Energy Economics 70 (2018) 396–420

Contents lists available at ScienceDirect

Energy Economics

journal homepage: www.elsevier.com/locate/eneco

Day-ahead electricity price forecasting with high-dimensional structures: Univariate vs. multivariate modeling frameworks

Florian Ziel[a,*], Rafał Weron[b]

[a] Faculty of Business Administration and Economics, University Duisburg-Essen, Essen, Germany
[b] Department of Operations Research, Wrocław University of Technology, Wrocław, Poland

ARTICLE INFO

Article history:
Received 14 October 2016
Received in revised form 1 December 2017
Accepted 11 December 2017
Available online xxxx

JEL classification:
C14
C22
C51
C53
Q47

Keywords:
Electricity price forecasting
Day-ahead market
Univariate modeling
Multivariate modeling
Forecast combination
Regression
Variable selection
Lasso

ABSTRACT

We conduct an extensive empirical study on short-term electricity price forecasting (EPF) to address the long-standing question if the optimal model structure for EPF is univariate or multivariate. We provide evidence that despite a minor edge in predictive performance overall, the multivariate modeling framework does not uniformly outperform the univariate one across all 12 considered datasets, seasons of the year or hours of the day, and at times is outperformed by the latter. This is an indication that combining advanced structures or the corresponding forecasts from both modeling approaches can bring a further improvement in forecasting accuracy. We show that this indeed can be the case, even for a simple averaging scheme involving only two models. Finally, we also analyze variable selection for the best performing high-dimensional lasso-type models, thus provide guidelines to structuring better performing forecasting model designs.

1. Introduction

There is no consensus in the existing literature on short-term *electricity price forecasting* (EPF) as to the representation of the price series (see Weron, 2014, for a recent review). Should the modeling be implemented in a *multivariate* fashion, i.e., with separate but possibly interdependent models for each of the 24 (48 or more) load periods, or within a *univariate* framework, where one large model is constructed and the same set of parameters is used to produce one- to 24-step ahead predictions for all load periods of the next day?

Surprisingly, though, there are very few and very limited studies in the EPF literature where the univariate and multivariate frameworks are compared. Cuaresma et al. (2004) apply variants of AR(1) and general ARMA processes (including ARMA with jumps) to short-term EPF in the German EEX market. They conclude that specifications in which each hour of the day is modeled separately (i.e., a *multivariate* framework) present uniformly better forecasting properties than *univariate* time series models. More recently, Ziel (2016a) notes that, when we compare the forecasting performance of relatively simple time series models implemented either in a *multivariate* or a *univariate* framework, the latter generally perform better for the first half of the day, whereas the former are better in the second half of the day. However, there has been no through, empirical study to date, involving many fine-tuned specifications from both groups. With this paper we want to fill the gap and provide much needed evidence. In particular we want to address three pertinent questions:

1. Which modeling framework – *multivariate* or *univariate* – is better for EPF?
2. If one of them is better, is it better across all hours, seasons of the year and markets?
3. How many and which past values of the spot price process should be used in EPF models?

* Corresponding author.
E-mail addresses: florian.ziel@uni-due.de (F. Ziel), rafal.weron@pwr.edu.pl (R. Weron).

https://doi.org/10.1016/j.eneco.2017.12.016

The remainder of the paper is structured as follows. In Section 2 we thoroughly discuss the univariate and multivariate modeling frameworks, which are driven by different data-format perspectives. This is a crucial, conceptual part of the paper, which sets ground for the empirical analysis in the following Sections. In Section 3 we briefly describe the 12 price series used and present the *area hyperbolic sine* transform for stabilizing the variance of spot price data. In Section 4 we define 10 forecasting models representing eight model classes: (C1) the mean values of the past prices, (C2) similar-day techniques, (C3) sets of 24 parsimonious, interrelated autoregressive (AR) structures (so-called *expert* models), (C4) sets of 24 univariate AR models, (C5) vector autoregressive (VAR) models, (C6) sets of 24 parameter-rich, interrelated AR models estimated using the *least absolute shrinkage and selection operator* (i.e., lasso or LASSO; which shrinks to zero the coefficients of redundant explanatory variables), (C7) univariate AR models and (C8) univariate, parameter-rich AR models estimated using the lasso. In Section 5 we evaluate their performance on the basis of the Mean Absolute Error (MAE), the mean percentage deviation from the best (m.p.d.f.b.) model and using two variants of the Diebold and Mariano (1995) test for significant differences in the forecasting performance. We also discuss variable selection for the best performing lasso-type models. In Section 6 we wrap up the results and provide guidelines for energy modelers and forecasters. Finally, in the Appendixes we define the full set of 58 forecasting models considered in our empirical study (for clarity of exposition in Section 5 we report detailed results only for 10 representative models), provide formulas for alternative representations of some of the models, and summarize the predictive performance of all 58 models.

2. The univariate and multivariate modeling frameworks

Recall, that the day-ahead price series is a result of conducted once per day (usually around noon) auctions for the 24 h of the next day (Burger et al., 2007; Huisman et al., 2007; Weron and Ziel, forthcoming). Consequently, the electricity prices $P_{d,1}, \ldots, P_{d,24}$ for day d and hours $1, \ldots, 24$ are disclosed at once, and can be regarded as a *multivariate* time series of the 24-dimensional random vector $\boldsymbol{P}_d = [P_{d,1}, \ldots, P_{d,24}]'$. Next to the daily auction argument there are two other practical reasons for the multivariate modeling framework: (i) the demand forecasting literature, which has generally favored the multivariate framework for short-term predictions, and (ii) the fact that each load period (hour, half-hour) displays a rather distinct price profile, reflecting the daily variation of demand, costs, operational constraints and bidding strategies (Gianfreda et al., 2016; Karakatsani and Bunn, 2008; Shahidehpour et al., 2002). On the other hand, the electricity prices can be rewritten as one 'high-frequency' (hourly, half-hourly) *univariate* time series: $P_t = P_{24d+h} = P_{d,h}$, hence are prone to modeling within a univariate framework. The univariate approach is more popular in the engineering EPF literature, dominated by neural network models (see Aggarwal et al., 2009, for a review), but has its roots also in the traditional time series analysis of financial and commodity markets.

Both approaches have their proponents. For instance, Cuaresma et al. (2004), Misiorek et al. (2006), Zhou et al. (2006), Garcia–Martos et al. (2007), Karakatsani and Bunn (2008), Lisi and Nan (2014), Alonso et al. (2016), Gaillard et al. (2016), Hagfors et al. (2016), Maciejowska et al. (2016), Nowotarski and Weron (2016), Uniejewski et al. (2016), and Ziel (2016a), among others, advocate the use of sets of 24 (48 or more) models estimated independently for each load period, typically using Ordinary Least Squares (OLS). In the neural network literature, Amjady and Keynia (2009a), Marcjasz et al. (forthcoming) and Panapakidis and Dagoumas (2016), among others, use a separate network (i.e., a different parameter set) for each hour of the next day.

Studies where univariate statistical time series models are used include Nogales et al. (2002), Contreras et al. (2003), Conejo et al. (2005), Zareipour et al. (2006), Paraschiv et al. (2015) and Ziel et al. (2015a), while papers where neural networks are put to work include Rodriguez and Anders (2004), Amjady (2006), Pao (2007), Amjady et al. (2010), Abedinia et al. (2015), Kim (2015), Dudek (2016), Keles et al. (2016) and Rafiei et al. (2017), among others.

2.1. The multivariate modeling framework

The simplest, yet surprisingly often used structure for the 24-dimensional price time series is a *set of 24 univariate models*:

$$\begin{cases} P_{d,1} = f_1\left(P_{d-1,1}, P_{d-2,1}, \ldots\right) + \varepsilon_{d,1} & \longrightarrow \ \hat{P}_{d,1}, \\ \qquad\vdots & \qquad\vdots \\ P_{d,24} = f_{24}\left(P_{d-1,24}, P_{d-2,24}, \ldots\right) + \varepsilon_{d,24} & \longrightarrow \ \hat{P}_{d,24}, \end{cases} \tag{1}$$

where $\varepsilon_{d,h}$ is the innovation (noise) term for day d and hour h, and $f_h(\cdot)$ are some functions of the explanatory variables of the past prices in the same load period. A commonly raised argument in favor of this approach is that it is simple to implement, involves only a small number of parameters for each load period and hence is computationally non-demanding. The downside, however, is that the estimated set of models does not take into account the potentially important dependencies between the variables across the load periods. Still, by increasing the set of dependent explanatory variables such interrelationships can be added. For instance, Gaillard et al. (2016), Uniejewski et al. (2016) and Ziel (2016a) consider the previous day's price for midnight, i.e., $P_{d-1,24}$, as an explanatory variable in each of the 24 single models. Formally such a *set of 24 interrelated models* can be written as:

$$\begin{cases} P_{d,1} = f_1\left(P_{d-1,1}, P_{d-2,1}, \ldots, P_{d-1,24}, P_{d-2,24}, \ldots\right) + \varepsilon_{d,1} & \longrightarrow \ \hat{P}_{d,1}, \\ \qquad\vdots & \qquad\vdots \\ P_{d,24} = f_{24}\left(P_{d-1,1}, P_{d-2,1}, \ldots, P_{d-1,24}, P_{d-2,24}, \ldots\right) + \varepsilon_{d,24} & \longrightarrow \ \hat{P}_{d,24}, \end{cases} \tag{2}$$

which according to Chatfield (2000) and Diebold (2004) can be regraded as a *multivariate* model, since the dependency structure is interrelated.

It should be emphasized that both frameworks, defined by Eqs. (1) and (2), make explicitly (or implicitly) assumptions on the innovations for individual load periods. For instance, that for each hour $\varepsilon_{d,h}$ follows a normal distribution with zero mean, i.e., $\varepsilon_{d,h} \sim N(0, \sigma_h^2)$. However, they do not assume anything about the joint distribution of the innovations for different hours. To mitigate this unwanted feature, a *fully multivariate* modeling framework may be implemented which treats the price series as panel data:

$$\begin{bmatrix} P_{d,1} \\ \vdots \\ P_{d,24} \end{bmatrix} = \boldsymbol{f}\left(\begin{bmatrix} P_{d-1,1} \\ \vdots \\ P_{d-1,24} \end{bmatrix}, \begin{bmatrix} P_{d-2,1} \\ \vdots \\ P_{d-2,24} \end{bmatrix}, \ldots \right) + \begin{bmatrix} \varepsilon_{d,1} \\ \vdots \\ \varepsilon_{d,24} \end{bmatrix} \quad \longrightarrow \quad \begin{bmatrix} \hat{P}_{d,1} \\ \vdots \\ \hat{P}_{d,24} \end{bmatrix}. \tag{3}$$

The model structure may be identical to that in Eqs. (1)–(2), but it allows for a joint estimation for all load periods, e.g., via multivariate Least Squares, multivariate Yule-Walker equations (as in this paper) or maximum likelihood (Lütkepohl, 2005). Hence, there is an explicit (or implicit) joint distribution assumption on the error