

Extrapolating the long-term seasonal component of electricity prices for forecasting in the day-ahead market

Katarzyna Chęć¹
Bartosz Uniejewski¹
Rafał Weron¹

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

WORMS is a joint initiative of the Management Science departments
of the Wrocław University of Science and Technology,
Wyb. Wyspiańskiego 27, 50-370 Wrocław, Poland

Extrapolating the long-term seasonal component of electricity prices for forecasting in the day-ahead market

Katarzyna Chęć^a, Bartosz Uniejewski^a, Rafał Weron^{a,*}

^a*Department of Operations Research and Business Intelligence, Wrocław University of Science and Technology,
50-370 Wrocław, Poland*

Abstract

Recent studies provide evidence that decomposing the electricity price into the long-term seasonal component (LTSC) and the remaining part, predicting both separately, and then combining their forecasts can bring significant accuracy gains in day-ahead electricity price forecasting. However, not much attention has been paid to predicting the LTSC, and the last 24 hourly values of the estimated pattern are typically copied for the target day. To address this gap, we introduce a novel approach which extracts the trend-seasonal pattern from a price series extrapolated using price forecasts for the next 24 hours. We assess it using two 5-year long test periods from the German and Spanish power markets, covering the Covid-19 pandemic, the 2021/2022 energy crisis, and the war in Ukraine. Considering parsimonious autoregressive and LASSO-estimated models, we find that improvements in predictive accuracy range from 3% to 15% in terms of the root mean squared error and exceed 1% in terms of profits from a realistic trading strategy involving day-ahead bidding and battery storage.

Keywords: Electricity price forecasting, Long-term seasonal component, Day-ahead market, Combining forecasts, Trading strategy

1. Introduction

Seasonal decomposition is a fundamental technique in time series analysis and forecasting whose origins can be traced back to the middle of the 19th century (Hyndman and Athanasopoulos, 2021; Petropoulos et al., 2022). In its basic form, it decomposes the signal $P_{d,h}$ into three components: a deterministic slowly varying trend-cyclical component $T_{d,h}$ also called the *long-term seasonal component* (LTSC), a deterministic regularly repeating pattern $s_{d,h}$ called the (short-term) seasonal component, and the remaining part, i.e., the “residual”, also called the stochastic component $\eta_{d,h}$, in an additive fashion, i.e., $P_{d,h} = T_{d,h} + s_{d,h} + \eta_{d,h}$. Note that we use here the notation

*Corresponding author; email: rafal.weron@pwr.edu.pl

prevailing in the *electricity price forecasting* (EPF) literature, especially in day-ahead forecasting, where d refers to the day and h to the hour of delivery (Weron, 2014).

Although commonly used in multi-step ahead forecasting, seasonal decomposition has not been extensively utilized in day-ahead EPF until Nowotarski and Weron (2016) provided empirical evidence that it can be advantageous to decompose the electricity price $P_{d,h}$ into the deterministic LTSC and a stochastic component that includes the short-term seasonality, i.e., $Y_{d,h} = s_{d,h} + \eta_{d,h}$, predict both components separately, and then combine their forecasts: $\hat{P}_{d,h} = \hat{T}_{d,h} + \hat{Y}_{d,h}$. The rationale behind this approach is that the residuals obtained after seasonal decomposition better satisfy the assumptions underlying typical model architectures, including the regression-type models considered in this study. Interestingly, the approach has proven effective for both parsimonious autoregressive (AR; Afanasyev and Fedorova, 2019; Grossi and Nan, 2019; Shah et al., 2021; Zafar et al., 2022) and non-linear autoregressive neural network-type models (NARX; Marcjasz et al., 2019), as well as parameter-rich models estimated using the *least absolute shrinkage and selection operator* (LASSO; Jędrzejewski et al., 2021). However, to our best knowledge, all existing studies in the context of day-ahead EPF have relied on a naive prediction of the LTSC, where the last 24 hourly values of the estimated trend-seasonal pattern were simply copied for the target day.

To address this gap, we introduce a novel approach to extrapolating the LTSC for the next day. We show that extending the time series of electricity spot prices by the 24 hourly forecasts for the next day – obtained with a base model that does not involve seasonal decomposition – and then extracting the LTSC component for the 24 hours of the target day from the extended time series can lead to statistically significant accuracy gains compared to the naive approach used in the literature so far. We employ two techniques to extract the LTSC: a simple *moving average* (MA; see, e.g., Weron, 2014) and *wavelet smoothing* (also called *thresholding*; Percival and Walden, 2000). Moreover, we use averaging to capture the benefits of combining forecasts from different models (Petropoulos et al., 2024).

To allow for a thorough evaluation of the studied models, like Wagner et al. (2022), and in line with the best practices outlined in Lago et al. (2021), we use very long test sets from more than one market. Namely, we consider 9 years of data (2015-2023) from two major European power markets – EPEX in Germany and OMIE in Spain. The out-of-sample test period spans 5 years (2019-2023), covering the Covid-19 pandemic and the 2021-2022 energy crisis with skyrocketing prices of electricity. The day-ahead electricity price forecasts are obtained either using a parsimonious autoregressive expert model with exogenous variables (ARX; Billé et al., 2023; Gaillard et al., 2016; Maciejowska et al., 2021; Ziel and Weron, 2018) or a parameter-rich, LASSO-estimated regression model (LEAR; Lago et al., 2021; Wagner et al., 2022). We assess the significance of differences in predictive performance using the multivariate variant of the Diebold-Mariano (DM) test, as introduced by Ziel and Weron (2018). In line with a recent trend in the EPF literature (Maciejowska et al., 2023), we also consider a realistic trading strategy involving day-ahead bidding and battery storage in order to quantify the benefits in monetary terms.

The remainder of the paper is structured as follows. In Section 2 we describe the datasets, then in Section 3 we present the methodology, i.e., describe the ARX and LEAR models, discuss the methods used to estimate and extrapolate the LTSC, and briefly explain the transformation used to stabilize the variance of the stochastic component. In Section 4 we discuss the main results.

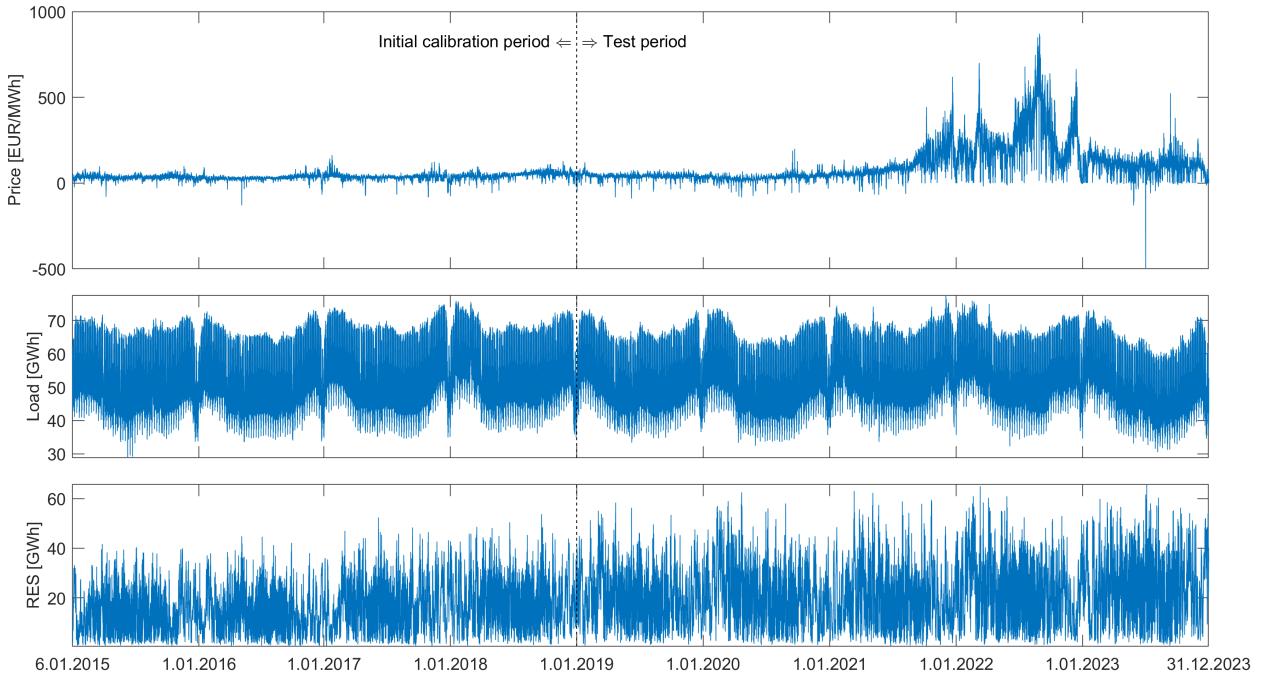


Figure 1: Electricity prices (*top*), day-ahead load forecasts (*middle*) and day-ahead RES generation forecasts (*bottom*) in the German electricity market (EPEX). The end of the initial calibration window is marked by the vertical dashed line.

Finally, in Section 5 we wrap up the findings and conclude.

2. Data

We consider data from two major European power markets – EPEX in Germany (see Figure 1) and OMIE in Spain (see Figure 2). The German market is one of the more studied ones, likely due to the central position in Europe, large wind and solar penetration amounting to over 45% of net electricity generation in 2023 (ENTSO-E, 2024), challenging price dynamics with abundant negative prices, and high trading liquidity in most segments (Caldana et al., 2017; Hagfors et al., 2016; Janczura, 2024; Kath and Ziel, 2018; Marcjasz et al., 2023; Narajewski and Ziel, 2020; Uniejewski, 2024). Although less studied in the literature (Diaz and Planas, 2016; García-Martos et al., 2015; Gianfreda et al., 2020; Lipiecki et al., 2024), the Spanish market is very interesting because of the rapid increase in the share of solar generation – from 5.1% of net electricity generation in 2015 to 16.5% in 2023 (ENTSO-E, 2024).

Both datasets are of hourly resolution and include day-ahead electricity prices as well as forecasts of electricity consumption and generation from *renewable energy sources* (RES; here solar and wind) downloaded from the ENTSO-E transparency platform (ENTSO-E, 2024). With the growing share of RES in the generation stack and the impact of its volume on electricity prices, day-ahead forecasts of solar and wind generation have become fundamental inputs to EPF models (Billé et al., 2023; Janczura and Wójcik, 2022; Westgaard et al., 2021). Both datasets have been preprocessed to handle missing/duplicate values that occur when switching to/from daylight

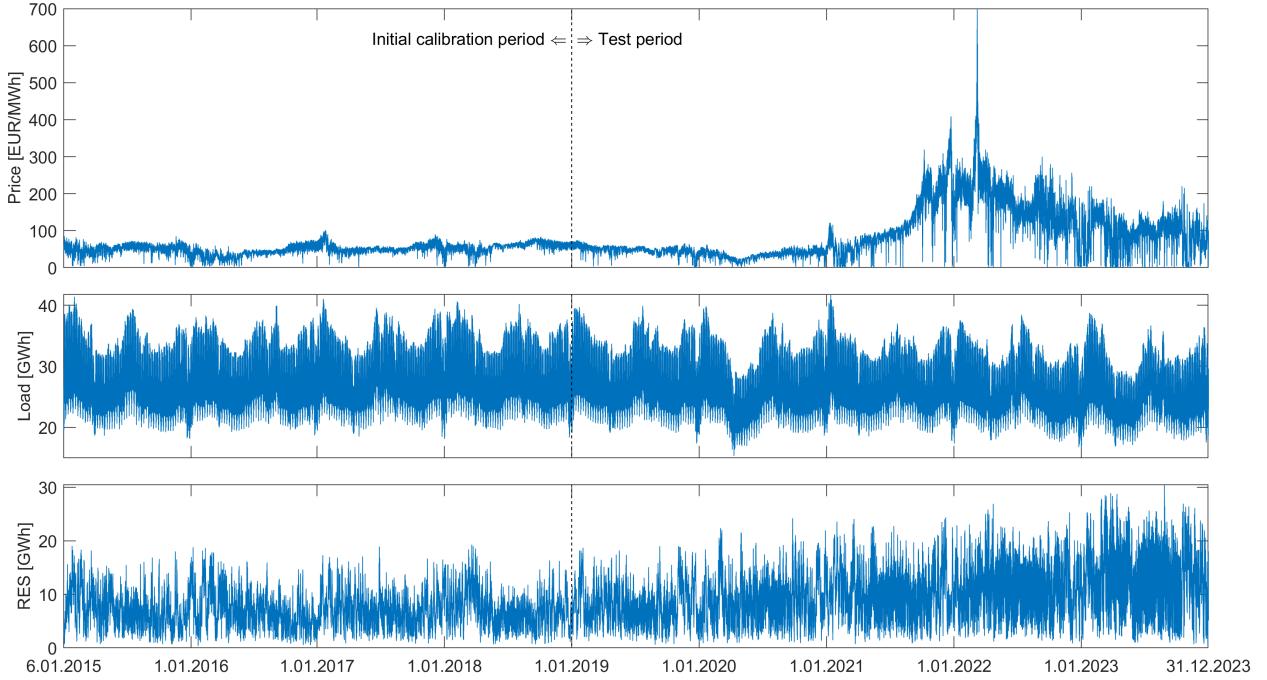


Figure 2: Electricity prices (*top*), day-ahead load forecasts (*middle*) and day-ahead RES generation forecasts (*bottom*) in the Spanish electricity market (OMIE). The end of the initial calibration window is marked by the vertical dashed line.

saving time – missing observations have been replaced by the arithmetic mean of two neighboring values, while duplicate ones have been replaced by their arithmetic mean (Weron, 2014). To make the marginal distributions of prices and fundamental variables less leptokurtic and thus more suitable for modeling, we use the area hyperbolic sine (asinh) transformation (Uniejewski et al., 2018, for details see Section 3.2). Note that unlike the logarithmic transformation that is popular in finance, the asinh can also handle negative values.

Furthermore, in line with the EPF literature, we use a rolling window scheme (Grossi and Nan, 2019; Lago et al., 2021; Marcjasz et al., 2018) to obtain day-ahead forecasts of electricity prices. Inspired by recent research which involved parameter-rich models (Marcjasz et al., 2023), we use a calibration window of 208 weeks or 1456 days. To obtain electricity price forecasts for the first day of the test period, i.e., 1.01.2019, the models are calibrated to data spanning from 6.01.2015 to 31.12.2018. Next, the calibration window is moved forward one day and the forecasts are made for 2.01.2019. This process continues until electricity price forecasts are obtained for the last day in the test set, i.e., 31.12.2023.

3. Methodology

In this study, we compute point forecasts of the 24 hourly prices for day d using the information known at ca. 8 am on day $d - 1$. Although the forecast horizon formally spans 16-40 hours, we actually make one step (day) ahead predictions in a multivariate modeling framework that uses a set of 24 interrelated models, one for each hour of day d (Ziel and Weron, 2018).

More specifically, we use a multi-stage procedure. First, the electricity price is decomposed into the LTSC and a stochastic component that includes the short-term seasonality: $P_{d,h} = T_{d,h} + Y_{d,h}$. Then we apply a so-called *variance stabilizing transformation* (VST; see Section 3.2 and Uniejewski et al., 2018) to $Y_{d,h}$ and obtain standardized series $y_{d,h}$. The latter is predicted using one of the models described in Section 3.1 to yield $\hat{y}_{d,h}$. In the last two stages, we apply the inverse VST and add the LTSC forecast to obtain the final price forecast $\hat{P}_{d,h}$:

$$P_{d,h} \xrightarrow[\text{Sec. 3.3}]{-\text{LTSC}} Y_{d,h} \xrightarrow[\text{Sec. 3.2}]{\text{VST}} y_{d,h} \xrightarrow[\text{Sec. 3.1}]{\text{Predict}} \hat{y}_{d,h} \xrightarrow[\text{Sec. 3.2}]{\text{VST}^{-1}} \hat{Y}_{d,h} \xrightarrow[\text{Sec. 3.5}]{+\text{LTSC}} \hat{P}_{d,h}. \quad (1)$$

Note that the data are transformed after removing the LTSC and not vice versa, as this has been found to yield more accurate predictions in EPF (Jędrzejewski et al., 2021). Note also that for the benchmark models that do not use seasonal decomposition, we apply the VST directly to the raw price series, thus omitting the first and last stages of the above procedure.

3.1. Forecasting with regression models

The day-ahead electricity price forecasts are obtained using two model classes. The first is a parsimonious autoregressive expert model with exogenous variables (ARX), originally proposed by Misiorek et al. (2006), later modified and compared in a number of EPF studies under different names and acronyms (Billé et al., 2023; Gaillard et al., 2016; Maciejowska et al., 2021; Maciejowska and Nowotarski, 2016; Taylor, 2021; Ziel, 2016; Ziel and Weron, 2018). In the ARX model the electricity price for day d and hour h is given by the following formula:

$$y_{d,h} = \underbrace{\beta_1 y_{d-1,h} + \beta_2 y_{d-2,h} + \beta_3 y_{d-7,h}}_{\text{autoregressive effects}} + \underbrace{\beta_4 y_{d-1,24}}_{\text{end-of-day}} + \underbrace{\beta_5 y_{d-1}^{\min} + \beta_6 y_{d-1}^{\max}}_{\text{non-linear effects}} + \underbrace{\beta_7 X_{d,h}^1 + \beta_8 X_{d,h}^2}_{\text{exogenous variables}} + \underbrace{\sum_{j=1}^7 \beta_{h,j+8} D_j}_{\text{daily dummies}} + \varepsilon_{d,h}, \quad (2)$$

where the first three regressors account for the autoregressive effects of the prices for the same hour on days $d-1$, $d-2$ and $d-7$, $y_{d-1,24}$ provides information on the last known price level, i.e., midnight of day $d-1$, y_{d-1}^{\max} and y_{d-1}^{\min} stand for the maximum and minimum price of the previous day, $X_{d,h}^1$ and $X_{d,h}^2$ are the exogenous variables – respectively the day-ahead forecasts of the system-wide load and RES generation, D_1, \dots, D_7 are daily dummies, and $\varepsilon_{d,h}$ is the noise term. The coefficients β_j are estimated using *ordinary least squares* (OLS).

The second model class is a parameter-rich LASSO-estimated regression (LEAR) introduced to the EPF literature by Uniejewski et al. (2016) and Ziel (2016), and later used by Lago et al. (2018), Maciejowska et al. (2023), Wagner et al. (2022) and Ziel and Weron (2018), among others.

Here, we consider a variant proposed by Lago et al. (2021), who coined the acronym LEAR:

$$\begin{aligned}
y_{d,h} = & \underbrace{\sum_{i=1}^{24} (\beta_{h,i} y_{d-1,i} + \beta_{h,i+24} y_{d-2,i} + \beta_{h,i+48} y_{d-3,i} + \beta_{h,i+72} y_{d-7,i})}_{\text{autoregressive effects}} \\
& + \underbrace{\sum_{i=1}^{24} (\beta_{h,i+96} X_{d,i}^1 + \beta_{h,i+120} X_{d-1,i}^1 + \beta_{h,i+144} X_{d-7,i}^1)}_{\text{exogenous variable \#1 and its lags}} \\
& + \underbrace{\sum_{i=1}^{24} (\beta_{h,i+168} X_{d,i}^2 + \beta_{h,i+192} X_{d-1,i}^2 + \beta_{h,i+216} X_{d-7,i}^2)}_{\text{exogenous variable \#2 and its lags}} + \underbrace{\sum_{j=1}^7 \beta_{h,240+j} D_j}_{\text{daily dummies}} + \varepsilon_{d,h},
\end{aligned} \tag{3}$$

where the first 96 regressors are autoregressive terms that include prices from all hours of days $d - 1$, $d - 2$, $d - 3$ and $d - 7$, and the following 2×72 regressors are all hourly values of the exogenous variables X^1 and X^2 for days d , $d - 1$ and $d - 7$. Such a model structure allows all cross-hour dependencies to be incorporated into the price forecasts. The LEAR model is estimated using the least absolute shrinkage and selection operator (LASSO) of Tibshirani (1996), which automatically selects the most relevant regressors for predicting $y_{d,h}$. Although many different regularization methods have been proposed in the literature, Uniejewski (2024) identified LASSO as a parsimonious yet robust and well-performing variant in a comprehensive EPF evaluation study.

3.2. Transforming the time series

Along with seasonal decomposition, data transformation is the primary preprocessing technique in time series analysis (Hyndman and Athanasopoulos, 2021). Its purpose is to remove known sources of variation, make the data more consistent across the sample, and – particularly in EPF – allow the handling of close to zero or negative values (Diaz and Planas, 2016; Janczura, 2024; Kath and Ziel, 2018). For comprehensive comparisons of different variance stabilizing transformations (VSTs) see Ciarreta et al. (2022), Shi et al. (2021) and Uniejewski et al. (2018), and for a discussion of the order of applying seasonal decomposition and VSTs before model calibration see Jędrzejewski et al. (2021).

Following Lago et al. (2021) and Ziel and Weron (2018), we use here the *area hyperbolic sine* (asinh) transformation with the (median, MAD) normalization:

$$y_{d,h} = \text{asinh} \left(\frac{Y_{d,h} - \text{Med}_\tau}{1.4826 \cdot \text{MAD}_\tau} \right), \tag{4}$$

where

$$\text{asinh}(x) = \log \left(x + \sqrt{x^2 + 1} \right), \tag{5}$$

Med_τ is the median and MAD_τ is the median absolute deviation of $Y_{d,h}$ in the calibration window τ , and 1.4826 is the inverse of the 75th percentile of the standard normal distribution; this factor ensures asymptotic normal consistency with the standard deviation (Uniejewski et al., 2018).

Once $\hat{y}_{d,h}$ is computed, we apply the inverse transformation:

$$\hat{Y}_{d,h} = 1.4826 \cdot \text{MAD}_\tau \cdot \sinh(\hat{y}_{d,h}) + \text{Med}_\tau, \tag{6}$$

where \sinh is the hyperbolic sine. As Narajewski and Ziel (2020) note, the latter is not the correct inverse transformation since, given random variable X , $\mathbb{E}\sinh(X)$ does not have to equal $\sinh(\mathbb{E}X)$. Nevertheless, the difference between Eq. (6) and the correct inverse VST is not substantial and is generally ignored in the literature.

3.3. Extracting the LTSC

When Nowotarski and Weron (2016) introduced the seasonal component approach, they used two techniques for extracting the long-term seasonal component (LTSC): the Hodrick-Prescott filter (HP; for sample applications in EPF see, e.g., Caldana et al., 2017; Lisi and Nan, 2014; Marcjasz et al., 2019; Zafar et al., 2022) and wavelet smoothing. Here, we replace the HP filter with a much simpler but equally effective simple moving average.

In the *moving average* (MA) approach the LTSC is approximated by averaging observations within a window centered at $t = 24d + h$:

$$\hat{T}_{d,h} = \hat{T}_t = \frac{1}{m} \sum_{j=-k}^k P_{t+j}, \quad (7)$$

where $m = 2k + 1$ is the width of the window. The latter affects the shape of the trend-seasonal component, i.e., low m yields a more volatile series, while high m smooths out the fluctuations. In our study, we consider five smoothing levels: 1, 7, 28, 56, and 91 days, corresponding to window sizes ranging from $m = 24 + 1$ (about 1 day) to $m = 91 \cdot 24 + 1$ (about 3 months). This allows us to obtain models that react differently to sudden changes in market conditions. Such an approach is similar to averaging across calibration windows of different lengths in macroeconomics (Pesaran and Timmermann, 2007) or in EPF (Hubicka et al., 2019).

On the other hand, *wavelet smoothing* (also called *thresholding*; Percival and Walden, 2000) applies the discrete wavelet transform to decompose the original series into a sum of the so-called approximation series S_J capturing the general trend and a set of detail series D_1, D_2, \dots, D_J . At each step $j = 1, \dots, J$ of this iterative procedure, details D_j of a given frequency are removed to yield a smoother, but twice shorter signal S_j . In this study, we consider five smoothing levels S_J with $J \in (5, 7, 9, 10, 11)$, corresponding to time scales ranging from $2^5 = 32$ hours (or “1 day”) to $2^{11} = 2048$ hours (or “3 months”). Similar to the MA approach, the use of different smoothing levels allows us to capture changes in market conditions. Following Nowotarski and Weron (2016), we use the Daubechies wavelet family of order 24. For sample applications of wavelet smoothing in EPF see Afanasyev and Fedorova (2019), Grossi and Nan (2019), Jędrzejewski et al. (2021) Lisi and Nan (2014) and Marcjasz et al. (2019), among others.

3.4. Combining forecasts

The contemporary forecasting literature agrees that combining predictions from different models generally improves forecasting accuracy (Atiya, 2020; Petropoulos et al., 2022). The same has been reported for energy forecasting (Hong et al., 2020), and EPF in particular (Berrisch and Ziel, 2024; Hubicka et al., 2019; Lago et al., 2021; Nitka and Weron, 2023). In this study, we use simple arithmetic averaging to combine forecasts obtained with (i) different levels of decomposition and (ii) different methods of decomposing the data. As a result, we consider three different classes of forecasting models and denote them by suffixes. The average the forecasts obtained for the

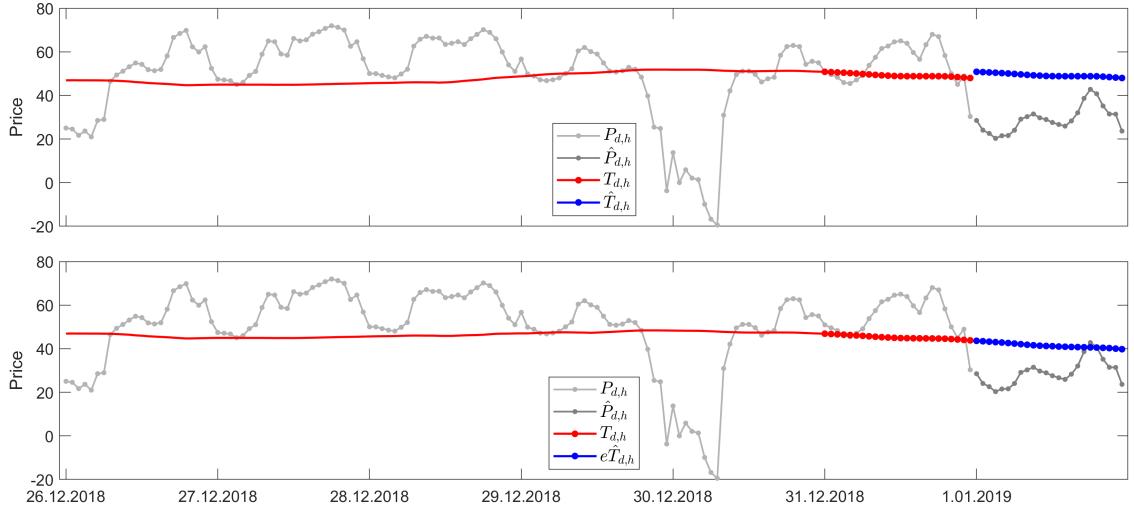


Figure 3: Comparison of the naive LTSC day-ahead forecast ($\hat{T}_{d,h}$; top panel) and the introduced in this study extrapolated LTSC forecast ($e\hat{T}_{d,h}$; bottom panel). In both cases the LTSC is calculated using the moving average (MA) method with a 7-day smoothing window ($m = 24 \cdot 7 + 1$).

five moving average windows is denoted by **-MA**. Similarly, the average of the forecasts obtained for the five different levels of wavelet smoothing is denoted by **-S**. Finally, the average of all 10 forecasts is denoted by **-MAS**. Although more sophisticated ensembling methods have been considered in the literature, the simple arithmetic average has been repeatedly shown to be robust and competitive (Raviv et al., 2015; Marcjasz et al., 2018; Uniejewski and Maciejowska, 2023).

3.5. Extrapolating the LTSC

To our best knowledge, all studies that use the seasonal component approach introduced by Nowotarski and Weron (2016) rely on a naive prediction of the LTSC, where the last 24 hourly values of the estimated trend-seasonal pattern are simply copied for the target day. Depending on the pool of individual forecasts, we denote this approach by **SCARX-*** or **SCLEAR-***, where $*$ takes the value MA, S or MAS, see Section 3.4, when referring to the model defined by Eq. (2) or Eq. (3), respectively.

In this paper, we introduce a new approach to predicting the LTSC for the next day by extrapolating the input price vector before calculating the LTSC. The proposed algorithm consists of the following steps:

1. Compute the 24 hourly electricity price forecasts for the next day using the model without seasonal decomposition, i.e., ARX or LEAR.
2. Append these forecasts to the prices in the calibration window, increasing the length of the calibration window by 24 observations.
3. Use the extrapolated price vector as the input to one of the approaches to extracting the LTSC described in Section 3.3.

By following this procedure, we obtain the LTSC forecast for the target day along with the seasonal component of the prices in the calibration window. Therefore, this approach preserves the

continuity of the LTSC vector, which transitions smoothly from the calibration period to the target day, see Figure 3. In our study, the seasonal component models that use the extrapolated LTSC approach are denoted by a prefix **e**: **eSCARX-*** for ARX-based models and **eCLEAR-*** for LEAR-based models.

Note that we also considered a simpler method to extrapolate the LTSC, where the price series is extended for the next 24 hours using a naive forecast, i.e., $\hat{P}_{d,h} = P_{d-1,h}$, not an ARX or LEAR forecast as in steps 1-2 of the above algorithm. However, this approach resulted in inferior performance and its results are not reported here.

4. Results

We evaluate the results using two metrics of prediction accuracy and a statistical test to assess the statistical significance of the results obtained (Lago et al., 2021; Petropoulos et al., 2022). For ensuring the reliability of our forecasts, we use 5-year long test periods, i.e., from 1.01.2019 to 31.12.2023, for both datasets.

4.1. Prediction accuracy

We use the two most popular error metrics in the EPF literature (Ciarreta et al., 2022; Lago et al., 2021; Maciejowska et al., 2023; Weron, 2014; Ziel, 2016) to measure the prediction accuracy of point forecasts, namely the Mean Absolute Error:

$$\text{MAE} = \frac{1}{24N_d} \sum_{d=1}^{N_d} \sum_{h=1}^{24} |P_{d,h} - \hat{P}_{d,h}|, \quad (8)$$

and the Root Mean Square Error (RMSE)

$$\text{RMSE} = \sqrt{\frac{1}{24N_d} \sum_{d=1}^{N_d} \sum_{h=1}^{24} (P_{d,h} - \hat{P}_{d,h})^2}, \quad (9)$$

where $N_d = 1826$ stands for the number of days in the out-of-sample test period, i.e., 1.01.2019–31.12.2023, see Figures 1 and 2, and $P_{d,h}$ and $\hat{P}_{d,h}$ respectively denote the actual and predicted price for day d and hour h .

4.1.1. Performance across the whole test period

In Tables 1 and 2 we report the MAE and RMSE errors for the considered models over the 5-year test period. Note, that cells are colored (red → high, green → low) independently for each market (EPEX, OMIE) and model class (ARX, LEAR). For instance, in Table 1 the MAE of 11.309 is red because it is the highest (i.e., worst) value among LEAR-class models for the OMIE dataset. In both tables columns labeled “%chng.” show the percentage difference between the error (MAE or RMSE) for an extrapolated LTSC-type model (eSCARX, eCLEAR) and a naive LTSC-type model (SCARX, SCLEAR). Recall from Section 3.4 that suffixes -MA, -S and -MAS denote combined predictions for three pools of individual forecasts.

Several important conclusions can be drawn:

Table 1: Mean Absolute Errors (MAE) for the considered models over the 5-year test period (01.01.2019–31.12.2023). Cells are colored (red → high, green → low) independently for each market (EPEX, OMIE) and model class (ARX, LEAR). Columns labeled “%chng.” show the percentage difference between the MAE for an eSC-type model and an SC-type model in the two columns on the left.

LTSC	EPEX						OMIE					
	ARX			LEAR			ARX			LEAR		
None	17.849			14.761			12.056			11.309		
SCARX	eSCARX	%chng.	SCLEAR	eSCLEAR	%chng.	SCARX	eSCARX	%chng.	SCLEAR	eSCLEAR	%chng.	
-MA	18.373	17.011	-7.70%	15.809	14.116	-11.33%	11.706	11.373	-2.89%	11.029	10.416	-5.71%
-S	18.182	17.297	-4.99%	16.155	14.062	-13.87%	11.569	11.130	-3.87%	10.914	10.167	-7.09%
-MAS	18.041	16.870	-6.71%	15.717	13.820	-12.86%	11.556	11.163	-3.46%	10.874	10.167	-6.72%

Table 2: Root Mean Square Errors (RMSE) for the considered models over the 5-year test period (01.01.2019–31.12.2023). Cells are colored (red → high, green → low) independently for each market (EPEX, OMIE) and model class (ARX, LEAR). Columns labeled “%chng.” show the percentage difference between the RMSE for an eSC-type model and an SC-type model in the two columns on the left.

LTSC	EPEX						OMIE					
	ARX			LEAR			ARX			LEAR		
None	31.531			26.861			20.409			19.690		
SCARX	eSCARX	%chng.	SCLEAR	eSCLEAR	%chng.	SCARX	eSCARX	%chng.	SCLEAR	eSCLEAR	%chng.	
-MA	34.352	31.555	-8.49%	30.112	26.927	-11.18%	20.472	19.810	-3.29%	19.660	18.616	-5.45%
-S	33.735	31.443	-7.04%	30.717	26.436	-15.01%	20.267	19.471	-4.00%	19.299	18.192	-5.91%
-MAS	33.567	31.008	-7.93%	29.877	26.130	-13.40%	20.207	19.504	-3.54%	19.239	18.179	-5.66%

- All extrapolated LTSC-based models (eSCARX, eSCLEAR) beat their naive LTSC counterparts (SCARX, SCLEAR) by a wide margin.
- The extrapolated LTSC-based models (eSCARX, eSCLEAR) are able to outperform the corresponding benchmark models without seasonal decomposition.
- The moving average (MA) approach is confirmed to be a very useful tool for identifying the LTSC.
- The LEAR-based models always outperform the corresponding ARX-based models.

In more detail, for models based on the parsimonious ARX structure, the reductions range from 2.89% to 7.70% in MAE and from 3.29% to 8.49% in RMSE. For the parameter-rich LEAR-type models, the improvements are even more striking – from 5.71% to 13.87% in MAE and from 5.45% to 15.01% in RMSE. This shows the high effectiveness of the method introduced in this paper compared to the naive approach used so far. This is especially true for the German EPEX market, where much higher improvements in accuracy are observed, see the 7th column in Tables 1 and 2. Interestingly, the improvements with respect to the classical ARX and LEAR models that do not use seasonal decomposition are higher for the Spanish OMIE market and reach up to 10.65% for SCLEAR-S and -MAS models in terms of the MAE and 7.98% for the SCLEAR-MAS model in terms of the RMSE; not reported in the Tables. Apparently, the SCARX and SCLEAR models, which naively extrapolate the LTSC for the next day, cannot cope with the extreme electricity price

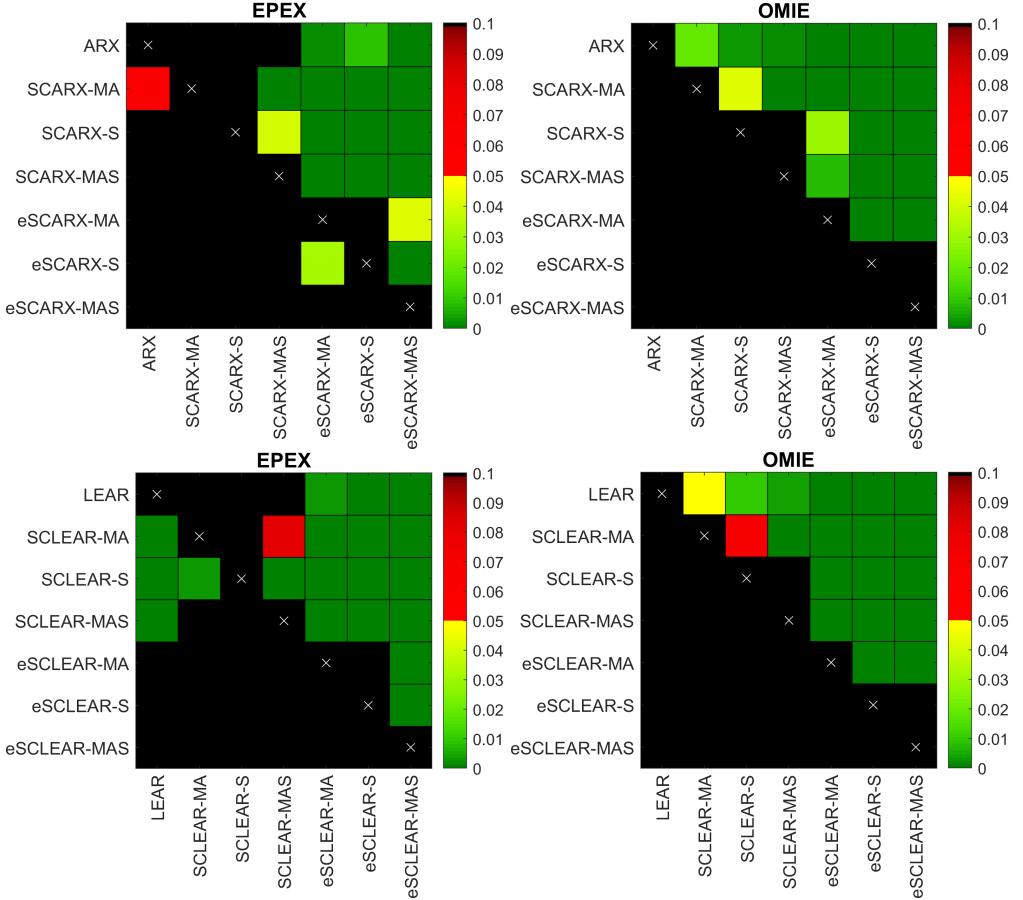


Figure 4: Results of the Diebold-Mariano (DM) test with the multivariate loss differential series (10) and the $\|\cdot\|_1$ norm, separately for the two datasets (EPEX – left, OMIE – right) and model classes (ARX – top, LEAR – bottom). Like in Lago et al. (2021) and Ziel and Weron (2018), we use a heat map to indicate the range of the p -values. The closer they are to zero (dark green) the more significant is the difference between the forecasts of a model on the X -axis (better) and the forecasts of a model on the Y -axis (worse); p -values ≥ 0.10 are marked in black.

dynamics in Germany over the studied period, see Figure 1.

In the case of the EPEX dataset, the best performing averaging approach is -MAS, which combines all MA- and wavelet-based forecasts. It is worth noting that only the extrapolated LTSC approach is able to outperform the corresponding benchmarks in Germany, as the SC-type models provide less accurate price forecasts. For the OMIE market, all three averaging schemes -MA, -S, and -MAS show good performance. This time, both the naive and the extrapolated LTSC approaches outperform the corresponding benchmarks, although the latter are more accurate.

4.1.2. Statistical significance

Following Ziel and Weron (2018), we use the multivariate variant of the Diebold and Mariano (1995) test (DM) to assess the statistical significance of differences in predictive performance. This version of the DM test is an asymptotic z -test of the hypothesis that the mean of the ‘daily’

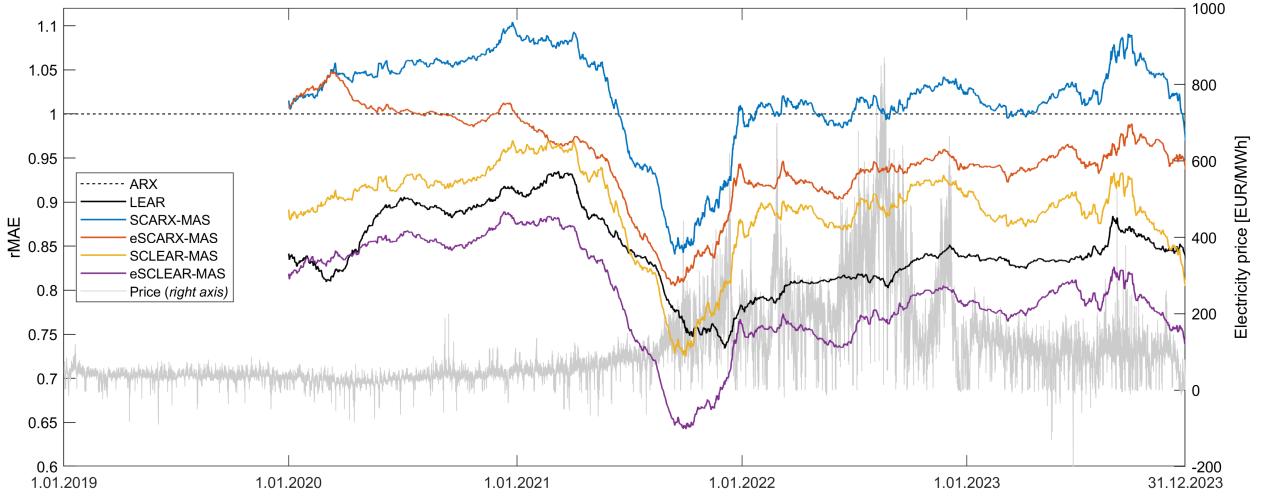


Figure 5: Rolling 365-day *relative mean absolute errors* (rMAE; *left axis*) for the German (EPEX) market with respect to MAE of the ARX model, see Eq. (11). Values for 1.01.2020 are the relative MAEs for 1.01.2019–31.12.2019, values for 2.01.2020 are the relative MAEs for 2.01.2019–1.01.2020, etc. The gray curve is the mean daily electricity price (in EUR/MWh; *right axis*) in the depicted period.

or ‘multivariate’ loss differential series (Maciejowska et al., 2023; Narajewski and Ziel, 2020):

$$\Delta_d^{A,B} = \|\varepsilon_d^A\|_r - \|\varepsilon_d^B\|_r, \quad (10)$$

is zero, where $\|\varepsilon_d^Z\|_r = (\sum_{h=1}^{24} |\varepsilon_{d,h}^Z|^r)^{1/r}$ is the r -th norm of the 24-dimensional vector ε_d^Z of out-of-sample errors for model Z . In our study, we use the $\|\cdot\|_1$ norm, i.e., set $r = 1$. Naturally, we assume that the loss differential series is covariance stationary.

In Figure 4 we present the results using four chessboards, separately for the two datasets (EPEX, OMIE) and the two model classes (ARX, LEAR). Like in Lago et al. (2021) and Ziel and Weron (2018), each chessboard uses a heat map to indicate the range of the p -values: green and yellow for $p < 0.05$, red for $0.05 \leq p < 0.10$, and black for $0.10 \leq p$.

For both markets, the eSCARX-based models significantly outperform the ARX, and similarly, the eSCLEAR-based models outperform the LEAR benchmark. On the other hand, only in the case of the OMIE market do the SCARX- and SCLEAR-based models provide more accurate forecasts than the respective benchmark models, but at a lower level of significance.

Furthermore, all models with extrapolated LTSC significantly outperform the corresponding naive LTSC-based models and for all averaging schemes, i.e., -MA, -S, and -MAS. In the majority of cases, the -MAS approach outperforms both the -MA and -S averaging schemes, which is not surprising since -MAS models use a larger and more diverse pool of individual forecasts. Overall, the eSCLEAR-MAS model is significantly better than any of the other models, with the exception of eSCLEAR-S for the OMIE market, but even in that case it is not significantly worse.

4.1.3. Temporal performance

In Figures 5 and 6 we plot the rolling 365-day *relative mean absolute errors* (rMAE) for the best performing models, i.e., based on the richest pool of individual forecasts labeled -MAS, with

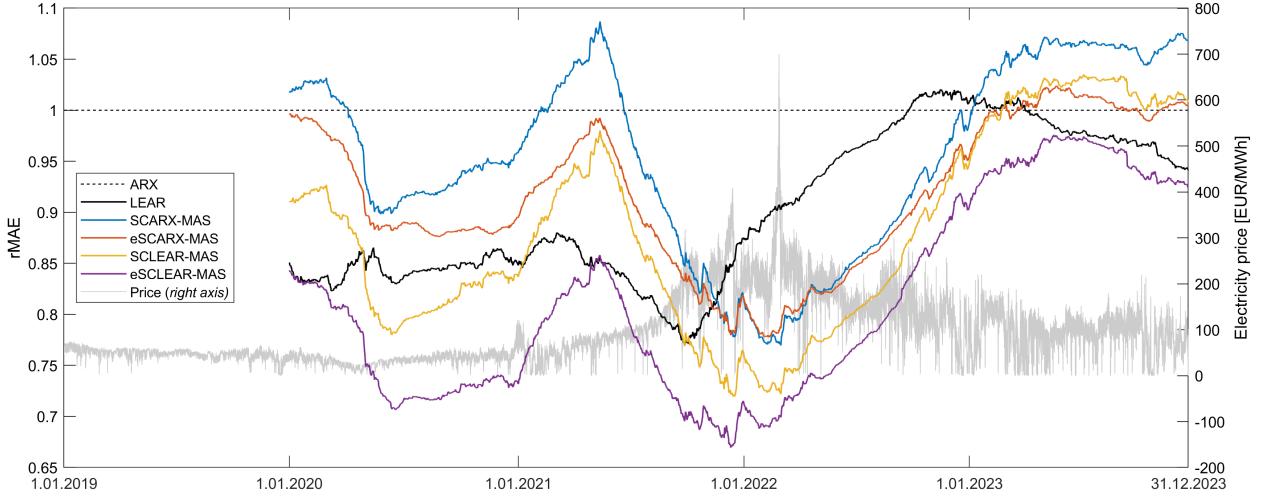


Figure 6: Rolling 365-day *relative mean absolute errors* (rMAE; *left axis*) for the Spanish (OMIE) market with respect to MAE of the ARX model, see Eq. (11). Values for 1.01.2020 are the relative MAEs for 1.01.2019–31.12.2019, values for 2.01.2020 are the relative MAEs for 2.01.2019–1.01.2020, etc. The gray curve is the mean daily electricity price (in EUR/MWh; *right axis*) in the depicted period.

respect to the error of the ARX model:

$$\text{rMAE}(\delta) = \frac{\text{MAE}_{\text{model}}(\delta)}{\text{MAE}_{\text{ARX}}(\delta)} = \frac{\frac{1}{24 \cdot 365} \sum_{d=\delta-365}^{\delta-1} \sum_{h=1}^{24} |P_{d,h} - \hat{P}_{\text{model},d,h}|}{\frac{1}{24 \cdot 365} \sum_{d=\delta-365}^{\delta-1} \sum_{h=1}^{24} |P_{d,h} - \hat{P}_{\text{ARX},d,h}|}, \quad (11)$$

where $P_{d,h}$, $\hat{P}_{\text{model},d,h}$ and $\hat{P}_{\text{ARX},d,h}$ respectively denote the actual price, its *model*-derived forecast and its ARX-derived forecast for day d and hour h , and $\delta = 1.01.2020, 2.01.2020, \dots, 31.12.2023$. Note that if *model* = ARX then $\text{rMAE} \equiv 1$, see the dashed horizontal lines. Values of rMAE below 1 indicate better performance than that of the benchmark ARX model, while values above 1 indicate worse performance.

Several important conclusions can be drawn from Figures 5 and 6:

- Throughout the test period, the eSC-type models (red and purple lines) achieve higher accuracy than the corresponding SC-type models (blue and yellow lines). This is true for both ARX-based and LEAR-based models.
- SCARX-MAS (blue line) is the worst performing model for most of the time, being outperformed not only by other SC-type models, but also sometimes by the ARX benchmark (dashed horizontal line).
- The best performing model is the eSCLEAR-MAS (purple line) for both the EPEX and OMIE markets. It is the only one that consistently outperforms the LEAR benchmark, except for a very brief period (365-day windows ending in February and March 2020) in Germany.
- In the case of the EPEX market, the overall ranking of the forecasting models barely changes over time.
- For OMIE market the performance of the LEAR benchmark changes rapidly over time. It outperforms all models except eSCLEAR-MAS until the end of 2021, then it provides one of the less accurate predictions in the year 2022 and regains the second place in 2023.

Table 3: Trading profits for the entire out-of-sample test period expressed as fractions of the profit for the crystal ball strategy, i.e., $\Pi_{CB} = 110\,897.27$ EUR for EPEX and $\Pi_{CB} = 55\,236.55$ EUR for OMIE. Columns labeled “%chng.” show the percentage difference between the profit for an eSC-type model and an SC-type model in the two columns on the left.

LTSC	EPEX						OMIE					
	ARX			LEAR			ARX			LEAR		
None	0.881			0.888			0.830			0.794		
–MAS	0.876	0.886	1.08%	0.880	0.890	1.07%	0.825	0.834	1.04%	0.818	0.829	1.27%

4.2. Economic evaluation

To assess the practical value of reducing price forecasting errors for decision-makers, we now consider a realistic trading strategy in the day-ahead market. More specifically, we assume that we own a 1.25 MWh battery, which for technical reasons cannot be discharged below 0.25 MWh (or 20% of the nominal capacity) and its efficiency of charging as well as discharging is ca. 90% (which corresponds to a ca. 80% efficiency for a single charge-discharge cycle; Sikorski et al., 2019). Originally proposed by Uniejewski and Weron (2021), the strategy involves placing a bid to buy electricity when prices are low and charge the battery, and simultaneously placing a bid to sell electricity when prices are high after discharging the battery; see also Marcjasz et al. (2023) and Nitka and Weron (2023) who used variants of this strategy.

The maximum profit for day d , called the profit of the *crystal ball* strategy, is given by:

$$\Pi_{CB,d} = 0.9P_{d,h_2} - \frac{1}{0.9}P_{d,h_1}, \quad (12)$$

where h_1 and h_2 respectively are the hours with the lowest and the highest price of day d , and $P_{d,h}$ is the actual price for day d and hour h . Clearly, *model*-derived forecasts will yield a profit of

$$\Pi_{model,d} = 0.9P_{d,\widehat{h}2} - \frac{1}{0.9}P_{d,\widehat{h}1} \leq \Pi_{CB,d}, \quad (13)$$

where $\widehat{h}1$ and $\widehat{h}2$ respectively are the hours with the lowest and the highest *model*-predicted prices for day d , regardless of the actual price forecasts $\hat{P}_{d,h}$ for this day. In other words, the *model*-derived forecasts will yield a profit equal to that of the crystal ball strategy only if the *model* correctly predicts the hours with the lowest and highest prices of the day, i.e., if $\widehat{h}1 = h_1$ and $\widehat{h}2 = h_2$.

In Table 3 we report the profits obtained for selected models, expressed as fractions of the profit of the crystal ball strategy for the whole out-of-sample test period, i.e., $\Pi_{CB} = \sum_d \Pi_{CB,d}$, where the summation is over $d = 1.01.2019, 2.01.2019, \dots, 31.12.2023$. In our case, $\Pi_{CB} = 110\,897.27$ EUR for EPEX and $\Pi_{CB} = 55\,236.55$ EUR for OMIE. Like in Section 4.1.3, we only consider models based on the richest pool of individual forecasts (labeled –MAS). The following conclusions can be drawn:

- For both the German and Spanish markets, the eSC-type models yield the highest profits within each group (ARX, LEAR).
- The best performing model for EPEX is the eCLEAR-MAS model, while for OMIE the more parsimonious eSCARX-MAS model.
- The differences in profits are more pronounced for the Spanish than for the German market.
- On the other hand, the differences in profits are much less pronounced than the differences in predictive performance reported in Tables 1 and 2.

The latter observation suggests that while the considered models predict price levels much better than the benchmarks, the benchmarks are able to identify the hours with the lowest and highest prices of the day almost as well as the eSC-type models.

5. Conclusions

We have introduced a novel approach to predicting the long-term seasonal component (LTSC) for the next day, which is a fundamental input to the seasonal component (SC) approach introduced by Nowotarski and Weron (2016). Considering parsimonious autoregressive (ARX) and LASSO-estimated autoregressive (LEAR) models, we have provided evidence that improvements in predictive accuracy from using the proposed extrapolated SC-type (i.e., eSC-type) models compared to using a naive prediction of the LTSC can be as high as 15% for the German EPEX market (in terms of the RMSE; 14% in terms of the MAE) over a 5-year test period covering the Covid-19 pandemic, the 2021/2022 energy crisis, and the war in Ukraine.

Interestingly, the improvements with respect to the classical ARX and LEAR models that do not use seasonal decomposition are higher for the Spanish OMIE market and reach up to 10% (in terms of the MAE; and 8% in terms of the RMSE); all the differences are statistically significant, as measured by the multivariate variant of the Diebold-Mariano test (Ziel and Weron, 2018). Apparently, the models that naively extrapolate the LTSC for the next day do not cope well with the extreme electricity price dynamics in Germany over the period studied and are outperformed by the classical ARX and LEAR benchmarks.

Furthermore, in line with a recent trend in the electricity price forecasting (EPF) literature (Maciejowska et al., 2023), we have considered a realistic trading strategy involving day-ahead bidding and battery storage in order to quantify the benefits in monetary terms. Although for both the German and Spanish markets the eSC-type models yield the highest profits within each group (ARX, LEAR), the differences in profits are much less pronounced than the differences in predictive performance. This may be an indication that while the considered models predict price levels much better than the benchmarks, the benchmarks are able to identify the hours with the lowest and highest prices of the day almost as well as the eSC-type models.

The latter suggests that computing probabilistic forecasts and considering quantile-based trading strategies (Marcjasz et al., 2023; Nitka and Weron, 2023; Uniejewski and Weron, 2021) may lead to more significant performance improvements. Similarly, studying other energy markets where the seasonal component plays an important role, such as natural gas, may provide evidence that the proposed approach extends beyond electricity markets. Investigating the impact of seasonal decomposition on other machine learning models, including deep neural networks, may open

new avenues of research. Finally, although the use of simple averaging was effective in this study, more sophisticated ensembling techniques can be considered. All this, however, is left for future research.

Overall, our results highlight the importance of seasonal decomposition and accurate day-ahead forecasting of the trend-seasonal pattern of electricity prices. The proposed approach is robust and ensures good performance even under extremely volatile market conditions. Due to the use of forecast averaging, it does not require ex-ante selection of the LTSC parameters (width of the moving average, wavelet decomposition level). Its simplicity and low computational requirements make it a perfect tool for daily market operations, both for point forecasting tasks and as a reliable source of inputs for probabilistic forecasting models and risk management applications (Gneiting et al., 2023; Lipiecki et al., 2024).

Acknowledgments

This work was partially supported by the National Science Center (NCN, Poland) through grants No. 2018/30/A/HS4/00444 (to K.C.), No. 2023/49/N/HS4/02741 (to B.U.) and No. 2021/43/I/HS4/02578 (to R.W.).

References

Afanasyev, D.O., Fedorova, E.A., 2019. On the impact of outlier filtering on the electricity price forecasting accuracy. *Applied Energy* 236, 196–210.

Atiya, A.F., 2020. Why does forecast combination work so well? *International Journal of Forecasting* 36, 197–200.

Berrisch, J., Ziel, F., 2024. Multivariate probabilistic CRPS learning with an application to day-ahead electricity prices. *International Journal of Forecasting* 40, 1568–1586.

Billé, A., Gianfreda, A., Del Grosso, F., Ravazzolo, F., 2023. Forecasting electricity prices with expert, linear, and nonlinear models. *International Journal of Forecasting* 39, 570–586.

Caldana, R., Fusai, G., Roncoroni, A., 2017. Electricity forward curves with thin granularity: Theory and empirical evidence in the hourly EPEXspot market. *European Journal of Operational Research* 261, 715–734.

Ciarreta, A., Muniain, P., Zarraga, A., 2022. Do jumps and cojumps matter for electricity price forecasting? evidence from the German-Austrian day-ahead market. *Electric Power Systems Research* 212, 108144.

Diaz, G., Planas, E., 2016. A note on the normalization of Spanish electricity spot prices. *IEEE Transactions on Power Systems* 31, 2499–2500.

Diebold, F.X., Mariano, R.S., 1995. Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13, 253–263.

ENTSO-E, 2024. Entso-e transparency platform. <https://transparency.entsoe.eu>. Accessed: 2024-11-09.

Gaillard, P., Goude, Y., Nedellec, R., 2016. Additive models and robust aggregation for GEFCom2014 probabilistic electric load and electricity price forecasting. *International Journal of Forecasting* 32, 1038–1050.

García-Martos, C., Caro, E., Sánchez, M.J., 2015. Electricity price forecasting accounting for renewable energies: optimal combined forecasts. *Journal of the Operational Research Society* 66, 871–884.

Gianfreda, A., Ravazzolo, F., Rossini, L., 2020. Comparing the forecasting performances of linear models for electricity prices with high RES penetration. *International Journal of Forecasting* 36, 974–986.

Gneiting, T., Lerch, S., Schulz, B., 2023. Probabilistic solar forecasting: Benchmarks, post-processing, verification. *Solar Energy* 252, 72–80.

Grossi, L., Nan, F., 2019. Robust forecasting of electricity prices: Simulations, models and the impact of renewable sources. *Technological Forecasting and Social Change* 141, 305–318.

Hagfors, L., Kamperud, H., Paraschiv, F., Prokopczuk, M., Sator, A., Westgaard, S., 2016. Prediction of extreme price occurrences in the German day-ahead electricity market. *Quantitative Finance* 16, 1929–1948.

Hong, T., Pinson, P., Wang, Y., Weron, R., Yang, D., Zareipour, H., 2020. Energy forecasting: A review and outlook. *IEEE Open Access Journal of Power and Energy* 7, 376–388.

Hubicka, K., Marcjasz, G., Weron, R., 2019. A note on averaging day-ahead electricity price forecasts across calibration windows. *IEEE Transactions on Sustainable Energy* 10, 321–323.

Hyndman, R., Athanasopoulos, G., 2021. *Forecasting: Principles and Practice* (3rd ed.). OTexts, Melbourne. URL: <http://otexts.org/fpp3>.

Janczura, J., 2024. Expectile regression averaging method for probabilistic forecasting of electricity prices. *Computational Statistics* 39. Forthcoming.

Janczura, J., Wójcik, E., 2022. 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, 106015.

Jędrzejewski, A., Marcjasz, G., Weron, R., 2021. Importance of the long-term seasonal component in day-ahead electricity price forecasting revisited: Parameter-rich models estimated via the lasso. *Energies* 14, 3249.

Kath, C., Ziel, F., 2018. The value of forecasts: Quantifying the economic gains of accurate quarter-hourly electricity price forecasts. *Energy Economics* 76, 411–423.

Lago, J., De Ridder, F., De Schutter, B., 2018. Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms. *Applied Energy* 221, 386–405.

Lago, J., Marcjasz, G., De Schutter, B., Weron, R., 2021. Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy* 293, 116983.

Lipiecki, A., Uniejewski, B., Weron, R., 2024. Postprocessing of point predictions for probabilistic forecasting of day-ahead electricity prices: The benefits of using isotonic distributional regression. *Energy Economics* 139, 107934.

Lisi, F., Nan, F., 2014. Component estimation for electricity prices: Procedures and comparisons. *Energy Economics* 44, 143–159.

Maciejowska, K., Nitka, W., Weron, T., 2021. Enhancing load, wind and solar generation for day-ahead forecasting of electricity prices. *Energy Economics* 99, 105273.

Maciejowska, K., Nowotarski, J., 2016. A hybrid model for GEFCom2014 probabilistic electricity price forecasting. *International Journal of Forecasting* 32, 1051–1056.

Maciejowska, K., Uniejewski, B., Weron, R., 2023. Forecasting electricity prices, in: *Oxford Research Encyclopedia of Economics and Finance*. Oxford University Press, pp. 1–34. doi:10.1093/acrefore/9780190625979.013.667.

Marcjasz, G., Narajewski, M., Weron, R., Ziel, F., 2023. Distributional neural networks for electricity price forecasting. *Energy Economics* 125, 106843.

Marcjasz, G., Serafin, T., Weron, R., 2018. Selection of calibration windows for day-ahead electricity price forecasting. *Energies* 11, 2364.

Marcjasz, G., Uniejewski, B., Weron, R., 2019. On the importance of the long-term seasonal component in day-ahead electricity price forecasting with NARX neural networks. *International Journal of Forecasting* 35, 1520–1532.

Misiorek, A., Trück, S., Weron, R., 2006. Point and interval forecasting of spot electricity prices: Linear vs. non-linear time series models. *Studies in Nonlinear Dynamics & Econometrics* 10, Article 2.

Narajewski, M., Ziel, F., 2020. Econometric modelling and forecasting of intraday electricity prices. *Journal of Commodity Markets* 19, 100107.

Nitka, W., Weron, R., 2023. Combining predictive distributions of electricity prices. Does minimizing the CRPS lead to optimal decisions in day-ahead bidding? *Operations Research and Decisions* 33, 103–116.

Nowotarski, J., Weron, R., 2016. On the importance of the long-term seasonal component in day-ahead electricity price forecasting. *Energy Economics* 57, 228–235.

Percival, D.B., Walden, A.T., 2000. *Wavelet Methods for Time Series Analysis*. Cambridge University Press.

Pesaran, M., Timmermann, A., 2007. Selection of estimation window in the presence of breaks. *Journal of Econometrics* 137, 134–161.

Petropoulos, F., Apiletti, D., Assimakopoulos, V., *et al.*, 2022. Forecasting: Theory and practice. *International Journal of Forecasting* 38, 705–871.

Petropoulos, F., Laporte, G., Aktas, E., *et al.*, 2024. Operational research: Methods and applications. *Journal of the*

Operational Research Society 75, 423–617.

Raviv, E., Bouwman, K.E., van Dijk, D., 2015. Forecasting day-ahead electricity prices: Utilizing hourly prices. *Energy Economics* 50, 227–239.

Shah, I., Akbar, S., Saba, T., Ali, S., Rehman, A., 2021. Short-term forecasting for the electricity spot prices with extreme values treatment. *IEEE Access* 9, 105451–105462.

Shi, W., Wang, Y., Chen, Y., Ma, J., 2021. An effective two-stage electricity price forecasting scheme. *Electric Power Systems Research* 199, 107416.

Sikorski, T., Jasinski, M., Ropuszynska-Surma, E., *et al.*, 2019. A case study on distributed energy resources and energy-storage systems in a virtual power plant concept: Economic aspects. *Energies* 12, 4447.

Taylor, J., 2021. Evaluating quantile-bounded and expectile-bounded interval forecasts. *International Journal of Forecasting* 37, 800–811.

Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society B* 58, 267–288.

Uniejewski, B., 2024. Regularization for electricity price forecasting. *Operations Research and Decisions* 34, 267–286.

Uniejewski, B., Maciejowska, K., 2023. LASSO principal component averaging: A fully automated approach for point forecast pooling. *International Journal of Forecasting* 39, 1839–1852.

Uniejewski, B., Nowotarski, J., Weron, R., 2016. Automated variable selection and shrinkage for day-ahead electricity price forecasting. *Energies* 9, 621.

Uniejewski, B., Weron, R., 2021. Regularized quantile regression averaging for probabilistic electricity price forecasting. *Energy Economics* 95, 105121.

Uniejewski, B., Weron, R., Ziel, F., 2018. Variance stabilizing transformations for electricity spot price forecasting. *IEEE Transactions on Power Systems* 33, 2219–2229.

Wagner, A., Ramentol, E., Schirra, F., Michaeli, H., 2022. Short- and long-term forecasting of electricity prices using embedding of calendar information in neural networks. *Journal of Commodity Markets* 28, 100246.

Weron, R., 2014. Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting* 30, 1030–1081.

Westgaard, S., Fleten, S.E., Negash, A., Botterud, A., Bogaard, K., Verling, T., 2021. Performing price scenario analysis and stress testing using quantile regression: A case study of the Californian electricity market. *Energy* 214, 118796.

Zafar, U., Kellard, N., Vinogradov, D., 2022. Multistage optimization filter for trend-based short-term forecasting. *Journal of Forecasting* 41, 345–360.

Ziel, F., 2016. Forecasting electricity spot prices using LASSO: On capturing the autoregressive intraday structure. *IEEE Transactions on Power Systems* 31, 4977–4987.

Ziel, F., Weron, R., 2018. Day-ahead electricity price forecasting with high-dimensional structures: Univariate vs. multivariate modeling frameworks. *Energy Economics* 70, 396–420.