \alpha\lambda \\ \frac{\lambda^2(1 + \alpha)}{2} & \text{for } |\beta_i| > \alpha\lambda \end{cases} \quad (11)$$

The additional parameter $\alpha \geq 2$ takes one of the values from the grid $A = \{10, 20, 30\}$.

3.2.11. Penalties comparison

In this study, the considered penalties can be classified into three groups, each representing a different approach to regularization. The first group includes standard methods widely used in the literature, represented by LASSO and ridge regression. These methods are based on simple ℓ^1 or ℓ^2 norms. Ridge regression uses a quadratic shrinkage factor that shrinks all β_i parameters toward zero, but not exactly to zero, making it suitable for models with a small number of inputs. Conversely, LASSO focuses on both parameter shrinkage and variable selection, making it well suited for parameter-rich models.

The second group includes methods that attempt to combine or enhance the capabilities of the standard approaches. These include adaptive LASSO, elastic net, FLASH, and LQ regularization. The comparison between them is shown in the left panel of Figure 2. Note that Adaptive LASSO and FLASH regularization are not included in the plot. Adaptive LASSO, which depends on the q parameter, has the same shape as LASSO, LQ, or ridge, and differs only in the scale, which depends on the value of β^* . FLASH, on the other hand, is not defined as a penalty, but rather as an optimization algorithm, so it cannot be easily plotted.




Figure 2. Shapes of different penalty functions (for $\lambda = 1$) of the considered regularization types depending on the value of the estimated weight (β). For clarity, the plot is divided into two panels. The left panel shows penalty functions for the elastic net and LQ regularizations compared to standard approaches such as LASSO and ridge regression. The right panel displays the penalties for all concave functions considered. Note that adaptive LASSO and FLASH regularization are not included in the plot (see Section 3.2.11)

The third group goes in a different direction and is represented by methods such as Clipped LASSO, Concave PF, MC+ or SCAD. These approaches use concave penalties, which means that the growth rate of the penalties decreases as the absolute value of β increases. In particular, for Clipped LASSO, MC+, and SCAD, the penalty function becomes constant after a certain threshold. This allows the weights corresponding to less important variables to shrink to zero very quickly, while not penalizing β_i too much for the most important inputs. The right panel of figure 2 illustrates the comparison between these concave penalties, giving an insight into their characteristics.

4. Results

Linear measures such as mean absolute error provide a simple and easy-to-understand assessment of the average forecasting error. However, linear measures have limitations such as the inability to capture the variance of errors which can lead to misleading conclusions. Therefore, in this study, we use the root mean square errors (RMSE). The measure is defined by the following equation:

$$\text{RMSE} = \sqrt{\frac{1}{24D} \sum_{d=1}^D \sum_{h=1}^{24} \hat{\varepsilon}_{d,h}^2}. \quad (12)$$

where, $\hat{\varepsilon}_{d,h} = P_{d,h} - \hat{P}_{d,h}$ is the forecasting error for day d and hour h .

4.1. OLS

Table 1 shows the results of the ARX and fARX models estimated using the OLS method. It can be seen that, surprisingly, the less complex ARX model produces more accurate results for the OMIE market in terms of RMSE, and only slightly worse for the EPEX market. The relatively poor performance of the parameter-rich model is due to the overfitting problem.

Table 1. RMSE of the forecast obtained with two considered baseline models estimated with OLS

EPEX		OMIE	
ARX	fARX	ARX	fARX
13.55	12.65	9.78	10.35

Bold text highlights the best model within each market.

The result of the OLS-estimated models presented in Table 1 should be considered as a reference point for other regularization-based estimation techniques.

4.2. Regularization results

4.2.1. Adaptive

Table 2. RMSE of the prediction obtained with two considered baseline models estimated with adaptive LASSO

	EPEX				OMIE			
	ARX		fARX		ARX		fARX	
	CV	BIC	CV	BIC	CV	BIC	CV	BIC
1	<u>14.61</u>	15.67	12.12	15.39	<u>9.79</u>	10.20	9.63	10.26
1.5	<u>15.06</u>	15.84	12.13	15.32	<u>9.81</u>	10.26	9.62	10.27
2	15.28	15.78	12.19	15.42	9.84	10.28	9.58	10.33
CV	14.68		12.17		9.79		<u>9.62</u>	

The best result within each model is underlined, while the best model within each market is highlighted in bold.

Table 2 shows the RMSE of the ARX and fARX models estimated using the adaptive LASSO method. The results for each market are divided into two sub-columns referring to two model structures (ARX and fARX). In addition, for each model, the results are reported for two λ selection methods. Finally, the first three rows refer to a model with a fixed q parameter, and the last row (CV) presents the result of the fully automated approach when both λ and q are selected by cross-validation. Some important conclusions can be drawn from the results presented in the Table 2:

- The parameter-rich fARX model outperforms the parsimonious ARX model in both markets, although the difference is more pronounced in the EPEX market.
- Cross-validation is a better method in terms of prediction accuracy for selecting the λ parameter than BIC for adaptive LASSO, especially when using the parameter-rich fARX model.

- Among all considered values of the additional parameter, $q = 1$ performs best on average in terms of RMSE. However, the fully automated approach is only slightly worse (by less than 0.5%) than the optimal value of the parameter. This indicates that CV can be successfully applied to select values of both parameters (λ and q).

4.2.2. Clipped LASSO

Table 3. RMSE of the prediction obtained with two considered baseline models estimated with clipped LASSO

	EPEX				OMIE			
	ARX		fARX		ARX		fARX	
	CV	BIC	CV	BIC	CV	BIC	CV	BIC
0.5	13.93	15.28	11.96	17.56	<u>9.74</u>	10.18	9.39	12.03
1.0	<u>13.91</u>	15.20	11.95	17.48	<u>9.80</u>	10.53	9.39	12.17
1.5	<u>13.91</u>	15.23	11.95	17.45	9.81	10.58	9.38	12.25
CV	13.91		11.95		9.77		<u>9.39</u>	

The best result within each model is underlined, while the best model within each market is highlighted in bold.

Table 3 reports the RMSE of the models estimated using the clipped LASSO method. The results are presented in a format similar to the tables in the previous sections. In particular, the results for each market are reported for two model structures (ARX and fARX) and two methods of α selection. There are several important conclusions that can be drawn from the results presented in the Table 3:

- Similar to the adaptive LASSO, the fARX model outperforms the parsimonious ARX model.
- Considering the method used to select the λ parameter, the CV beats the BIC by a large margin.
- The optimal value of the parameter γ varies across models and markets. This time the fully automated approach performs best for EPEX market and only slightly worse (for OMIE) than the best ex-post value of the parameter.

4.2.3. Concave PF

Table 4. RMSE of the prediction obtained with two considered baseline models estimated with concave PF

	EPEX				OMIE			
	ARX		fARX		ARX		fARX	
	CV	BIC	CV	BIC	CV	BIC	CV	BIC
5	13.93	15.37	11.97	17.67	9.81	10.46	9.42	11.87
15	13.92	15.22	11.96	17.59	<u>9.81</u>	10.51	9.41	11.99
25	<u>13.91</u>	15.22	11.96	17.54	9.81	10.54	9.41	11.92
CV	13.93		11.96		9.82		9.41	

The best result within each model is underlined, while the best model within each market is highlighted in bold.

Table 4 presents the RMSE of the forecasts obtained with models estimated using the concave PF regularization method. Some important conclusions can be drawn from the results presented in Table 4:

- As previously, the most accurate forecasts are produced by the parameter-rich model (fARX).
- In terms of selecting the value of the λ parameter, cross-validation clearly outperforms BIC, especially in the case of the fARX baseline model where BIC struggles to select the optimal λ value.
- The value $k = 25$ performs best on average in terms of predictive accuracy, outperforming all other considered values of the additional parameter. The fully automated approach shows only a small performance loss compared to the optimal value of the parameter (less than 0.1%), indicating that it is a good performing alternative.

4.2.4. Elastic net

Table 5. RMSE of the prediction obtained with two considered baseline models estimated with elastic net

	EPEX				OMIE			
	ARX		fARX		ARX		fARX	
	CV	BIC	CV	BIC	CV	BIC	CV	BIC
0.25	<u>13.81</u>	14.94	11.92	20.33	9.81	10.55	9.35	12.62
0.50	13.84	14.87	<u>11.95</u>	18.35	9.80	10.61	9.34	12.03
0.75	13.87	15.02	11.96	18.16	<u>9.79</u>	10.48	9.31	11.76
CV	13.87		11.96		9.79		9.32	

The best result within each model is underlined, while the best model within each market is highlighted in bold.

Table 5 displays the RMSE of the elastic net estimated models. The results for each market are presented in the same table format as in the previous sections. These results allow us to draw some important conclusions:

- The fARX model produces the most accurate forecasts. However, the ARX model performs better when the value of the λ parameter is selected with BIC.
- Cross-validation clearly outperforms BIC, with the difference being more visible for fARX model.
- The values $\alpha = 0.25$ and $\alpha = 0.75$ perform best for EPEX and OMIE, respectively. The fully automated approach tends to produce very similar results (in terms of prediction accuracy) to the $\alpha = 0.75$ model.

4.2.5. FLASH

Table 6 shows the RMSE of the forecasts obtained with models estimated using the FLASH regularization. The results are presented in a format similar to the tables in the previous sections. In particular, for each market, the results are divided into two model structures (ARX and fARX) and two methods of λ selection. From the results in Table 6, some important conclusions can be drawn:

- The fARX clearly outperforms the ARX model.
- Regarding the choice of the λ parameter, cross-validation outperforms BIC by a substantial margin, especially for the EPEX market.
- The value $\gamma = 0.25$ performs best on average for both markets. The fully automated approach tends to produce forecasts that are only slightly less accurate than the best model.

Table 6. RMSE of the prediction obtained with two considered baseline models estimated with FLASH algorithm

	EPEX				OMIE			
	ARX		fARX		ARX		fARX	
	CV	BIC	CV	BIC	CV	BIC	CV	BIC
0.25	13.92	15.35	<u>11.97</u>	15.73	<u>9.83</u>	10.52	<u>9.40</u>	10.69
0.5	13.90	15.33	<u>11.98</u>	15.21	<u>9.85</u>	10.45	<u>9.41</u>	10.53
0.75	<u>13.81</u>	15.29	11.97	14.61	9.92	10.41	9.46	10.32
CV	13.82		11.98		9.92		9.42	

The best result within each model is underlined, while the best model within each market is highlighted in bold.

4.2.6. LASSO

Table 7. RMSE of the forecast obtained with two considered baseline models estimated with LASSO regularization

	EPEX				OMIE			
	ARX		fARX		ARX		fARX	
	CV	BIC	CV	BIC	CV	BIC	CV	BIC
	<u>13.91</u>	15.23	<u>11.94</u>	17.74	<u>9.81</u>	10.57	<u>9.39</u>	12.63

The best result within each model is underlined, while the best model within each market is highlighted in bold.

The RMSE of the models estimated using LASSO are presented in Table 7. The results are divided into two model structures (ARX and fARX) and two methods of selecting λ values for each market. Note that the LASSO regularization has only one parameter, so both CV and BIC are fully automated approaches. The results in Table 7 lead to several important conclusions:

- The fARX model produces the most accurate price predictions, but ARX outperforms fARX in some cases.
- Cross-validation beats BIC by a large margin for selecting the λ parameter, especially for the fARX model.

4.2.7. Lq

Table 8. RMSE of the prediction obtained with two considered baseline models estimated with adaptive LASSO.

	EPEX				OMIE			
	ARX		fARX		ARX		fARX	
	CV	BIC	CV	BIC	CV	BIC	CV	BIC
1.25	13.85	13.54	<u>11.92</u>	32.68	9.79	<u>9.77</u>	<u>9.33</u>	28.15
1.5	13.79	13.54	<u>11.92</u>	29.42	9.80	<u>9.77</u>	9.34	22.40
1.75	13.74	<u>13.54</u>	11.95	22.70	9.82	9.77	9.39	16.13
CV	13.85		11.93		9.78		9.34	

The best result within each model is underlined, while the best model within each market is highlighted in bold.

The RMSE of the predictions estimated with LQ regularization is shown in Table 8. The presentation of the results follows a similar format to the previous sections, where the results are divided into two model structures (ARX and fARX) and two methods of selecting the value of the λ parameter for each market. The first three rows in each panel correspond to a model with a fixed parameter α , while the last row (CV) represents the results of the fully automated approach. From the results in Table 8, several important conclusions can be drawn:

- The fARX model produces the best prediction in terms of forecast accuracy. However, it performs very weakly when BIC is used to select the value of the λ parameter.
- The performance of the two methods for selecting the optimal value of the λ parameter depends on the structure of the model. In particular, BIC performs better for the parsimonious ARX model but fails to select λ for the more complex fARX model.
- The value $q = 1.25$ performs best on average regarding RMSE. The fully automated approach performs only slightly worse than the optimal *ex-post* model for the fARX model, but substantially worse for the ARX structure.

4.2.8. MC+

Table 9. RMSE of the prediction obtained with two considered baseline models estimated with MC+

	EPEX				OMIE			
	ARX		fARX		ARX		fARX	
	CV	BIC	CV	BIC	CV	BIC	CV	BIC
1	13.98	15.59	12.00	16.27	9.83	10.16	9.41	10.52
3	13.93	15.35	11.99	16.98	<u>9.81</u>	10.47	9.41	11.21
5	<u>13.92</u>	15.27	11.98	17.11	<u>9.81</u>	10.45	9.40	11.36
CV	13.99		<u>12.00</u>		9.84		<u>9.40</u>	

The best result within each model is underlined, while the best model within each market is highlighted in bold.

Table 9 shows the RMSE of the models estimated by the MC+ regularization. The format of the table is the same as in the previous sections. There are some important conclusions that can be drawn from the results in the Table 9:

- The fARX model outperforms the ARX model on average for both methods of selecting the tuning parameter λ .
- Cross-validation consistently outperforms BIC for λ value selection by a large margin, especially in the EPEX market.
- The optimal value of the parameter γ varies across models, but on average the $\gamma = 5$ performs best. While the fully automated approach is slightly outperformed by the optimal parameter values in each case, it remains a reasonable alternative. This suggests that using a fully automated approach can lead to satisfying results without having to rely on expert parameter choices.

4.2.9. Ridge

Table 10. RMSE of the forecast obtained with two considered baseline models estimated with ridge regularization

EPEX				OMIE			
ARX		fARX		ARX		fARX	
CV	BIC	CV	BIC	CV	BIC	CV	BIC
13.70	<u>13.53</u>	<u>12.02</u>	22.79	9.83	<u>9.77</u>	<u>9.39</u>	15.99

The best result within each model is underlined, while the best model within each market is highlighted in bold.

In Table 10, we present the RMSE of the models estimated by ridge regression. The results are divided into two model structures (ARX and fARX) and two methods of selecting λ values for each market. Ridge regression, similar to LASSO, has only one parameter, and thus both CV and BIC are fully automated approaches. There are several important conclusions that can be drawn from the results presented in the Table 10:

- The parameter-rich fARX model produces the most accurate forecasts, but in the ARX outperforms fARX.
- In terms of selecting the value of the parameter λ , CV outperforms the information criteria for the fARX model, but BIC performs better for the parsimonious ARX model.

4.2.10. SCAD

Table 11. RMSE of the prediction obtained with two considered baseline models estimated with SCAD regularization

	EPEX				OMIE			
	ARX		fARX		ARX		fARX	
	CV	BIC	CV	BIC	CV	BIC	CV	BIC
10	<u>13.71</u>	14.75	12.18	15.00	9.98	10.13	<u>9.47</u>	10.10
20	13.82	15.03	12.15	14.85	9.91	10.12	<u>9.51</u>	10.22
30	13.86	14.91	<u>12.13</u>	15.06	<u>9.85</u>	10.10	9.56	10.38
CV	13.76		12.15		9.97		9.51	

The best result within each model is underlined, while the best model within each market is highlighted in bold.

Table 11 presents the RMSE of the models estimated using the SCAD regularization. The presentation of the results follows a similar format to the previous sections, where the results are divided into two model structures (ARX and fARX) and two methods of selecting the value of the λ parameter for each market. The first three rows in each panel correspond to a model with a fixed α parameter, while the last row (CV) represents the results of the fully automated approach. Several important conclusions can be drawn from the results presented in the Table 11:

- The fARX model outperforms the ARX model on average for both methods of selecting the tuning parameter λ .

- For selecting the value of the λ parameter, cross-validation consistently outperforms BIC by a large margin.
- It is impossible to identify a clear winner among models and markets regarding the optimal value of the α parameter, as it varies among them. Even though the fully automated approach is slightly inferior to the optimal parameter values in most cases, it is still a valuable option. This means that relying on a fully automated approach can yield satisfying results.

4.3. Summary of the results

Sections 4.2.1–4.2.10 present the results for each regularization technique separately. While the conclusions varied depending on the penalty function used, some common factors remained consistent across most regularization types. Based on the results in Tables 2–11, the most important conclusions can be summarized as follows:

- In general, the fARX model with cross-validation as the λ selection method produces the most accurate forecasts for both markets. Although BIC slightly outperforms CV in some cases, this is only true for the parsimonious ARX model, which on average produces less accurate forecasts compared to fARX.
- BIC is not recommended for selecting the tuning parameter λ , especially for models with a large number of parameters.
- For the vast majority of regularization types, the accuracy of the forecast obtained with the fully automated approach, where all parameters were selected by cross-validation, is very similar to the best model selected *ex-post*. This shows that the presented solutions are robust and can be applied to different markets and time series without pre-selecting their parameters.

In the previous sections, the accuracy of the forecasts was only compared within a certain penalty function. However, the purpose of this study is to identify the best type of regularization in the context of the EPF. To rank the penalty functions, Table 12 summarizes the results for all ten considered regularization types. Table 12 shows the percentage change compared to the OLS estimated model. The relative measure is defined as

$$r\text{RMSE}^{\text{reg}} = \frac{\text{RMSE}^{\text{reg}} - \text{RMSE}^{\text{OLS}}}{\text{RMSE}^{\text{OLS}}} \times 100\%$$

where RMSE^{reg} is the root mean squared error of the given regularization type. RMSE^{OLS} is the RMSE value for the predictions obtained with the corresponding OLS-estimated model.

Table 12 shows the percentage change compared to OLS estimation for the fully automated approaches and the best *ex-post* models. For each model and market, we present the result of the procedure when all parameters are selected with CV and the results underlined in Tables 2–11. Note that the negative value in Table 12 indicates that the model estimated with a given regularization type is better in terms of predictive accuracy than the same model estimated with OLS. Moreover, considering that the most common regularization technique in the EPF application is LASSO, the results of models that outperform LASSO-estimated equivalents are underlined.

Table 12. RMSE of the forecasts obtained using different regularization techniques compared to the models estimated using OLS [%]

	EPEX				OMIE			
	ARX		fARX		ARX		fARX	
	CV	best	CV	best	CV	best	CV	best
aLASSO	8.33	7.88	-3.83	-4.19	<u>0.13</u>	<u>0.10</u>	-7.13	-7.44
cLASSO	2.69	<u>2.65</u>	-5.56	-5.56	<u>0.00</u>	-0.36	<u>-9.32</u>	<u>-9.44</u>
CPF	2.82	<u>2.69</u>	-5.47	-5.49	0.41	0.37	-9.10	-9.15
EN	<u>2.41</u>	<u>1.97</u>	-5.48	<u>-5.79</u>	<u>0.19</u>	<u>0.17</u>	-10.00	-10.11
FLASH	<u>2.00</u>	<u>1.97</u>	-5.34	-5.39	1.51	0.52	-9.01	-9.25
LASSO		2.66		-5.64		0.32		-9.29
LQ	2.26	<u>-0.03</u>	-5.71	-5.84	<u>0.09</u>	<u>-0.06</u>	-9.80	<u>-9.91</u>
MC	3.27	2.76	-5.17	-5.31	0.66	0.40	-9.21	-9.23
Ridge		-0.11		-4.98		-0.09		-9.29
SCAD	<u>1.57</u>	<u>1.20</u>	-4.01	-4.12	1.94	0.72	-8.16	-8.53

The results for each model are presented in two columns. The first column (CV) shows the result of the model with all parameters automatically selected by CV. The second column (best) shows the result of the model with the λ parameter selected by CV or BIC (the better one selected ex-post) and the other selected from the grid according to *ex-post* performance. The best result in each column is bolded and any one better than LASSO is underlined.

Several important conclusions can be drawn from the results reported in Table 12:

- The most accurate forecasts for the parsimonious ARX model are obtained with a fully automated approach using ridge regression and the BIC method to select the tuning parameter λ .
- Overall, the OLS estimation is hard to beat for the ARX model. Only two types of regularization managed to produce more accurate forecasts for both markets, and only by a small margin. In particular, the ARX model estimated with the LQ regularization with the q parameter selected ex-post outperformed the OLS estimation by 0.03% and 0.06% for the EPEX and OMIE markets, respectively.
- On the other hand, for the parameter-rich fARX model, estimation with any regularization function outperforms the OLS-estimated model for both markets.
- In the case of the fARX model, the LASSO estimation is very hard to beat for other regularization functions. Only the elastic net and LQ are able to consistently outperform LASSO.
- Across all the regularization types, the LQ regularization performs best on average. It produces more accurate forecasts than the LASSO-estimated models in all cases.
- The second best regularization type in terms of forecast accuracy is the elastic net. In particular, it outperforms LASSO in seven out of eight cases.
- It is worth noting the very good performance of Clipped LASSO for the OMIE market. It performs best of all regularization types for ARX and outperforms LASSO for fARX. However, it is outperformed by other methods for the EPEX market.

5. Conclusions

The study investigated the use of different types of penalties for regularization in the context of the EPF. It compares the performance of ten regularization techniques, including adaptive LASSO, clipped

LASSO, concave potential function, elastic net, forward LASSO adaptive shrinkage, LASSO, LQ, mini-max concave PLUS, ridge regression, and smoothly clipped absolute deviation, using two data sets from the German EPEX SPOT and Iberian OMIE markets and two different model structures.

The results show that for both markets, the parameter-rich fARX model with cross-validation as the parameter selection method produces the most accurate forecasts. Conversely, using BIC to select the optimal λ value is not advisable for the majority of regularized models. The study recommends using cross-validation to select all parameters in a fully automated manner. The accuracy of such models is similar to the best models selected *ex-post*, and at the same time it does not require additional expert knowledge.

OLS estimation is hard to beat for the parsimonious ARX model, but ridge regression slightly outperforms the benchmark. For the parameter-rich fARX model, the estimation with any regularization function far outperforms the OLS-estimated model for both markets. The analysis shows that the LQ and elastic net techniques perform better than others on average in terms of RMSE. In particular, they were the only types of penalty functions that consistently produced more accurate forecasts than the most commonly used LASSO.

Acknowledgments

This work was partially supported by the Ministry of Science and Higher Education (MNiSW, Poland) through Diamond Grant No. 0199/DIA/2019/48.

References

- [1] AGRAWAL, R. K., MUCHAHARY, F., AND TRIPATHI, M. M. Ensemble of relevance vector machines and boosted trees for electricity price forecasting. *Applied Energy* 250 (2019), 540–548.
- [2] ANTONIADIS, A., AND FAN, J. Regularization of wavelet approximations. *Journal of the American Statistical Association* 96, 455 (2001), 939–967.
- [3] BARNES, A. K., AND BALDA, J. C. Sizing and economic assessment of energy storage with real-time pricing and ancillary services. In *2013 4th IEEE International Symposium on Power Electronics for Distributed Generation Systems (PEDG)* (Rogers, AR, USA, 2013), IEEE, pp. 1-7.
- [4] CIARRETA, A., MUNIAIN, P., AND ZARRAGA, A. Do jumps and cojumps matter for electricity price forecasting? Evidence from the German-Austrian day-ahead market. *Electric Power Systems Research* 212 (2022), 108144.
- [5] EFRON, B., HASTIE, T., JOHNSTONE, I., AND TIBSHIRANI, R. Least angle regression. *The Annals of Statistics* 32, 2 (2004), 407–499.
- [6] FAN, J., AND LI, R. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association* 96, 456 (2001), 1348–1360.
- [7] GAILLARD, P., GOUDE, Y., AND NEDELLEC, R. Additive models and robust aggregation for GEFCom2014 probabilistic electric load and electricity price forecasting. *International Journal of Forecasting* 32, 3 (2016), 1038–1050.
- [8] GRASMAIR, M., HALTMEIER, M., AND SCHERZER, O. Sparse regularization with l^q penalty term. *Inverse Problems* 24, 5 (2008), 055020.
- [9] HASTIE, T., TIBSHIRANI, R., AND WAINWRIGHT, M. *Statistical Learning with Sparsity: The Lasso and Generalizations*. CRC Press, 2015.
- [10] HOERL, A. E., AND KENNARD, R. W. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* 12, 1 (1970), 55–67.
- [11] HUBICKA, K., MARCJASZ, G., AND WERON, R. A note on averaging day-ahead electricity price forecasts across calibration windows. *IEEE Transactions on Sustainable Energy* 10, 1 (2019), 321–323.
- [12] JANCZURA, J., AND WÓJCIK, E. Dynamic short-term risk management strategies for the choice of electricity market based on probabilistic forecasts of profit and risk measures. The German and the Polish market case study. *Energy Economics* 110 (2022), 106015.
- [13] KATH, C., AND ZIEL, F. The value of forecasts: Quantifying the economic gains of accurate quarter-hourly electricity price forecasts. *Energy Economics* 76 (2018), 411–423.

- [14] LAGO, J., MARCJASZ, G., DE SCHUTTER, B., AND WERON, R. Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy* 293 (2021), 116983.
- [15] LUDWIG, N., FEUERRIEGEL, S., AND NEUMANN, D. Putting big data analytics to work: Feature selection for forecasting electricity prices using the LASSO and random forests. *Journal of Decision Systems* 24, 1 (2015), 19–36.
- [16] MACIEJOWSKA, K. Portfolio management of a small RES utility with a structural vector autoregressive model of electricity markets in Germany. *Operations Research and Decisions* 32, 4 (2022), 75–90.
- [17] MAYER, K., AND TRÜCK, S. Electricity markets around the world. *Journal of Commodity Markets* 9 (2018), 77–100.
- [18] MAZUMDER, R., FRIEDMAN, J. H., AND HASTIE, T. SparseNet: Coordinate descent with nonconvex penalties. *Journal of the American Statistical Association* 106, 495 (2011), 1125–1138.
- [19] McILHAGGA, W. penalized: A MATLAB toolbox for fitting generalized linear models with penalties. *Journal of Statistical Software* 72, 6 (2016), 1–21.
- [20] MIRAKYAN, A., MEYER-RENSCHHAUSEN, M., AND KOCH, A. Composite forecasting approach, application for next-day electricity price forecasting. *Energy Economics* 66 (2017), 228 – 237.
- [21] MUNIAIN, P., AND ZIEL, F. Probabilistic forecasting in day-ahead electricity markets: Simulating peak and off-peak prices. *International Journal of Forecasting* 36, 4 (2020), 1193–1210.
- [22] NARAJEWSKI, M., AND ZIEL, F. Ensemble forecasting for intraday electricity prices: Simulating trajectories. *Applied Energy* 279 (2020), 115801.
- [23] NARAJEWSKI, M., AND ZIEL, F. Optimal bidding in hourly and quarter-hourly electricity price auctions: Trading large volumes of power with market impact and transaction costs. *Energy Economics* 110 (2022), 105974.
- [24] NIKOLOVA, M. Local strong homogeneity of a regularized estimator. *SIAM Journal on Applied Mathematics* 61, 2 (2000), 633–658.
- [25] NITKA, W., AND WERON, R. Combining predictive distributions of electricity prices. Does minimizing the CRPS lead to optimal decisions in day-ahead bidding? *Operations Research and Decisions* 33, 3 (2023), 105–118.
- [26] PARK, M. Y., AND HASTIE, T. L_1 -Regularization path algorithm for generalized linear models. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 69, 4 (2007), 659–677.
- [27] RADCHENKO, P., AND JAMES, G. M. Improved variable selection with Forward-Lasso adaptive shrinkage. *The Annals of Applied Statistics* 5, 1 (2011), 427–448.
- [28] SCHMIDT, M., FUNG, G., AND ROSALES, R. *Optimization Methods for L_1 -Regularization*. Technical Report No. 19, University of British Columbia, 2009.
- [29] SERAFIN, T., UNIEJEWSKI, B., AND WERON, R. Averaging predictive distributions across calibration windows for day-ahead electricity price forecasting. *Energies* 12, 13 (2019), 2561.
- [30] TIBSHIRANI, R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 58, 1 (1996), 267–288.
- [31] UNIEJEWSKI, B., NOWOTARSKI, J., AND WERON, R. Automated variable selection and shrinkage for day-ahead electricity price forecasting. *Energies* 9, 8 (2016), 621.
- [32] UNIEJEWSKI, B., WERON, R., AND ZIEL, F. Variance stabilizing transformations for electricity spot price forecasting. *IEEE Transactions on Power Systems* 33 (2018), 2219–2229.
- [33] WERON, R. Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting* 30, 4 (2014), 1030–1081.
- [34] ZHANG, C.-H. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics* 38, 2 (2010), 894–942.
- [35] ZIEL, F. Forecasting electricity spot prices using LASSO: On capturing the autoregressive intraday structure. *IEEE Transactions on Power Systems* 31, 6 (2016), 4977–4987.
- [36] ZIEL, F. Iteratively reweighted adaptive lasso for conditional heteroscedastic time series with applications to AR-ARCH type processes. *Computational Statistics and Data Analysis* 100 (2016), 773–793.
- [37] ZIEL, F., STEINERT, R., AND HUSMANN, S. Efficient modeling and forecasting of electricity spot prices. *Energy Economics* 47 (2015), 98–111.
- [38] ZIEL, F., AND WERON, R. Day-ahead electricity price forecasting with high-dimensional structures: Univariate vs. multivariate modeling frameworks. *Energy Economics* 70 (2018), 396–420.
- [39] ZOU, H. The adaptive lasso and its oracle properties. *Journal of the American Statistical Association* 101, 476 (2006), 1418–1429.
- [40] ZOU, H., AND HASTIE, T. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 67, 2 (2005), 301–320.