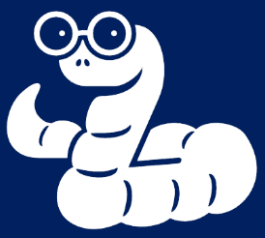

WORKing papers in Management Science

WORMS/20/17

Trading on short-term path forecasts of intraday electricity prices

Tomasz Serafin[1]
Grzegorz Marcjasz[1]
Rafał Weron[1]

[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

Trading on short-term path forecasts of intraday electricity prices

Tomasz Serafin[a], Grzegorz Marcjasz[a], Rafał Weron[a]

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

Abstract

We introduce a profitable trading strategy that can support decision-making in continuous intraday markets for electricity. It utilizes a novel forecasting framework, which generates prediction bands from a pool of path forecasts or approximates them using probabilistic price forecasts. The prediction bands then define a time-dependent price level that, when exceeded, indicates a good trading opportunity. Results for the German intraday market show that, in terms of the energy score, our path forecasts beat a well performing similar-day benchmark by over 25%. Moreover, they provide empirical evidence that the increased computational burden induced by generating realistic price paths is offset by higher trading profits. Still, the proposed approximate method offers a reasonable trade-off – it does not require generating path forecasts and yields only slightly lower profits.

Keywords: Intraday electricity market, Probabilistic forecast, Path forecast, Prediction bands, Energy score, Trading recommendations

1. Introduction

The landscape of European power trading is changing. With a constantly growing share of generation from wind, solar and other *renewable energy sources* (RES), ongoing market integration and active demand-side management, there is a clear tendency towards focusing on shorter time horizons. The workhorse of electricity trading in Europe – the *day-ahead* (DA) market with its uniform price auction conducted a day before delivery [1] – is gradually giving way to *intraday* (ID) trading. The intraday volumes for Germany have doubled between 2014 and 2018, and currently amount to nearly 20% of all electricity traded in the wholesale market [2]. The Cross-Border Intraday Project (XBID), inaugurated in June 2018, has further paved the way for a joint and efficient allocation of intraday capacities and easy access to a single European market [3, 4]. As a result, the ID trading volumes at the EPEX SPOT power exchange for Austria, France and the Netherlands increased in 2019 by ca. 20%, 30% and 55%, respectively [5]. With the bulk of intraday capacities being allocated via continuous trading, e.g., over 85% in Germany, this calls for the development of algorithms and computational tools that can provide support for decision-making in this rapidly developing and very specific market.

*Corresponding author

Email address: rafal.weron@pwr.edu.pl (Rafał Weron)

Preprint submitted to WORMS

Although the continuous intraday markets for electricity share similarities to continuously operated financial and commodity markets, the trading algorithms developed for the latter cannot be applied directly due to the very specific characteristics of intraday electricity markets. For instance, the German market allows participants to continuously trade in parallel 24 so-called *products*, corresponding to the delivery of electricity during each hour of the day; EPEX SPOT also offers 48 half- and 96 quarter-hourly products, but they are less liquid and hence are not studied here [5, 6]. Intraday trading for the next day starts at 15:00, i.e., three hours after the day-ahead auction closes, picks up volume as time passes and continues up to 5 minutes before the delivery. Due to these peculiar arrangements, also the *electricity price forecasting* (EPF) tools [7–12] are not very useful for participants of continuous intraday markets.

To address this gap, we consider a range of forecasting techniques and introduce a profitable trading strategy that can be readily applied in continuous intraday markets for electricity. Our contribution is twofold. Firstly, we propose a 3-step procedure (abbreviated '3S') for computing *path* – also called *trajectory* or *ensemble* – price forecasts of hourly products, which comprises:

1. computing point forecasts of the intraday price for different time horizons using a parameter-rich regression model estimated via the *least absolute shrinkage and selection operator* (LASSO), as in [13, 14],
2. applying quantile regression to obtain 99 percentiles approximating the predictive distributions at different time horizons, similarly to [15, 16],
3. using a Gaussian copula for capturing temporal dependencies in the generated trajectories, like in [17–19].

We assess the predictive accuracy of the obtained path forecasts in terms of the *energy score*, a strictly proper scoring rule for multivariate distributions [20], and show that our 3-step approach outperforms the similar-day (SD) method proposed in [4] by over 25%.

Secondly, we introduce a novel application of the so-called (*simultaneous*) *prediction bands* that encapsulate a path forecast's joint predictive distribution [21]. We either construct the bands directly from the path forecasts generated in step 3 or by adjusting quantile lines, i.e., lines that link the same quantiles of the predictive distributions computed in step 2, so that they yield the required simultaneous coverage probability; hence the name – *adjusted quantile lines* (AQL) approach. In either case, they allow us to estimate a time-dependent price level that, when exceeded, indicates a good trading opportunity. In this way, we provide a readily available methodology to support decision-making when trading in the continuous ID market for electricity. Our results show that the increased computational burden induced by generating realistic price paths is offset by higher trading profits compared to the SD and AQL approaches.

The remainder of this paper is structured as follows. In Section 2 we describe the German intraday market for electricity and present the data. In Section 3 we discuss the forecasting framework, including the LASSO-estimated regression model, the quantile regression-based approach to constructing probabilistic forecasts from point predictions and the methods used to compute path forecasts and prediction bands. In Section 4 we present the trading strategies and discuss ex-ante selection of the simultaneous coverage probability. Then, in Section 5, we evaluate the predictive performance in terms of the energy score and trading profits. Finally, in Section 6 we wrap up our findings and conclude.

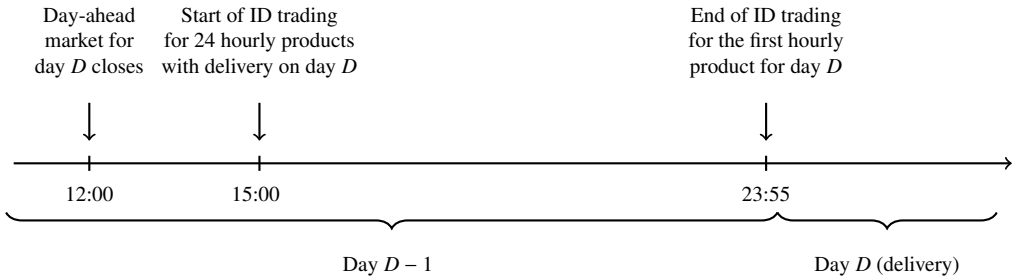

Figure 1: Timeline of events in the German ID market for hourly products. Intraday trading for all hours of day D starts at 15:00 on day $D-1$ and continues until 5 minutes before the delivery of the product. For instance, trading of the first hourly product, i.e., with delivery starting at 0:00, ends at 23:55 on day $D-1$.

2. Preliminaries

2.1. Market description

The German intraday market for electricity is one of the most developed, with trading volumes increasing on a year-to-year basis. Currently, nearly one fifth of all electricity traded in the wholesale market is allocated via ID transactions [2]. Continuously traded hourly products have by far the largest share – 40 TWh (or ca. 75%) in 2019, compared to 0.07 TWh for half-hourly and 6.7 TWh for quarter-hourly products, and 6.9 TWh for the ID auction for quarter-hours [5]. ID trading of hourly products for the next day starts at 15:00, i.e., three hours after the day-ahead auction closes, and continues up to 5 minutes before the delivery of each product, see Figure 1.

While for the DA market the definition of the electricity price is straightforward, the situation for the ID market is not that clear cut. Due to the extreme volatility of prices, taking the last quote as the product's price can be highly deceptive [22]. The most popular point of reference is the so-called *ID3 index*, which is the volume-weighed average price of all transactions that took place between three hours and half an hour before the delivery [13]. Although convenient to use, the ID3 index summarizes the whole market trading activity during last three hours with just a single number. Since over 70% of all transactions take place at this time and their intensity grows exponentially [22], a large share of information may be lost when just looking at the ID3 price.

Hence, in this study we describe the trading activity during the last three hours before the delivery using twelve 15-minute *volume-weighed average* (VWA) prices of all transactions that took place in these time intervals:

$$X_{d,h,t_j} = \frac{\sum_i x^i_{d,h,t_j} v^i_{d,h,t_j}}{\sum_i v^i_{d,h,t_j}}, \tag{1}$$

where x^i_{d,h,t_j} is the price of the i-th transaction for the product with delivery on day d and hour h, which took place in the j-th 15-minute sub-period t_j, v^i_{d,h,t_j} is the volume of the transaction and $j = 1, \ldots, 12$. Using this notation the ID3 index for day d and hour h can be expressed as:

$$\mathrm{ID3}_{d,h} = \frac{\sum_{j=1}^{12} \sum_i x^i_{d,h,t_j} v^i_{d,h,t_j}}{\sum_{j=1}^{12} \sum_i v^i_{d,h,t_j}}. \tag{2}$$

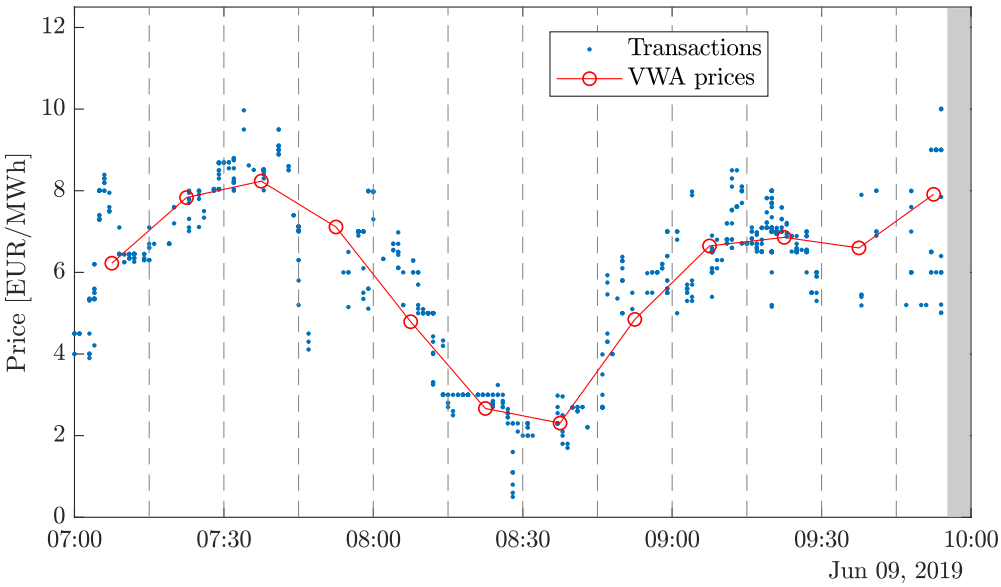

Figure 2: Illustration of the volume weighed average prices (VWA; red circles) constructed from individual transaction prices (volumes are not reported; blue dots) for the hourly product with delivery starting at 10:00 on 9 June 2019. Note, that during the last 5 minutes before the delivery trading activities are not permitted (gray area), hence the last VWA price is calculated based on the 10-minute interval.

The twelve 15-minute VWA prices defined in Eqn. (1) form a 'curve' that can be considered as a trajectory of the ID price for a particular product, as illustrated in Figure 2. In what follows we will try to forecast these curves and make trading recommendations based on these predictions. Using transaction data directly is not advisable, because of the extreme price volatility exhibited by individual transactions. Moreover, when trading significant volumes, one is more likely to divide the whole amount into orders with different prices.

2.2. Data

The transaction data at our disposal includes all trades (prices and volumes) that took place during the trading period for each of the 24 hourly ID products, for each day from 15.06.2017 to 29.09.2019. Apart from ID data, we consider two price-related and four fundamental time series of hourly resolution, each spanning from 15.06.2017 to 29.09.2019. The former two include day-ahead electricity prices and the ID3 index, see Figure 3. The fundamental series include the system-wide load (or consumption), country-wide wind generation and day-ahead forecasts of these two series. In Figure 4 only the forecasts are plotted; the actual values would be indistinguishable at this resolution.

The first dashed line in both Figures marks the end of the initial 364-day calibration window for point forecasts used in step 1 of our 3-step approach (abbreviated '3S'; see Section 3.2.1). This window is followed by three 91-day periods: the initial calibration window for probabilistic forecasts (step 2; see Section 3.2.2), the initial calibration window for path forecasts (step 3; see Section 3.2.3) and the initial calibration window for the ex-ante selection of the simultaneous coverage probability (see Section 4.2.2). The last dashed line also marks the beginning of the

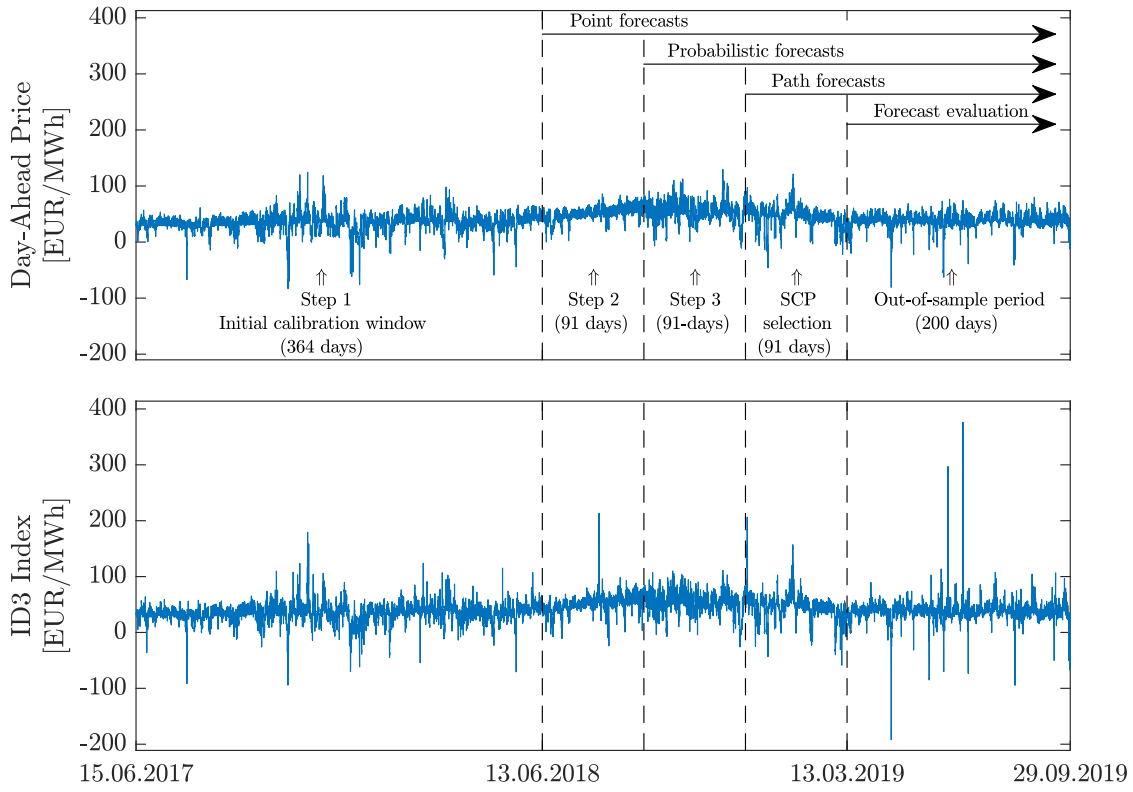

Figure 3: German day-ahead (*top*) and intraday (*bottom*) hourly prices spanning from 15.06.2017 to 29.09.2019. The dashed lines mark the ends of the initial calibration windows in step 1, 2 and 3 of the 3S procedure (see Section 3.2), and the initial calibration window for the ex-ante selection of the simultaneous coverage probability (see Section 4.2.2). The last dashed line also marks the beginning of the 200-day out-of-sample test period (13.03.2019-29.09.2019).

200-day out-of-sample test period, spanning from 13.03.2019 to 29.09.2019. In this study we use the so-called *rolling window framework* – every day the models are reestimated, forecasts for 24 hours of the next day are computed and all calibration windows are moved one day forward. The procedure is repeated until the forecasts for last day in the out-of-sample period are obtained.

3. Forecasting framework

For each hourly product with delivery in the 200-day out-of-sample test period (13.03.2019–29.09.2019), the whole forecasting procedure takes place 4 hours before the delivery to leave enough time for decision-making, see the timeline in Figure 5. The procedure yields path (and probabilistic) forecasts of the VWA prices for the 12 sub-periods $t_1, \ldots, t_{12}$, which are used for computing the simultaneous prediction bands (see Section 3.3) and eventually fed into the trading strategy described in Section 4.2. Two ways of obtaining path forecasts are considered: a similar-day (SD) approach suggested in [4] (see Section 3.1) and a 3-step (3S) approach which comprises: generating point forecasts (step 1; see Section 3.2.1), using them to compute quantile forecasts for 99 percentiles (step 2; see Section 3.2.2) and constructing path forecasts from the latter via a Gaussian copula (step 3; Section 3.2.3). The simultaneous prediction bands, that encapsulate the

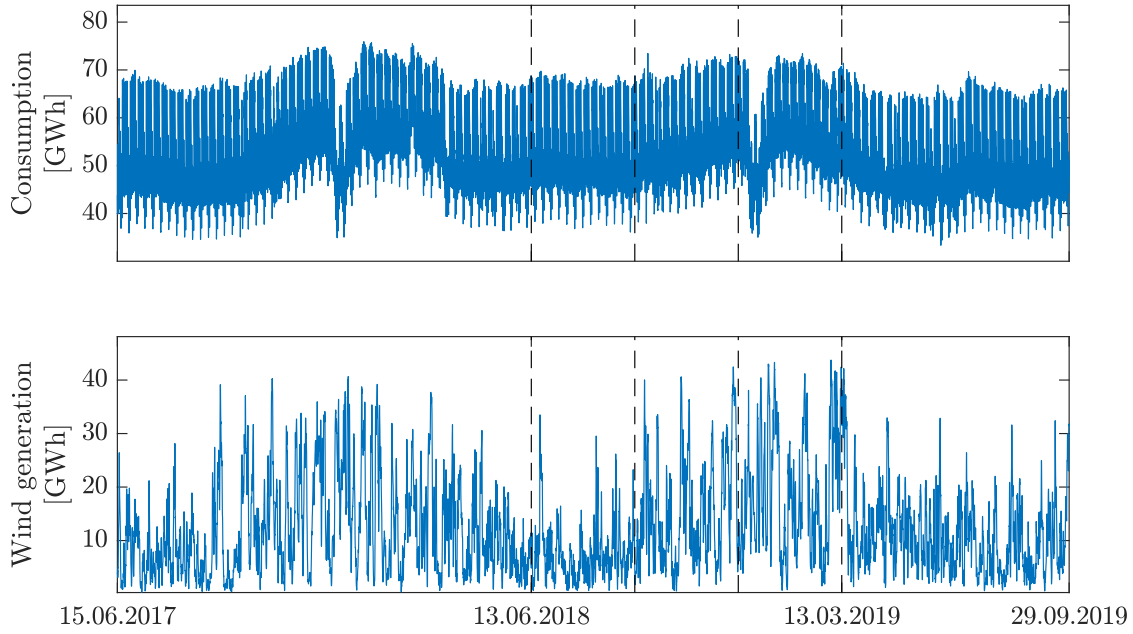

Figure 4: Day-ahead forecasts of the system-wide load (or consumption; *top*) and wind generation (*bottom*) at hourly resolution, spanning from 15.06.2017 to 29.09.2019. For explanation of the dashed lines see Figure 3.

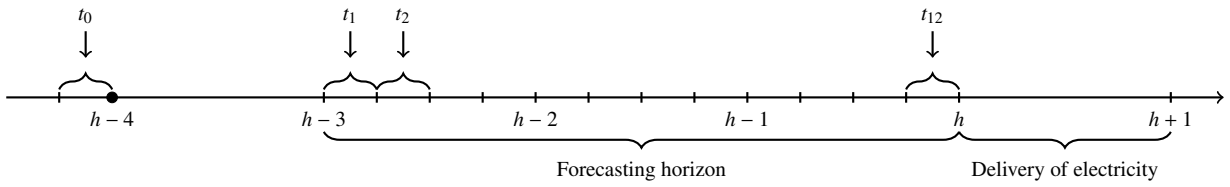

Figure 5: Timeline of the forecasting framework. Predictions for all twelve 15-minute sub-periods $t_1, \ldots, t_{12}$ spanning the last three hours of trading are computed 4 hours before the delivery starts (denoted by •). At this time the last known VWA price is X_{d,h,t_0} or the volume weighted average (VWA) of transactions in sub-period t_0, i.e., between 4:15 and 4 hours before the delivery.

joint predictive distribution for the 12 sub-periods $t_1, \ldots, t_{12}$, are either directly computed from the SD/3S-generated path forecasts or approximated based on the probabilistic predictions from step 2 (a method we dub *adjusted quantile lines*, AQL; see Section 3.3.2). For clarity, the whole forecasting framework is illustrated in Figure 6.

3.1. Similar-day path forecasts

A relatively well-performing, similar-day technique of generating trajectories was suggested by Narajewski and Ziel [4]. We slightly modify it and use it as a benchmark. The technique assumes that the price dynamics, i.e., the changes in the VWA prices over the consecutive trading periods, follow exactly the same dynamics as a randomly chosen historical price path. Formally, the method can be written as $\Delta X_{d,h,t_j} = \Delta X_{d^*,h,t_j}$, where $\Delta X_{d,h,t_j}$ denotes the difference between the VWA prices in the t_j-th and the preceding 15-minute sub-periods, and d^* is a randomly selected past day. It is important to note two things: d^* is selected once for all 15-minute sub-periods t_j,

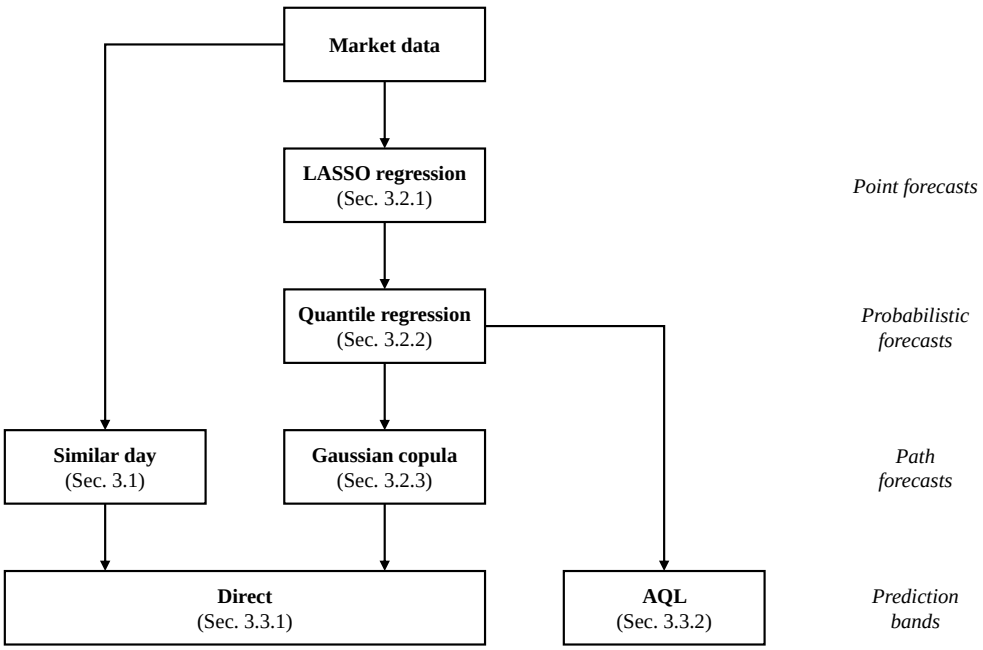

Figure 6: Flowchart of the forecasting framework described in Section 3.

i.e., the whole path is considered, and h is fixed, i.e., the sample path reflects the price evolution of the same hourly product.

To obtain the VWA price path from the increments, a starting point has to be set. Unlike in [4], we do not make the unrealistic assumption that the price 3 hours before the delivery is known (when making the forecast 4 hours before the delivery). Instead, we rely on predicted values. Hence, our similar-day path forecasts are obtained by independently sampling price increment series and choosing the starting point as a randomly selected quantile from the predictive distribution of the VWA price in the first sub-period t_1. Such an approach additionally allows us to obtain many non-duplicate trajectories, which is important for generating prediction bands (see Section 3.3).

3.2. 3-step approach to computing path forecasts

3.2.1. Generating point forecasts

The baseline regression model for the VWA price in the j-th 15-minute period t_j before the delivery on day d and hour h is given by:

$$X_{d,h,t_j} = \beta_0 + \underbrace{\sum_{i=4}^{24} \beta_{i-3} \mathrm{ID3}_{d,h-i} + \sum_{i=0}^{24} \beta_{22+i} \mathrm{DA}_{d,h-i}}_{\text{past ID3 and past/forward-looking DA prices}} + \underbrace{\sum_{i=0}^{24} \beta_{47+i} \widehat{W}_{d,h-i} + \beta_{72} W_{d,h-4} + \beta_{73} W_{d,h-24}}_{\text{wind generation forecasts and past values}}$$
$$+ \underbrace{\sum_{i=0}^{24} \beta_{74+i} \widehat{L}_{d,h-i} + \beta_{99} L_{d,h-4} + \beta_{100} L_{d,h-24}}_{\text{load forecasts and past values}} + \underbrace{\beta_{101} X_{d,h,t_0}}_{\text{last VWA price}} + \varepsilon_{d,h,t_j}, \tag{3}$$

where $\mathrm{ID3}_{d,h}$ is the value of the ID3 price index for day d and hour h, $\mathrm{DA}_{d,h}$ is the day-ahead price for day d and hour h, $\widehat{W}_{d,h}$ and $W_{d,h}$ are the day-ahead predicted and actual wind generation for day d and hour h, respectively, $\widehat{L}_{d,h}$ and $L_{d,h}$ are the day-ahead predicted and actual system-wide load for day d and hour h, respectively, and X_{d,h,t_0} is the last known VWA price, i.e., the VWA price of all transactions between 4 hours and 15 minutes and 4 hours before the delivery. As Narajewski and Ziel [22] and Marcjasz et al. [14] argue, including the latter regressor significantly improves the forecasts. Note, that when writing $\mathrm{ID3}_{d,h-i}$, we refer to the price i hours before day d and hour h, not to the price on day d and hour $h - i$, since the latter index may be negative. Note also, that for each day d and hour h we estimate the β_i's in Eqn. (3) independently for each t_j, i.e., we do not generate 12-step ahead forecasts but 12 times compute one-step ahead predictions, each time for a different horizon t_j, see Figure 5.

As in Marcjasz et al. [23], all variables in Eqn. (3) are preprocessed before calibrating the model. Specifically, each input variable is normalized by subtracting the in-sample median and dividing by the in-sample median absolute deviation adjusted by the 75-th percentile of the standard normal distribution; normalization is performed independently for each of the input variables, as well as for the dependent variable. Next, the *area hyperbolic sine* is applied as the so-called variance stabilizing transformation [24]. However, unlike in the cited studies, the mathematically correct variant of the inverse transformation is used [22].

As in many recent electricity price forecasting studies [4, 13, 14, 25, 26], the regression model is estimated using the *least absolute shrinkage and selection operator* (LASSO) of Tibshirani [27]. LASSO implicitly performs feature selection by penalizing large coefficients and effectively setting some of the β_i's in Eqn. (3) to zero. In our study, the regularization parameter is chosen via cross validation using 3 folds and 50 automatically chosen values of the parameter, as implemented in the *scikit-learn* library for Python [28]. It is also worth mentioning, that although for each of the 12 sub-periods t_j we start from the same set of independent variables, LASSO may eliminate some of them and ultimately the models for different t_j's may use different regressors.

3.2.2. Computing quantile forecasts

Once the point forecasts are generated, we approximate the predictive distributions $\hat{F}_{d,h,t_j}$ at the 12 sub-periods $t_1, \ldots, t_{12}$ using quantile regression. The latter provides an effective and a commonly used in energy forecasting way of converting point forecasts to probabilistic ones [11, 29].

In what follows, we consider 99 percentiles, i.e., quantiles $q^{(\alpha)}_{d,h,t_j}$ with $\alpha = 0.01, 0.02, ..., 0.99$. This allows to approximate the entire predictive distribution relatively well. Due to numerical inefficiencies, however, the neighboring percentiles may be overlapping leading to so-called quantile crossing [30]. Hence, following [15, 31], the 99 quantile estimates are sorted to obtain monotonic quantile curves.

The idea of quantile regression-based methods is that the α-th quantile of the predicted variable (here: the VWA price X_{d,h,t_j}) can be represented as a linear combination of predictor variables (here: an intercept and the point prediction from the LASSO-regression in Section 3.2.1):

$$\hat{q}^{(\alpha)}_{d,h,t_j} = [1 \quad \hat{X}_{d,h,t_j}] \cdot \boldsymbol{w}_\alpha, \tag{4}$$

where $\boldsymbol{w}_\alpha$ is a vector of weights for quantile α, estimated by minimizing the so-called *pinball score* for each percentile [4, 11, 16, 20]. The whole process has to be repeated for each percentile and for each of the sub-periods t_j, yielding a total of 99 percentile forecasts for each of the 12 VWA prices. Note, that because the dataset at our disposal spans only $2\frac{1}{4}$ years, the window for calibrating quantile regression is much shorter (91 days, see Figure 3) than for the LASSO-estimated regression in Section 3.2.1. This, however, should not be a problem [31].

3.2.3. Capturing temporal dependencies

Following [17–19], to construct VWA price trajectories that capture temporal dependencies, we assume that after a suitable transformation, the prediction errors follow a multivariate Gaussian distribution. More complex dependence structures can be modeled either by other copula functions or by sampling from a multivariate empirical *cumulative distribution function* (CDF). The latter corresponds to the multivariate bootstrap where we draw the full residual vector at once, similarly as in the similar-day approach in Section 3.1.

If the probabilistic forecasts derived in step 2 (see Section 3.2.2) are reliable – and we assume they are – the observed proportions for each of the quantiles correspond to the nominal ones [32]. Hence, the random variable Z_{t_j} whose realization $Z^{d,h}_{t_j}$ on day d and hour h is given by:

$$Z^{d,h}_{t_j} = \Phi^{-1}\left(\hat{F}_{d,h,t_j}(X_{d,h,t_j})\right), \tag{5}$$

where $\Phi^{-1}(\cdot)$ is the inverse normal CDF, is normally distributed. Next, we assume that the random vector $\mathbf{Z} = (Z_{t_1}, \ldots, Z_{t_{12}})$ follows a multivariate normal distribution $\mathcal{N}(\mathbf{0}, \Sigma)$ with covariance matrix Σ. To estimate Σ, which contains information about the dependencies between prediction errors for all sub-periods, we use a (third; see Figures 3 and 4) 91-day calibration window. In our study, the covariance matrix is – unlike for the point and probabilistic predictions – reestimated on a daily (not hourly) basis; this allows to get reliable estimates from a relatively short calibration window. Moreover, for the above procedure to work, the predictive distribution $\hat{F}_{d,h,t_j}$ has to be a continuous function. We achieve this by linearly interpolating between the 99 quantile forecasts obtained from Eqn. (4). We also add two extreme points – the maximum and the minimum past price – to serve as quantiles of order $\alpha = 0$ and 1.

To generate M path forecasts for hour h and day d, each spanning all 12 sub-periods $t_1, \ldots, t_{12}$, we first randomly draw M realizations of the multivariate normal random variable $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \hat{\Sigma})$,

where $\hat{\Sigma}$ is the estimated covariance matrix. Next, for each $i = 1, \ldots, M$, we take the i-th realization $\mathbf{Z}^{d,h,i} = (Z^{d,h,i}_{t_1}, \ldots, Z^{d,h,i}_{t_{12}})$ and apply the inverse transformation:

$$\widetilde{X}^i_{d,h,t_j} = \hat{F}^{-1}_{d,h,t_j}\left(\Phi(Z^{d,h,i}_{t_j})\right) \tag{6}$$

to obtain the i-th price path forecast $\widetilde{X}^i_{d,h} = \left(\widetilde{X}^i_{d,h,t_1}, \ldots, \widetilde{X}^i_{d,h,t_{12}}\right)$. In the empirical study in Section 5 we use $M = 10^5$.

3.3. Determining prediction bands

While a prediction interval (or a predictive distribution) reflects VWA price uncertainty at a single point in time, a (*simultaneous*) *prediction band* accounts for the temporal dynamics of the whole price path $X_{d,h,t_1}, \ldots, X_{d,h,t_{12}}$. More precisely, a prediction band [21]:

$$\left(B^L_{d,h}, B^U_{d,h}\right) \equiv \left\{\left(B^L_{d,h,t_j}, B^U_{d,h,t_j}\right) : j = 1, \ldots, 12\right\}, \tag{7}$$

with coverage probability $1 - \alpha$ satisfies:

$$\mathbb{P}\left(B^L_{d,h,t_j} \leq X_{d,h,t_j} \leq B^U_{d,h,t_j}, \forall_j\right) = 1 - \alpha, \tag{8}$$

where $B^L_{d,h}$ is the lower and $B^U_{d,h}$ the upper edge. In other words, the prediction band is such that $(1 - \alpha)\%$ of the whole VWA price paths lie inside it.

In Section 4 we will construct trading strategies for an energy producer who wants to sell generated electricity. Therefore, only the upper edge of the band will be of interest. Without loss of generality, we can set $B^L_{d,h} \equiv -\infty$ and in what follows refer to $B^U_{d,h}$ satisfying:

$$\mathbb{P}\left(X_{d,h,t_j} \leq B^U_{d,h,t_j}, \forall_j\right) = 1 - \alpha, \tag{9}$$

as the simultaneous prediction band with coverage probability $1 - \alpha$.

3.3.1. Direct approach

In this study, we use two ways of estimating prediction bands. The first approach, dubbed *direct*, utilizes a pool of $M = 10^5$ generated price paths, either using the similar-day (see Section 3.1) or our 3-step procedure (see Section 3.2). When estimating $B^U_{d,h}$ with coverage probability $1-\alpha$, the procedure identifies extreme prices at each time point and discards price paths containing these values, and continues until $\alpha\%$ of the paths are removed. Then, the prediction band is formed by linking the maximum values of the non-removed paths at each time point t_j. Note, that the direct approach is applied to both similar-day (denoted by Direct_{SD}) and 3S-based (denoted by Direct_{3S}) price paths, see Figure 6.

3.3.2. Adjusted quantile line (AQL) approach

The second approach builds upon the concept of the Bonferroni correction, in a sense that quantile forecasts are 'corrected' so that the resulting quantile line has the desired simultaneous coverage probability, i.e., is a prediction band. Hence the name *adjusted quantile line* (AQL) approach.

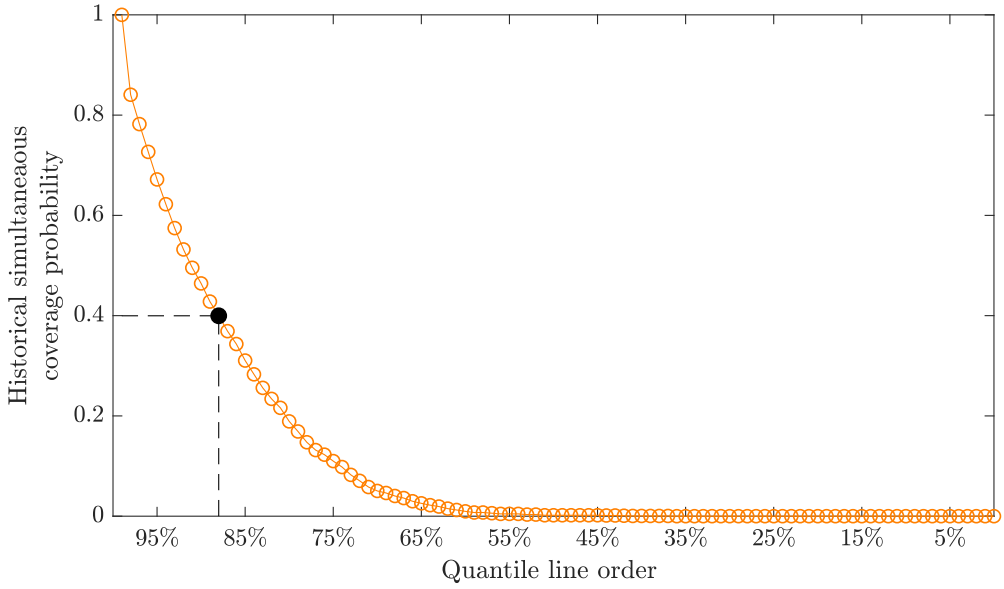

Figure 7: Simultaneous coverage probability (SCP) as a function of the quantile line order. This particular functional dependence is used to generate the AQL prediction bands for the first day in the out-of-sample test window. The dashed lines indicate that a quantile line of order 89% yields SCP=0.4.

Obviously, a quantile line of order $1-\alpha$ does not have the desired simultaneous coverage probability, since it is constructed pointwise without taking into account temporal dependencies. The correction we propose is based on the historical simultaneous coverage of quantile lines of order $1-\alpha$, for $\alpha = 0\%, 1\%, \ldots, 99\%, 100\%$, obtained from quantile forecasts in the 91-day period preceding the moment of forecasting, for an illustration see Figure 7.

First, the average historical coverage of all quantile lines for a given period is calculated, i.e., for each quantile line we check what is the percentage of actual VWA price trajectories that exceeded the line. Then, we compare the obtained results with the desired coverage probability $1-\alpha$. This way, two quantile lines are chosen, one of order $1-\alpha_1$, with historical coverage lower than $1-\alpha$, and one of order $1-\alpha_2$, with historical coverage higher than $1-\alpha$. Linearly interpolating between the two quantile lines we obtain the AQL, which approximately exhibits the desired simultaneous coverage probability of $1-\alpha$. Note, that the historical coverage probability is calculated jointly for all 24 hours of the day. In Figure 8 we illustrate the three types of prediction bands – Direct_{3S}, Direct_{SD} and AQL – for a sample day.

4. Trading strategies

To evaluate the price forecasts in economic terms, we consider a range of trading strategies that reflect a real-life market environment. We take the position of a company that generates electricity from intermittent renewable energy sources (RES) or is a trader that aggregates volumes generated by small RES producers, like in [33]. The exact generated volume is not known in advance, however, the closer is the delivery the more accurate are the predictions. We assume that the company locks-in a large share of generation via day-ahead transactions, and leaves a small

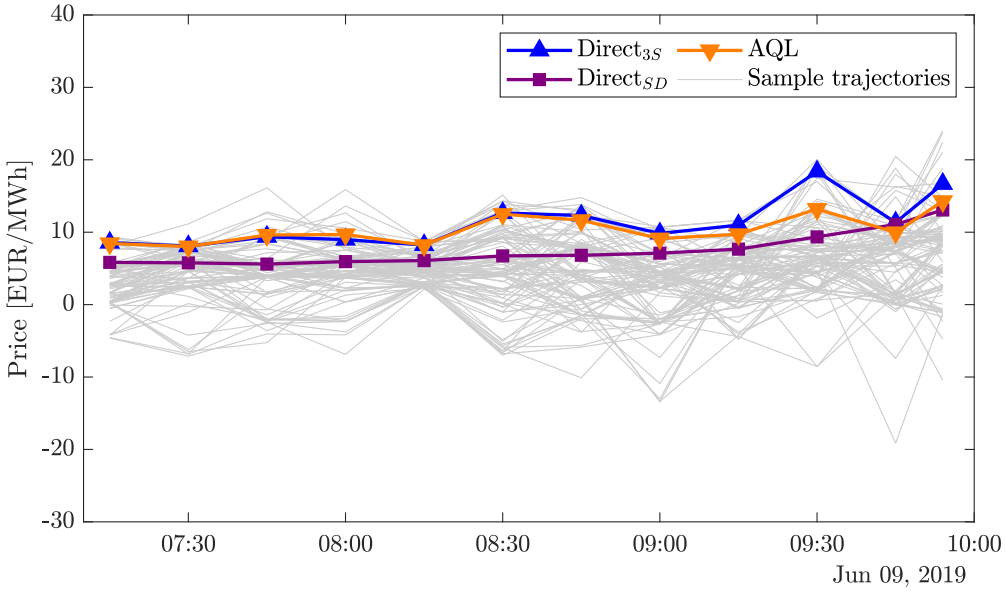

Figure 8: Illustration of the three types of prediction bands – Direct$_{3S}$ (*blue*), AQL (*orange*) and Direct$_{SD}$ (*purple*) – for a sample day (9 June 2019). Gray lines represent generated VWA price paths.

amount to be sold in the German intraday market. For simplicity, 1 MW of electricity for each hour, every single day. Furthermore, we assume that the company is a price taker and its impact on VWA prices (and imbalance volumes) is negligible. Finally, we ignore transaction costs, since they are dependent on individual arrangements and may vary. Given the above, the decision-making process concerns selecting the time(s) and the price(s) at which we place limit orders to sell 1 MW of electricity in the ID market.

4.1. Naive strategies

We use two naive strategies as benchmarks. In the first one, dubbed *Naive1*, 1 MW of electricity is sold for the market price 4 hours before the delivery. In the second, dubbed *Naive2*, the electricity is sold at the last price in the market, i.e., the VWA price closest to the delivery. Note, that both naive strategies do not involve forecasting nor decision-making.

4.2. Prediction band-based strategies

The trading strategy based on simultaneous prediction bands is straightforward and intuitive. Once the actual VWA price path crosses (from below) the prediction band, the electricity is sold in the market. More precisely, prediction bands determine the price of the limit order which is placed in the market every 15 minutes. We assume that if the actual VWA price from a certain 15-minute interval exceeds the level determined by the prediction band, our limit order gets filled. Otherwise the order is adjusted and the price limit is set to the next point on the prediction band: $B^U_{d,h,t_j} \rightarrow B^U_{d,h,t_{j+1}}$. If the VWA price does not cross the prediction band at any time point $t_1, \ldots, t_{12}$ during the trading period, see Figure 5, the electricity is sold at the last VWA price, as in the *Naive2* strategy.

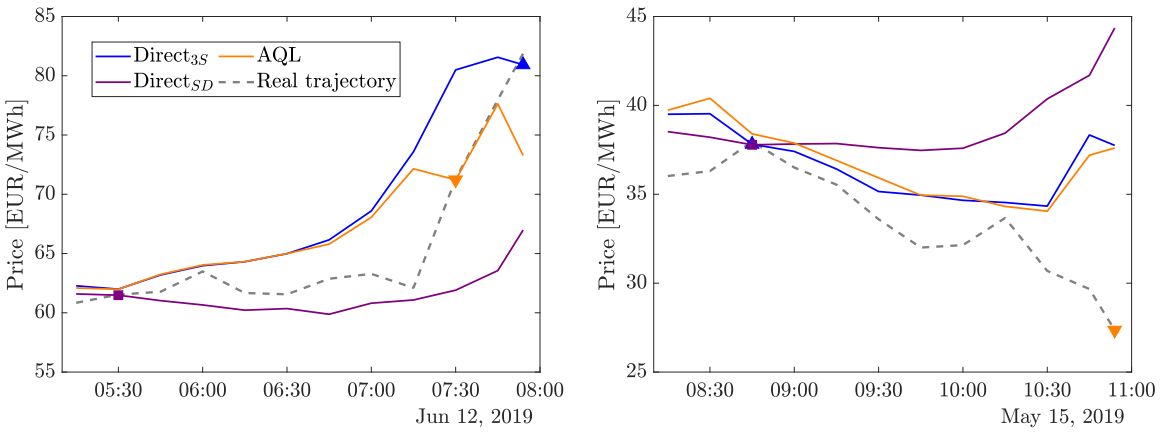

Figure 9: Exemplary trading situations. Color lines represent the three types of prediction bands – Direct$_{3S}$ (*blue*), AQL (*orange*) and Direct$_{SD}$ (*purple*), all with SCP = 30%. Color markers indicate transaction moments and prices.

4.2.1. Fixed simultaneous coverage probability

Obviously, for different levels of the simultaneous coverage probability (SCP) the trading recommendations will be different. In Section 5 we will analyze strategies for different values of the SCP. For the moment being, assume that SCP = 30% is the right choice. Two exemplary trading situations are illustrated in Figure 9. The three types of prediction bands – Direct$_{3S}$, AQL and Direct$_{SD}$ – in general yield different transactions times and prices. In both panels, the Direct$_{S3}$ and AQL-based prediction bands match the behavior of the actual VWA price trajectories much better than the Direct$_{SD}$-based ones. For instance, for 12 June 2019 they forecast a spike in prices towards the end of the trading period and electricity is sold closer to the delivery, respectively for approximately 82 EUR and 71 EUR, compared to only 61 EUR for the similar-day approach. On the other hand, for 15 May 2019 the actual VWA price trajectory decreases over the whole trading period. Here, the prediction bands based on the direct approach are hit early and yield a better transaction price, ca. 37 EUR, while the AQL-based band is never crossed and electricity is sold for around 34 EUR.

4.2.2. Ex-ante selection of the simultaneous coverage probability

The optimal value of the SCP changes in time and across the hours. Therefore, in order for our method to be applicable in a real market setting, the selection of SCP has to be automated. In this study, we use an adaptive approach to select ex-ante the 'optimal' simultaneous coverage probability. An additional, 91-day long rolling calibration window is used to measure the historical profits from each trading strategy for 19 different values of the SCP, i.e., 5%, 10%, ..., 95%, and the best performer is selected for constructing prediction bands for the next day. In Section 5 we will present results for both the ex-ante selection as well as for a range of fixed values of the SCP.

5. Results

In what follows, we evaluate the price forecasts in two different ways. First, in Section 5.1 we assess the predictive accuracy of the obtained path forecasts in terms of the *energy score*, a

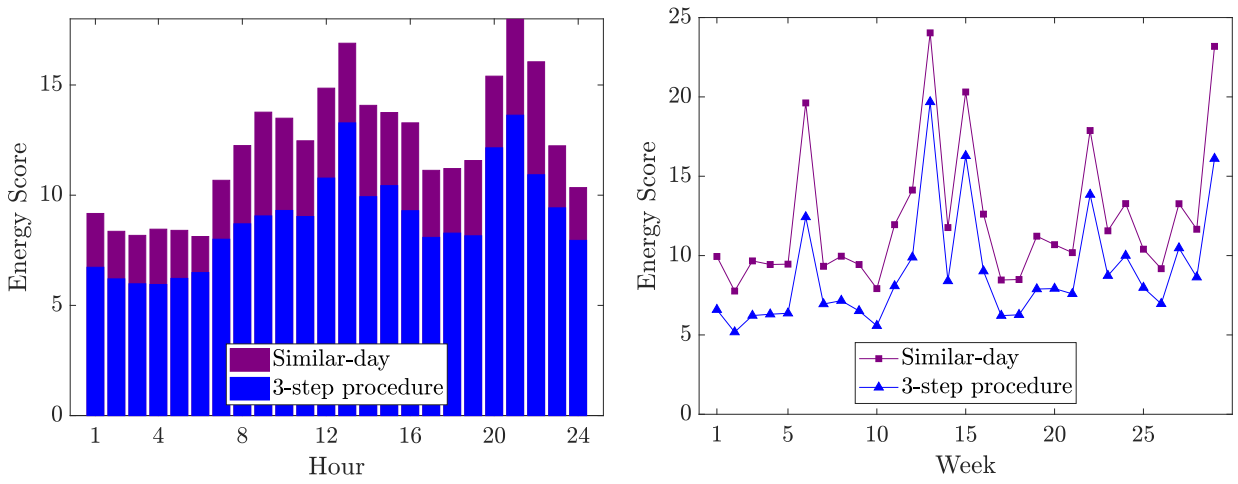

Figure 10: Energy score for the 3-step approach and the similar-day method. The score is averaged – either across all days in the out-of-sample period and plotted for the 24 hours of a day (*left panel*) or across all hours in a week and plotted for the 29 weeks of the out-of-sample period (*right panel*).

strictly proper scoring rule for multivariate distributions. Next, in Section 5.2 we evaluate the path forecasts indirectly by comparing the three prediction band-based trading strategies in terms of generated profits. In both cases the out-of-sample test period spans 200 days, see Figures 3 and 4.

5.1. Energy score

The *energy score* is defined by [20]:

$$ES_{d,h} = \frac{1}{M}\sum_{i=1}^{M}\left\|\widetilde{X}^{i}_{d,h} - X_{d,h}\right\|_2 - \frac{1}{M(M-1)}\sum_{i=1}^{M-1}\sum_{l=i+1}^{M}\left\|\widetilde{X}^{i}_{d,h} - \widetilde{X}^{l}_{d,h}\right\|_2, \tag{10}$$

where $\widetilde{X}^{i}_{d,h} = \left(\widetilde{X}^{i}_{d,h,t_1},\ldots,\widetilde{X}^{i}_{d,h,t_{12}}\right)$ is the i-th path forecast for day d and hour h, $X_{d,h}$ is the corresponding actual VWA price path and $M = 10^5$ is the number of generated paths, see Section 3.2.3 for details. Since Eqn. (10) assesses the accuracy of path forecasts for just one particular day and one hourly product, in Figure 10 the results are averaged – either across all days in the out-of-sample period and plotted for the 24 hours of a day (*left panel*) or across all hours in a week and plotted for the 29 weeks of the out-of-sample period (*right panel*). Looking at Figure 10, we can clearly see that in terms of the energy score our 3-step approach (see Section 3.3) significantly outperforms the similar-day method of Narajewski and Ziel [4]; on average by over 25%.

5.2. Trading profits

We define the trading profit of the energy company as the sum of gains $G_{d,h}$ from selling 1 MW of electricity every hour h and day d in the 200-day out-of-sample test period:

$$\text{profit} = \sum_{d=1}^{N}\sum_{h=1}^{24} G_{d,h}. \tag{11}$$

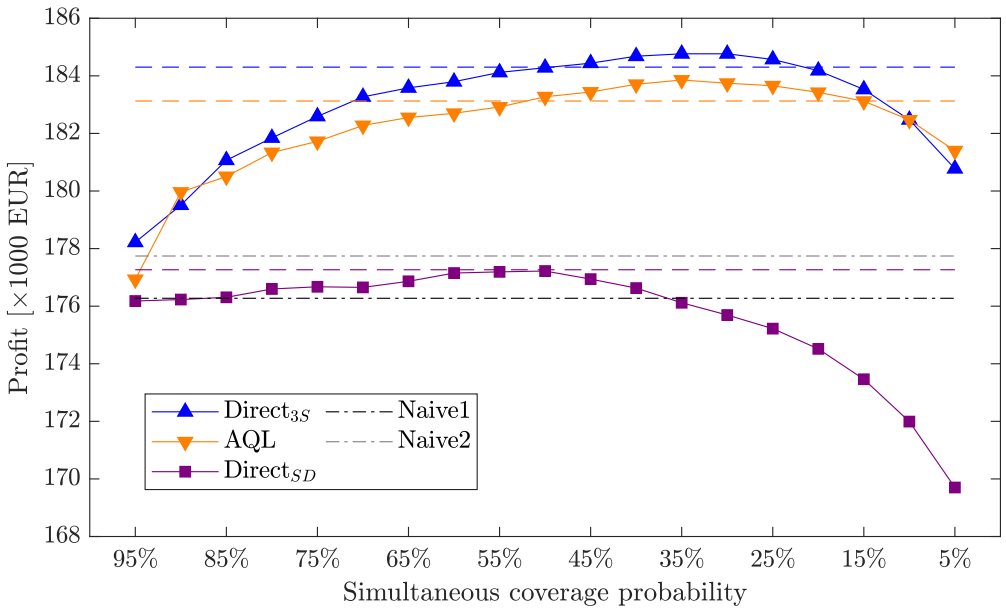

Figure 11: Profits from using the three prediction band-based (Direct$_{3S}$, AQL and Direct$_{SD}$) and two naive (*Naive1*, *Naive2*) trading strategies. In the former case, markers indicate profits for the 19 different values of the simultaneous coverage probability (SCP; see Section 4.2.1) and dashed lines profits for the ex-ante selected SCP (see Section 4.2.2).

The profits from using the three prediction band-based (Direct$_{3S}$, AQL and Direct$_{SD}$) and two naive (*Naive1*, *Naive2*) trading strategies are plotted in Figure 11. In the former case, markers indicate profits for the 19 different values of the simultaneous coverage probability (SCP; see Section 4.2.1) and dashed lines profits for the ex-ante selected SCP (see Section 4.2.2).

We can observe that the proposed strategy utilizing Direct$_{3S}$-based prediction bands outperforms both naive strategies by a considerable margin, except for very high values of the SCP; in the latter case the prediction bands tend to be too 'wide' and the majority of VWA price trajectories do not cross the bands. The AQL-based prediction bands generally yield only slightly lower profits and in two cases even outperform the direct approach. On the other hand, the similar-day approach significantly underperforms in comparison to the Direct$_{3S}$- and AQL-based strategies, and is at par with the naive ones. Its poor performance is likely due to the fact that paths are randomly drawn from past data, and in many cases the method 'predicts' market movements in the opposite direction to the ones observed in reality.

Furthermore, the performance strongly depends on the choice of the SCP – the profits for all strategies have an inverted (and tilted) cup shape. For the Direct$_{3S}$- and AQL-based strategies the maximum is reached at SCP=30-35%, while for the Direct$_{SD}$-based at SCP=50-60%. However, these maximums are only known ex-post. If we follow the ex-ante SCP selection algorithm of Section 4.2.2, then for the former two strategies the performance is suboptimal – roughly by 0.2% for the Direct$_{3S}$-based and by 0.4% for the AQL-based. Interestingly, for the similar-day approach, the ex-ante selection even slightly outperforms the maximum profit for a fixed SCP.

6. Conclusions

In this paper we have addressed an existing literature gap and introduced a profitable trading strategy that can support decision-making in continuous intraday markets for electricity. It utilizes a novel forecasting framework, which generates so-called (simultaneous) prediction bands from a pool of path forecasts or approximates them using probabilistic price forecasts. The prediction bands are then used to issue trading recommendations – they define a time-dependent price level that, when exceeded, indicates a good opportunity to sell electricity.

Using data from the German intraday market, we have shown that – in terms of the energy score – our path forecasts beat a well performing similar-day benchmark [4] by over 25%, both for all hours of the day (when aggregated across all days in the out-of-sample period) and for all weeks in the out-of-sample period (when aggregated across all hours in a week). Moreover, analyzing trading strategies based on three methods of generating prediction bands, we have provided empirical evidence that the increased computational burden induced by generating realistic price paths is offset by higher trading profits. Nevertheless, the proposed approximate method (dubbed AQL) offers a reasonable trade-off – it does not require generating path forecasts and yields only slightly lower profits.

The approach proposed in this paper can be further improved by providing yet more accurate path forecasts. Potentially, this could be achieved by using a more realistic temporal dependence structure, for instance, based on a Student-t copula [34], including heteroscedasticity or non-linearity in the point forecasting model [35, 36], or utilizing intraday fundamental information, e.g., very short-term wind and solar generation forecasts. The proposed methodology could be also (adapted and) tested in other electricity markets.

Acknowledgments

This work was partially supported by the German Research Foundation (DFG, Germany) and the National Science Center (NCN, Poland) through BEETHOVEN grant No. 2016/23/G/HS4/01005 (to T.S.), the Ministry of Science and Higher Education (MNiSW, Poland) through grant No. 0219/DIA/2019/48 (to G.M.) and the National Science Center (NCN, Poland) through grant No. 2018/30/A/HS4/00444 (to R.W.).

References

[1] K. Mayer, S. Trück, Electricity markets around the world, Journal of Commodity Markets 9 (2018) 77–100.
[2] Bundesnetzagentur, Monitoring report 2019, 2019. Available online: https://www.bundesnetzagentur.de.
[3] C. Kath, Modeling intraday markets under the new advances of the cross-border intraday project (XBID): Evidence from the german intraday market, Energies 12 (2019) 4339.
[4] M. Narajewski, F. Ziel, Ensemble forecasting for intraday electricity prices: Simulating trajectories, Applied Energy 279 (2020) 115801.
[5] EPEX, Annual Report 2019, 2020. Available online: http://www.epexspot.com.
[6] R. Kiesel, F. Paraschiv, Econometric analysis of 15-minute intraday electricity prices, Energy Economics 64 (2017) 77–90.
[7] R. Weron, Electricity price forecasting: A review of the state-of-the-art with a look into the future, International Journal of Forecasting 30 (2014) 1030–1081.

[8] D. Keles, J. Scelle, F. Paraschiv, W. Fichtner, Extended forecast methods for day-ahead electricity spot prices applying artificial neural networks, Applied Energy 162 (2016) 218–230.
[9] A. Doostmohammadi, N. Amjady, H. Zareipour, Day-ahead financial loss/gain modeling and prediction for a generation company, IEEE Transactions on Power Systems 32 (2017) 3360–3372.
[10] J. Lago, F. De Ridder, B. De Schutter, Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms, Applied Energy 221 (2018) 386–405.
[11] J. Nowotarski, R. Weron, Recent advances in electricity price forecasting: A review of probabilistic forecasting, Renewable and Sustainable Energy Reviews 81 (2018) 1548–1568.
[12] P. Raña, J. Vilar, G. Aneiros, On the use of functional additive models for electricity demand and price prediction, IEEE Access 6 (2018) 9603–9613.
[13] B. Uniejewski, G. Marcjasz, R. Weron, Understanding intraday electricity markets: Variable selection and very short-term price forecasting using LASSO, International Journal of Forecasting 35 (2019) 1533–1547.
[14] G. Marcjasz, B. Uniejewski, R. Weron, Beating the naïve – combining LASSO with naïve intraday electricity price forecasts, Energies 13 (2020) 1667.
[15] K. Maciejowska, J. Nowotarski, A hybrid model for GEFCom2014 probabilistic electricity price forecasting, International Journal of Forecasting 32 (2016) 1051–1056.
[16] B. Uniejewski, G. Marcjasz, R. Weron, On the importance of the long-term seasonal component in day-ahead electricity price forecasting: Part II – Probabilistic forecasting, Energy Economics 79 (2019) 171–182.
[17] P. Pinson, H. Madsen, H. A. Nielsen, G. Papaefthymiou, B. Klöckl, From probabilistic forecasts to statistical scenarios of short-term wind power production, Wind Energy 12 (2009) 51–62.
[18] S. Chai, Z. Xu, Y. Jia, Conditional density forecast of electricity price based on ensemble ELM and logistic EMOS, IEEE Transactions on Smart Grid 10 (2018) 3031–3043.
[19] T. Janke, F. Steinke, Probabilistic multivariate electricity price forecasting using implicit generative ensemble post-processing, in: Proceedings of the International Conference on Probabilistic Methods Applied to Power Systems – PMAPS 2020, 2020, p. 9183687.
[20] T. Gneiting, A. Raftery, Strictly proper scoring rules, prediction, and estimation, Journal of the American Statistical Association 102 (2007) 359–378.
[21] O. Jorda, M. Marcellino, Path forecast evaluation, Journal of Applied Econometrics 25 (2010) 635–662.
[22] M. Narajewski, F. Ziel, Econometric modelling and forecasting of intraday electricity prices, Journal of Commodity Markets 19 (2020) 100107.
[23] G. Marcjasz, B. Uniejewski, R. Weron, Probabilistic electricity price forecasting with NARX networks: Combine point or probabilistic forecasts?, International Journal of Forecasting 36 (2020) 466–479.
[24] B. Uniejewski, R. Weron, F. Ziel, Variance stabilizing transformations for electricity spot price forecasting, IEEE Transactions on Power Systems 33 (2018) 2219–2229.
[25] F. Ziel, Forecasting electricity spot prices using LASSO: On capturing the autoregressive intraday structure, IEEE Transactions on Power Systems 31 (2016) 4977–4987.
[26] T. Janke, F. Steinke, Forecasting the price distribution of continuous intraday electricity trading, Energies 12 (2019) 4262.
[27] R. Tibshirani, Regression shrinkage and selection via the lasso, Journal of the Royal Statistical Society B 58 (1996) 267–288.
[28] F. Pedregosa, G. Varoquaux, A. Gramfort, et al., Scikit-learn: Machine learning in Python, Journal of Machine Learning Research 12 (2011) 2825–2830.
[29] T. Hong, P. Pinson, Y. Wang, R. Weron, D. Yang, H. Zareipour, Energy forecasting: A review and outlook, IEEE Open Access Journal of Power and Energy 7 (2020) 376–388.
[30] V. Chernozhukov, I. Fernandez-Val, A. Galichon, Quantile and probability curves without crossing, Econometrica 73 (2010) 1093–1125.
[31] T. Serafin, B. Uniejewski, R. Weron, Averaging predictive distributions across calibration windows for day-ahead electricity price forecasting, Energies 12 (2019) 256.
[32] T. Gneiting, F. Balabdaoui, A. Raftery, Probabilistic forecasts, calibration and sharpness, Journal of the Royal Statistical Society B 69 (2007) 243–268.
[33] C. Kath, W. Nitka, T. Serafin, T. Weron, P. Zaleski, R. Weron, Balancing generation from renewable energy

sources: Profitability of an energy trader, Energies 13 (2020) 205.
[34] K. Ignatieva, S. Trück, Modeling spot price dependence in Australian electricity markets with applications to risk management, Computers and Operations Research 66 (2016) 415–433.
[35] A. Ciarreta, P. Muniain, A. Zarraga, Modeling and forecasting realized volatility in German-Austrian continuous intraday electricity prices, Journal of Forecasting 36 (2017) 680–690.
[36] I. Oksuz, U. Ugurlu, Neural network based model comparison for intraday electricity price forecasting, Energies 12 (2019) 4557.