

RESEARCH ARTICLE

WILEY

Distributional modeling and forecasting of natural gas prices

Jonathan Berrisch  | Florian Ziel 

House of Energy Markets and Finance,
University of Duisburg-Essen, Essen,
Germany

Correspondence

Jonathan Berrisch, House of Energy
Markets and Finance, University of
Duisburg-Essen, Universitätsstr. 2, 45141
Essen, Germany.
Email: Jonathan.Berrisch@uni-due.de

Abstract

We examine the problem of modeling and forecasting European day-ahead and month-ahead natural gas prices. For this, we propose two distinct probabilistic models that can be utilized in risk and portfolio management. We use daily pricing data ranging from 2011 to 2020. Extensive descriptive data analysis shows that both time series feature heavy tails and conditional heteroscedasticity and show asymmetric behavior in their differences. We propose state-space time series models under skewed, heavy-tailed distributions to capture all stylized facts of the data. They include the impact of autocorrelation, seasonality, risk premia, temperature, storage levels, the price of European Emission Allowances, and related fuel prices of oil, coal, and electricity. We provide rigorous model diagnostics and interpret all model components in detail. Additionally, we conduct a probabilistic forecasting study with significance tests and compare the predictive performance against literature benchmarks. The proposed day-ahead (month-ahead) model leads to a 13% (9%) reduction in out-of-sample continuous ranked probability score (CRPS) compared with the best performing benchmark model, mainly due to adequate modeling of the volatility and heavy tails.

KEYWORDS

gas prices, heavy-tailed distribution, probabilistic forecasting, risk premium, state-space models

1 | INTRODUCTION

Natural gas is one of the most important energy sources worldwide. In Europe, its relevance increased in recent years because it is one of the cleanest fuels, and thus it can be used as a transition fuel during the transition towards a renewable-based energy system. Increasing prices for carbon emissions accelerate this development (Safari et al., 2019).

Today's natural gas markets are liberalized to a great extent. In consequence, natural gas became a financial

commodity and short-term planning became increasingly important (Hong et al., 2020). That is, accurate price forecasts are attractive for various market participants like natural gas field operators, energy companies, and storage operators. These forecasts are utilized for commodity trading, portfolio and risk management, energy sustainability planning, and predictive maintenance (Petropoulos et al., 2022).

The literature on natural gas price forecasting grew rapidly in the last decade. A great range of different

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2022 The Authors. *Journal of Forecasting* published by John Wiley & Sons Ltd.

methods is applied: fundamental models (esp. Hubbert-Type), econometric models, statistical machine learning models, and hybrids thereof (Elgharbawy & Sayed, 2020; Tamba et al., 2018). Recent studies concerning price forecasting focused on comparing the performance of different methods rather than exploring fundamental relationships between input variables. We want to highlight the major contributions chronologically.

Nguyen and Nabney (2010) and Saghi and Rezaee (2021) combined wavelet transformations with various models such as linear regression, support vector regression, artificial neural networks (ANNs), and generalized autoregressive conditional heteroscedasticity (GARCH) to forecast natural gas prices. The considered linear models outperformed the nonlinear counterparts like the multilayer perceptron models. Salehnia et al. (2013) focused on variable selection and proposed using the gamma test for this. Afterward, they used several forecasting models, including local linear regression and NNs, to predict natural gas prices at different time scales. Their results favor the ANN that did not use any external inputs. The importance of prior variable selection was later revisited by Čeperić et al. (2017). They showed that prior variable selection is crucial for obtaining accurate forecasts using ANNs and support vector regressions. Otherwise, classical time series models outperform these machine learning methods. Geng et al. (2017) explored the relationship between crude oil and natural gas prices using linear and nonlinear Granger causality. They find that the oil price linearly causes the natural gas price. Su et al. (2019) used linear regression, support vector machines, and boosting to predict daily natural gas spot prices. They followed Čeperić et al. (2017) by using heating oil prices, storage capacities, and natural gas imports as explanatory variables. The LSboost algorithm, which minimizes the squared error, outperformed the linear regression and the support vector machine. An ANN outperformed gradient boosting in a previous study. However, that study focused on forecasting monthly natural gas prices, and an even broader set of explanatory variables was used (Su et al., 2019). Herrera et al. (2019) compared a random forest model, an ANN and a hybrid model which uses five more traditional econometric methods for forecasting monthly natural gas prices. The authors did not use any external regressors. The results were clearly in favor of the random forest model. Siddiqui (2019) forecasted daily natural gas spot prices using an autoregressive integrated moving average (ARIMA) model and an autoregressive ANN. The ANN outperformed the ARIMA model, and the model developed in Salehnia et al. (2013). Livieris et al. (2020) used a combination of advanced deep learning methods. In particular, they combined convolutional

layers with long short-term memory (LSTM) layers to predict the day-ahead natural gas prices. They also did not use any external regressors. Jianliang et al. (2020) reviewed three common models, that is, support vector regression, LSTM, and the improved pattern sequence similarity search. Afterward, they developed a weighted model that combines the three methods mentioned above. The results showed that the weighted model surpassed the other considered model. Wang et al. (2021) also uses recurrent neural networks. However, here GRU-based networks are used to forecast short-term natural gas prices.

The determinants of the natural gas price were studied by Wang et al. (2019) using dynamic model averaging. Their results show an evident decline in the impact of oil from 2001 to 2018 while fundamental factors like demand, supply, and the weather became increasingly important. The results also revealed an increasing impact of financial markets. This effect was captured by the impact of variables related to financial speculation. Further, Tang et al. (2019) compared the predictive power of google search data and news sentiment to predict the natural gas spot price. Their results showed that using Google search data can improve forecasts that otherwise solely rely on past prices. Recently, Wang et al. (2020) and Ferrari et al. (2021) build probabilistic forecasting models for several commodities, including natural gas. Similarly, Drachal (2021) used several commodities in a point-forecasting setting. Gao et al. (2021) applied time varying models with stochastic volatility, Markov switching models, and hybrids thereof to predict monthly natural gas prices. The results show that embedding *t*-distributed errors leads to improved forecasting performance, in particular EU markets.

There is also relevant literature on forecasting natural gas demand. Chen et al. (2020) developed a hybrid model to forecast hourly natural gas demand at 96 distribution nodes across Germany. They combined autoregressive (AR) models with convolutional ANNs to reduce overfitting on the one hand but allow nonlinear effects on the other hand. Karabiber and Xydis (2020) also considered ANNs alongside ARIMA models as well as the TBATS model (De Livera et al., 2011). They considered many exogenous regressors such as Fourier terms, temperature, the day type (workdays, weekends, holidays), wind speed, solar radiation, electricity prices, gas prices, and biogas production. The Fourier terms, the temperature, and the lagged consumption were included in all models. The best-performing method combined an ARIMA model with an ANN. This model surpassed the aforementioned TBATS model.

Finally, there is research on forecasting the volatility of natural gas prices. Lv and Shan (2013) used a variety

of linear and nonlinear GARCH type models, including models that capture asymmetric effects. Moreover, they used two different distributional assumptions: normality assumption and skewed-Student- t . The latter consistently outperformed the normality assumption, regardless of the chosen model. This confirms prior research stating that asset returns often feature heavy tails. Hence using the normal distribution is inappropriate in most cases (Engle & Patton, 2007). However, the results did not indicate asymmetric effects. Chkili et al. (2014) also considered asymmetries and long memory effects in the volatility process of natural gas prices and other commodities. They found asymmetric effects in all examined commodities, which contrasts the results of Lv and Shan (2013). Hailemariam and Smyth (2019) investigated how supply and demand shocks affect natural gas price volatility. They found that the effects greatly depend on the market regime and that several structural breaks influenced the volatility response. The global gas market liberalization altered the gas price volatility in particular. Liang et al. (2021) used a GARCH-MIDAS model to estimate the information content of uncertainty indices on natural gas prices volatility.

This paper contributes to the literature by creating probabilistic price forecasting models for natural gas day-ahead and month-ahead products. We propose state-space time series models assuming a heavy-tailed error distribution. The model development builds on a rigorous descriptive data analysis and results from previous research. The proposed models consider many aforementioned and novel effects and significantly outperform popular models presented in the literature. We provide an extensive discussion and validation of each model component and a detailed analysis of the probabilistic predictive performance.

The remainder of this paper is structured as follows. Section 2 presents an overview of the European gas market, the considered price data, and many stylized facts. The proposed models are presented and discussed in Section 3. Section 4 contains the forecasting study design and benchmark models. Section 5 presents extensive model diagnostics and discusses the predictive performance. Section 6 concludes.

2 | MARKET DATA AND STYLIZED FACTS

We use daily prices of future markets for natural gas in Europe. We consider the two most liquid products: the day-ahead and month-ahead product. The choice of the trading hub does not matter much because European natural gas markets are integrated to a great extent;

therefore, prices are highly correlated. Figure 1 presents this correlation for the day-ahead time series. That is, we chose pricing data from the Title Transfer Facility (TTF) Hub because it provides sufficient historical data and is rated the most important trading hub in Europe.¹ A detailed discussion about European natural gas market integration can be found in Hamie et al. (2021). We use volume-weighted average prices for the day-ahead product and closing prices for the month-ahead product. The day-ahead product represents the delivery of one MW natural gas equally each hour from 06:00 (CET) on the next calendar day until 06:00 (CET) of the following day. The month-ahead contract represents the delivery of one MW natural gas equally each hour from 06:00 (CET) on the first day of the month until 06:00 (CET) on the first day of the next month.

Figure 2 presents an overview of both analyzed time series. The data ranges from January 2011 to March 2020. The time series of the day-ahead product (green) features some large upward spikes, which lead to a heavy-tailed return distribution. Those spikes are likely caused by unexpected high demand. This pattern is not present in the month-ahead product because it is not suitable for offsetting these unexpected short-term shocks. We also observe volatility clusters in both processes, indicating that the conditional volatility is not constant. There is neither an apparent trend nor a cycle present in the data.

We fitted a generalized Student- t distribution to the differenced time series to obtain a tail index estimate to examine the heavy tails further. We denote the estimates by τ . They are presented in Table 1 along with other descriptive statistics and the p values of an augmented Dickey-Fuller (ADF) and a Philipps-Perron (PP) unit root test. The tail-index estimate of the day-ahead time series is particularly low, with a value merely larger than 2. This indicates that the second moment hardly exists. Therefore standard least-squares methods likely lead to unreliable results due to the considerable influence of particular observations. This underlines that the normality assumption does not hold.

Figure 3 shows the sample partial autocorrelation function (PACF) of both differenced time series, as well as the PACF of their absolute, positive, and negative parts. The PACF of the differenced day-ahead time series features some significance at lower lags, while the last significant lag of the differenced month-ahead time series is lag 21. The monthly rollover of the product likely causes the latter. This suggests, that the order of the process is 21 if the underlying process would follow an AR process. This pattern is also apparent in the sample PACF of

¹Review of the Gas Hub Assessment, European Federation of Energy Traders (EFET), 2019.




FIGURE 1 Natural gas day-ahead prices at different European trading points and their sample correlation

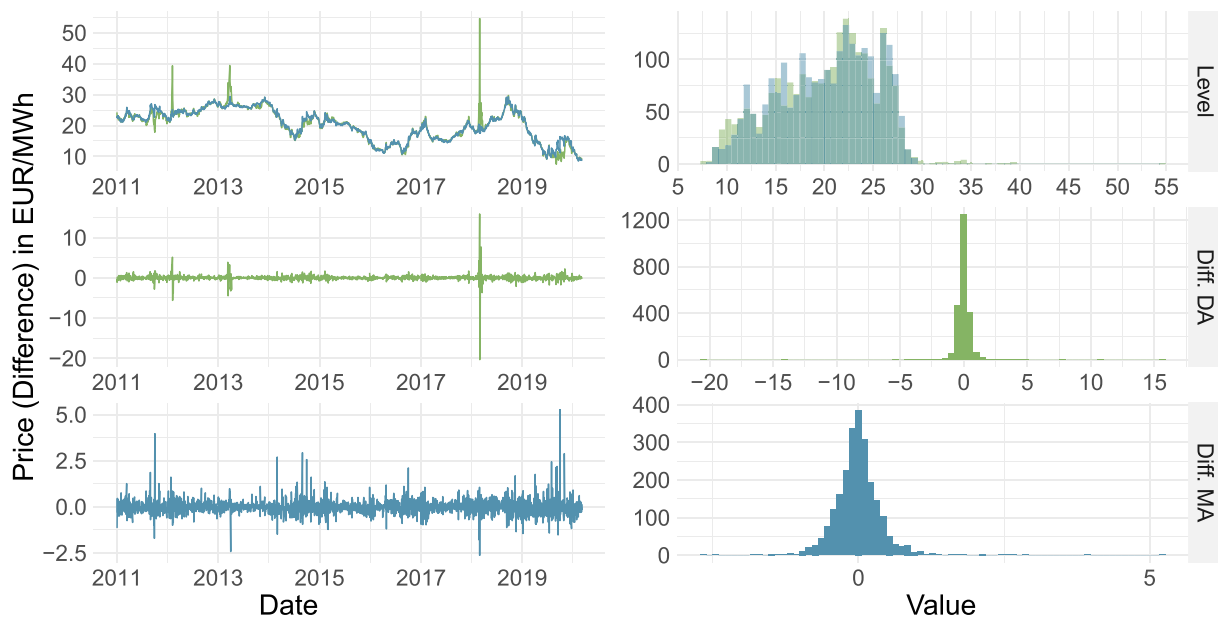


FIGURE 2 Original and differenced day-ahead (green) and month-ahead (blue) time series of natural gas prices (left) with corresponding histograms (right)

TABLE 1 Descriptive statistics of the day-ahead time series, τ denotes the degrees of freedom from a fitted generalized Student- t distribution, PP and ADF denote the p value of the Phillips-Perron and augmented Dickey-Fuller unit root test, respectively

Statistic	Mean	St. dev.	Min	Pctl(25)	Median	Pctl(75)	Max	τ	PP	ADF
Level (DA)	20.15	5.25	7.57	16.25	20.99	24.00	54.52	-	<0.01	0.18
First differences (DA)	-0.01	0.86	-20.31	-0.23	-0.01	0.21	15.77	2.26	<0.01	<0.01
Level (MA)	20.25	4.92	8.53	16.25	21.00	24.25	29.35	-	0.48	0.52
First differences (MA)	-0.01	0.42	-2.62	-0.21	-0.02	0.18	5.25	2.86	<0.01	<0.01




FIGURE 3 Sample PACF of the differenced day-ahead (green) and month-ahead time series (blue) and transformations thereof up to Lag 40 with 5% confidence intervals




FIGURE 4 Mean absolute differences (left) and mean differences of the day-ahead (green) and month-ahead (blue) time series, per weekday (top) and trading day (bottom)

the absolute differences. Thus, the rollover likely influences the mean and the conditional variance of the month-ahead prices. The PACF of the absolute differences of the day-ahead and month-ahead products are also significantly different from zero at lower lags, indicating short-term ARCH or GARCH effects in the conditional variance. Comparing the negative and positive parts of the differences reveals some disparity for both products. That disparity likely indicates leverage effects.

That is, modeling asymmetric effects in the volatility process probably improves the forecasting performance.

Figure 4 depicts the mean absolute difference and mean difference for both products on different weekdays (day-ahead) and days of the month (month-ahead). The day-ahead product features increased absolute differences on Mondays, likely caused by the increased amount of information that emerges over the weekend. The price change from Friday to Monday reflects the uncertainty

concerning the whole weekend, which is generally higher than the uncertainty only concerning one day (Mu, 2007). A suitable volatility model should consider this effect since it will likely improve the probabilistic forecasting performance. The mean absolute difference of the month-ahead product is increased on the first trading day of the month. The latter indicates increased conditional volatility. This is intuitive since the product roll-over causes the underlying delivery period to shift by 1 month. This effect will also be addressed by the proposed month-ahead model presented in Section 3.2.

However, the positive mean difference on the first trading day of the month is noteworthy. One possible explanation for this pattern is the uncertainty concerning the delivery month. While the uncertainty is lowest on the last day of the month (because it is close to delivery), it is highest on the first of every month. Market participants require risk premia as compensation for such risks. The risk-premia demanded by buyers and sellers usually cancel out. However, if the risk-aversion is distributed asymmetrically, the risk premia would differ between both sides of the market. Thus, the risk premia do not cancel out, and the prices will be affected. In particular, if the supply side is more risk-averse, it requires higher risk premia than the seller side, which leads to a price increase. Therefore, the sudden price increase that we observe on the first day of the month is possibly caused by asymmetric risk aversion. This also implies that the sudden price increase will gradually decay as the uncertainty concerning the delivery month diminishes. Both effects will be considered and discussed in Sections 3 and 5. Roncoroni and Brik (2017) discuss risk hedging strategies in natural gas markets.

Figure 5 shows the deviation of the monthly averages from the annual rolling average per year (transparent) and the average monthly deviation from the annual average (bold nontransparent). Both time series feature a seasonal pattern with increased prices in the colder winter months. This pattern is likely caused by the increased natural gas demand in winter, caused by heating. The pattern persists because of storage costs which prevent it from being fully diminished by arbitrage transactions (Chaton et al., 2008; Mirantes et al., 2012).

The natural gas price is likely also influenced by various external factors, as discussed in Section 1. A natural starting point is considering the prices of related energy products like oil, coal, and power. The literature review indicates that oil (Geng et al., 2017) and coal (Li et al., 2017) prices influence the natural gas price (Papież & Śmiech, 2011). This is intuitive since coal, oil, and natural gas can be used for power generation. Therefore, they are substitutable to some degree. Asche et al. (2006) found an influence of power prices on the natural gas prices. However, the direction of those dependencies may change if the underlying relations change, for example, caused by political actions.

Figure 6 presents the external regressors that we consider. In particular, we take into account the month-ahead Rotterdam coal price, the three-, six-, and nine-month-ahead Brent crude oil price, and the German day-ahead and month-ahead power base and power peak prices. Additionally, we consider the daily average Germany temperature, European natural gas storage levels (in percent), and daily spot prices of European Emission Allowances (EUA).




FIGURE 5 Monthly deviation from the annual moving average per year (transparent) and averaged over all years (bold, solid) for the day-ahead (left) and month-ahead (right) time series




FIGURE 6 Time series plot of: month-ahead Rotterdam coal price (Coal), EUA price, X-month-ahead Brent crude oil prices, day-ahead power base and power peak prices of Germany (PWR Base/Peak (DA/MA)), European natural gas storage levels (in percent) and daily averages of the German temperature

		Day-Ahead											
Cor(y_t, x_t)		0.48	-0.33	0.69	0.70	0.71	0.72	0.44	0.45	0.48	0.50	-0.05	-0.18
Cor($y_t, x_t y_{t-1}$)		0.07	-0.04	0.08	0.09	0.09	0.09	0.05	0.06	0.05	0.07	0.02	-0.01
		Month-Ahead											
Cor(y_t, x_t)		0.51	-0.29	0.74	0.75	0.76	0.77	0.45	0.52	0.49	0.58	0.07	-0.11
Cor($y_t, x_t y_{t-1}$)		0.07	-0.03	0.06	0.06	0.06	0.07	0.05	0.05	0.04	0.05	0.03	0.05
		Coal	EU Emission Allowances	Oil (1-Month-Ahead)	Oil (3-Month-Ahead)	Oil (6-Month-Ahead)	Oil (9-Month-Ahead)	PWR Base (DA)	PWR Base (MA)	PWR Peak (DA)	PWR Peak (MA)	Storage Level (Percent)	Temperature

FIGURE 7 Correlation between day-ahead and month-ahead natural gas prices and various related variables, unconditioned (top row) and conditioned on the past observed natural gas price (bottom row)

For EUA the direction in which it influences natural gas prices is unclear. Since using natural gas produces carbon emissions, high EUA prices will render the usage of natural gas unattractive. This effect likely leads to a price reduction. On the other hand, high EUA prices render the power generation using natural gas more attractive compared to the power generation using coal since natural gas produces fewer carbon footprint than coal. This substitution effect leads to increased demand for natural gas due to the fuel switch (Lu et al., 2012; Pratson et al., 2013). Hence, natural gas prices should rise. Which

of both effects mentioned above dominates has to be estimated. The results are discussed in Section 5.

Other factors that might influence natural gas prices through the influence on consumption are temperature (Karabiber & Xydis, 2020) and storage levels (Chaton et al., 2008). Both are directly related to the seasonal demand pattern.

Figure 7 gives more insights into the relation of the considered regressors to the natural gas day-ahead and month-ahead prices. It shows the correlation between the considered external regressors x_t and the natural gas

prices y_t unconditioned and conditioned on the level of the natural gas price at $t - 1$ denoted by y_{t-1} . With the temperature being the only exception, the correlations very are similar for the day-ahead and month-ahead products. Most of the regressors are significantly correlated to the natural gas price. However, most of this correlation diminishes when conditioning on the past level of y . We also observe that the correlation of oil and power prices merely varies across the considered products. Overall, the results suggest a positive relationship between the energy prices (coal, oil, and power) and a slightly negative influence of EUA prices. Further, the results suggest a minor role of storage levels for day-ahead and month-ahead prices and a small positive impact of temperature on month-ahead natural gas prices.

3 | EMPIRICAL MODELS

In this section, we discuss the proposed state-space time series models for the day-ahead and month-ahead product (Hyndman et al., 2008). Both models are developed in a stepwise procedure to incorporate all effects discussed previously. Only effects that improved the probabilistic forecasting accuracy measured by the out-of-sample Continuous Ranked Probability Score (CRPS) were kept. We want to emphasize here that we chose the CRPS as the target to conduct the model selection based on forecasting performance. As day-ahead and month-ahead products feature distinct properties to some extent, we discuss them separately.

3.1 | The day-ahead product model

The proposed day-ahead model is given in Equations (1a) to (1f):

$$\begin{aligned}
 y_t &= \underbrace{\ell_{t-1} + \frac{\psi_1}{5} \sum_{s=t-A-2}^{t-A+2} y_s}_{\text{Seasonal}} + \underbrace{\zeta_1 \text{Coal}_{t-1}}_{\text{Coal}} + \underbrace{\zeta_2 \text{EUA}_{t-1}}_{\text{EUA}} + \underbrace{\zeta_3 \text{PWR}_{t-1}}_{\text{Power}} + d_t, \\
 \ell_t &= \underbrace{\ell_{t-1} + \lambda d_t}_{\text{Exp. Smoothing}}, \\
 d_t &= \underbrace{\varphi d_{t-1} + \theta \varepsilon_{t-1} + \varepsilon_t}_{\text{ARMA Errors}}, \\
 \varepsilon_t &\sim \text{SST}(0, \sigma_t, \nu, \tau), \\
 \sigma_t &= \begin{cases} \delta \tilde{\sigma}_t & \text{Monday} \\ \tilde{\sigma}_t & \text{Else} \end{cases}, \\
 \tilde{\sigma}_t &= \underbrace{\omega + \alpha |\varepsilon_{t-1}| + \beta \sigma_{t-1}}_{\text{Absolute Value GARCH}} + \underbrace{\gamma |\varepsilon_{t-1}| \mathbb{1}(\varepsilon_{t-1} < 0)}_{\text{Leverage}}. \\
 &\quad \text{TGARCH}
 \end{aligned} \tag{1a}$$

The model exploits the autoregressive structure with an exponential smoothing process (1a and 1b) with autoregressive moving average (ARMA) errors denoted by d_t

(De Livera et al., 2011). The price at time t is denoted by y_t and ℓ_t denotes the level component of the state space model. The seasonal structure is implemented using a 1-week moving average, which is lagged by 1 year. One week usually contains five observations since the day-ahead product is only traded on working days. We also considered using seasonal lags or trigonometric Fourier terms to incorporate seasonality. However, the averaged lag proved superior. Moreover, we extend the level equation (1a) of the model by the coal price (coal), the price of EUA, and the day-ahead peak power price (PWR). Furthermore, we considered using a higher order ARMA process to model the error term. However, this did not lead to an improvement in CRPS.

The volatility process was modeled using a threshold GARCH (TGARCH) process (Bollerslev et al., 2010). Compared with the traditional GARCH model, it reduces the influence of large errors by modeling the conditional standard deviation instead of the conditional variance. Thereby, the conditional standard deviation depends on a constant, its own lagged values, the lagged absolute errors, and a leverage effect for negative errors. Refraining from squaring the errors is particularly relevant for natural gas prices due to the heavy-tails, which we discussed in Section 2.

We address the increased volatility on Mondays by leveraging the estimated standard deviation on Mondays. This approach is necessary to avoid increasing the variance on other weekdays as well. The latter would be the case when using a Monday dummy (due to the recursive structure of the TGARCH process). We also considered using a higher order TGARCH process. However, this did not improve the model performance.

We assume the errors to be skewed-student-t distributed, that is, $\varepsilon_t \sim \text{SST}$. The skewed version of the generalized Student-t distribution was developed by Fernández and Steel (1998), who proposed a general approach for introducing skewness in symmetric distributions and applied that approach to the generalized Student-t distribution. This distribution was later reparameterized by Wurtz et al. (2006) such that μ denotes the mean and σ denotes the standard deviation. The density is defined as follows:

$$f_{\mu, \sigma, \nu, \tau}(y) = \begin{cases} \frac{c}{\sigma_0} \left[1 + \frac{\nu^2 y^2}{\tau} \right]^{-(\tau+1)/2} & \text{if } y < \mu_0, \\ \frac{c}{\sigma_0} \left[1 + \frac{y^2}{\nu^2 \tau} \right]^{-(\tau+1)/2} & \text{if } y \geq \mu_0, \end{cases}$$

$$\text{with: } \mu_0 = \mu - \sigma m / s,$$

$$\sigma_0 = \sigma / s,$$

$$c = 2\nu [(1 + \nu^2) B(1/2, \tau/2) \tau^{1/2}]^{-1},$$

$$m = \frac{2\tau^{1/2}(\nu - \nu^{-1})}{(\tau - 1)B(1/2, \tau/2)},$$

$$s = \sqrt{\frac{\tau}{\tau - 2} (\nu^2 + \nu^{-2} - 1) - m^2},$$

(2)

where $y, \mu \in \mathbb{R}, \sigma, \nu > 0$, and $\tau > 2$. Here, ν denote the skewness parameter and τ denotes the degrees of freedom. The distribution is symmetric if $\nu = 1$.

This distribution features heavier tails than the Gaussian distribution and is, therefore, better suited for this data. Moreover, it allows for skewness, which was also discussed in related literature and Section 2. Finally, we discuss the importance of each model component in detail in Section 5.

3.2 | The month-ahead product model

The month-ahead model considers all characteristics that were discovered in Section 2. Preliminary analysis showed that complex autoregressive models do not outperform simple random walk models. Hence, we extended a random walk model to incorporate all the characteristics of the product. The proposed month-ahead model is given in Equations (3a) to (3e).

$$\begin{aligned}
 y_t &= \varphi_0 + \underbrace{\Phi_t}_{\text{Risk Rollover}} + \underbrace{\psi_1 \bar{y}_{t-1Y}^M}_{\text{Seasonal}} + \underbrace{\zeta_1 \text{EUA}_{t-1}}_{\text{EUA}} + \underbrace{\zeta_2 \text{Oil}_{t-1}}_{\text{Oil}} + \underbrace{\zeta_3 \tilde{T}_{t-1}}_{\text{Temperature}} + \varepsilon_t, \\
 \Phi_t &= \begin{cases} \varphi_1 y_{t-1,2MA} & \text{First Trading Day per Month} \\ \underbrace{\varphi_1 y_{t-1}}_{\text{Rollover}} + \underbrace{\eta \frac{D_t^{\text{LTD}}}{D_t^M}}_{\text{Risk Deduction}} & \text{Else} \end{cases}, \\
 \varepsilon_t &\sim \text{SST}(0, \sigma_t, \nu, \tau), \\
 \sigma_t &= \begin{cases} \delta \tilde{\sigma}_t & \text{First Trading Day per Month} \\ \tilde{\sigma}_t & \text{Else} \end{cases}, \\
 \tilde{\sigma}_t &= \underbrace{\omega + \alpha |\varepsilon_{t-1}| + \beta \sigma_{t-1}}_{\text{Absolute Value GARCH}} + \underbrace{\gamma |\varepsilon_{t-1}| \mathbb{1}(\varepsilon_{t-1} < 0)}_{\text{Leverage}}. \\
 &\hspace{10em} \text{TGARCH}
 \end{aligned} \tag{3a}$$

We modeled the product rollover by using the most recent price of the two-month-ahead product $y_{t-1,2MA}$ as a predictor for the month-ahead product on the first trading day of each month. The model considers the most recent price of the month-ahead product y_{t-1} on all other days. Additionally, we added a risk component. The buildup of the risk premia does not need to be modeled because the approach mentioned above to address the model rollover accounts for it already. However, the risk premia gradually decay over the month. We assume this decay to be linear. That is, we weight the decay by the number of days since the last trading day D_t^{LTD} to account for a more significant risk reduction, for example, on Mondays. We also normalize this component using the total number of days in a given month D_t^M . Moreover, we included a seasonal component. We used the monthly

average lagged by one year \bar{y}_{t-1Y}^M . This approach has a clear advantage compared with, for example, Fourier terms: It accounts for the monthly structure of the product, that is, the rollover.

The mean process was further extended by the lagged price of European Emission Allowances, the lagged price of Oil, and the lagged and smoothed temperature denoted by EUA_{t-1} , Oil_{t-1} , and \tilde{T}_{t-1} , respectively. The smoothed temperature is defined as $\tilde{T}_t = 0.95 \tilde{T}_{t-1} + (0.05) T_t$. That is, we follow Gaillard et al. (2016) by smoothing the temperature exponentially with a persistence parameter of 0.95. Lastly, the model was supplemented by a TGARCH process similar to the process of the day-ahead model. However, the conditional standard deviation is elevated on the first trading day of every month instead of every Monday. We assume that the errors are SST distributed as in the day-ahead model.

4 | FORECASTING STUDY

4.1 | Forecasting study design and benchmarks

The validity and forecasting performance of the models was evaluated using one-step-ahead out-of-sample forecasts obtained by an expanding window forecasting study.

That is, the estimation window expands as the forecasting study progresses. We chose this approach to utilize an increasing number of observations for estimation. This approach is more suitable than a rolling window forecasting study due to the heavy-tailed data. We initialized the calibration window with 4 years of observations (exactly 1012/1012 days day-ahead/month-ahead), leaving more than 4 years for validation (exactly 1056/1079 days day-ahead/month-ahead). We estimated the models by minimizing their negative log-likelihood functions. We implemented them in Python along with the SST distribution, which we made available on PyPI (Berrisch, 2021). For minimization, we used sequential-least squares programming, which is provided by Scipy (Virtanen et al., 2020).

Moreover, we set up three different benchmark models from recent literature. Because all three models were only used to do point-forecasting, we had to set a distributional assumption ourselves. We chose to estimate each model with the common normality assumption and the T assumption, which better fits these data, as discussed above.

First, we replicated the time series benchmark model of Siddiqui (2019). They used an ARIMA(2,1,2) model. Because they did not specify how this model was

calibrated, we just chose the same order. Second, we replicated the random forest model of Herrera et al. (2019), which outperformed all other considered models by a great margin. However, we had to adapt the hyperparameters here because Herrera et al. (2019) applied this model to monthly data. That is, we fitted 500 trees like Herrera et al. (2019) and created 270 lag variables that correspond to roughly 1 year of observations, and considered 23 ($\sqrt{270 \times 2} \approx 23$) variables at each split in contrast to 7. Lastly, we fit a vector autoregressive model like Geng et al. (2017) to the differenced time series of natural gas and month-ahead oil prices. We determined the vector autoregressive (VAR) order by using the Akaike information criterion.

We evaluated the benchmarks using the same forecasting study design as described above. We estimated the missing distributional parameters by fitting a Normal and Student-*t*-distribution to the error terms. These parameters were also estimated using maximum likelihood. Thus, we receive $2 \times 3 = 6$ benchmark models in total.

4.2 | Evaluation measures

We use four scoring functions to evaluate the forecasting performance on the aforementioned out-of-sample forecasts. Namely, the mean absolute error (MAE), root mean squared error (RMSE), CRPS, and pinball loss. Table 2 presents these scoring functions formally. While MAE and RMSE are very popular, they do not evaluate the whole probabilistic forecast. The (R)MSE should be used only to evaluate mean forecasts. Similarly, the MAE is appropriate only for evaluating median forecasts (Gneiting, 2011).

Forecasting the whole probabilistic distribution yields a significant advantage: The possibility to obtain several predictors like the mean, the median, or other quantiles of interest without the need to specify the desired

quantity of interest beforehand. This also allows us to use multiple scoring functions appropriately by evaluating the respective quantities of the distribution. However, as discussed above, and RMSE are only suitable for evaluating the mean and median. In contrast, CRPS and pinball loss can be used for evaluating the full predictive distribution. Both are also strictly proper. This means that they reach a unique minimum when the probabilistic forecast equals the actual probability distribution. Hence they will identify the true model if it is considered (Gneiting & Raftery, 2007). The Pinball loss evaluates the performance of individual quantile forecasts. One can calculate the Pinball loss of a probabilistic forecast for arbitrary probabilities of choice using the quantile function. In this paper, the pinball loss is calculated on a probability grid containing all percentiles. The results are then plotted as a line graph, usually creating a bell shape. As with the other loss functions, the model performance is better the smaller the pinball loss is.

The pinball loss is directly related to the MAE and the CRPS. Calculating the pinball loss at the 50% probability yields the MAE/2. Further, we can use the pinball loss to approximate the CRPS:

$$\text{CRPS}(\hat{F}_t(z), y_t) = 2 \int_0^1 \text{PL}(\hat{Q}_t(p), y_t, p) dp.$$

That is, evaluating the pinball loss on a dense grid approximates the CRPS (Berrisch & Ziel, 2021). Both measures have their advantages. The CRPS, on the one hand, summarizes the probabilistic forecasting performance in one single scalar but does not offer insights into the forecasting performance of particular quantile forecasts (Nowotarski & Weron, 2018). The pinball loss, on the other hand, resolves the forecasting performance in more detail. However, presenting pinball losses over a grid with high-resolution requires graphical methods almost surely.

TABLE 2 Scoring functions used for evaluating the forecasting performance of mean ($\hat{\mu}_t$), median ($\tilde{\mu}_t$), *p*-quantile ($\hat{Q}_t(p)$) or distribution (\hat{F}_t) predictions and corresponding observations y_t

Name	Scoring function
Root mean squared error	$\text{RMSE}(y_t, \hat{\mu}_t) = \sqrt{\frac{1}{T} \sum_{t=1}^T (y_t - \hat{\mu}_t)^2}$
Mean absolute error	$\text{MAE}(y_t, \tilde{\mu}_t) = \frac{1}{T} \sum_{t=1}^T y_t - \tilde{\mu}_t $
Pinball loss for probability <i>p</i>	$\text{PL}(\hat{Q}_t(p), y_t, p) = \begin{cases} p(y_t - \hat{Q}_t(p)) & y_t \geq \hat{Q}_t(p) \\ (1-p)(\hat{Q}_t(p) - y_t) & y_t < \hat{Q}_t(p) \end{cases}$
Continuous ranked probability score	$\text{CRPS}(\hat{F}_t(z), y_t) = \int_{\mathbb{R}} (\hat{F}_t(x) - \mathbb{1}(z > y_t))^2 dx$

Additionally, we use the Diebold–Mariano (DM) for model comparison (Diebold & Mariano, 2002). The test statistic is defined as follows:

$$t_{DM} = \frac{\bar{\Delta}_{A,B}}{\sigma(\bar{\Delta}_{A,B})} \text{ with } \bar{\Delta}_{A,B} = \frac{1}{N} \sum_{i=1}^N (L(\varepsilon_{A,i}) - L(\varepsilon_{B,i})), \quad (4)$$

where $\bar{\Delta}_{A,B}$ denotes the loss differential of the two competing models A and B . L denotes the loss function, which depends on the forecasting errors $\varepsilon_{A,i}$ and $\varepsilon_{B,i}$, respectively. The null hypothesis of this test states that the loss differential $\bar{\Delta}_{A,B}$ is not different from zero. The DM is applied using the loss functions presented Table 2.

5 | RESULTS AND DISCUSSION

5.1 | Model diagnostics

We used several measures to assess the model validity before evaluating the predictive performance. First, we visually inspected the standardized mean residuals of both models. These are presented in Figure 8. A comparison to Figure 3 shows that both models exploit the

majority of the autocorrelation in the data. The line graph and the histogram do not reveal any striking patterns either.

We conducted the *energy test of independence* to test for independence of the standardized mean residuals and several selected quantities (Székely et al., 2007; Székely & Rizzo, 2013). We denote the standardized residuals by $z_t = \varepsilon_t / \sigma_t$. This test is based on the distance correlation \mathcal{R} , which only requires finite first moments to exist. Distance correlation measures any correlation between two random variables in arbitrary dimensions. That is, $\mathcal{R} = 0$ only holds if X and Y are independent. This is used to create the *energy test of independence*. The null hypothesis of this test is *independence*: $\mathcal{H}_0 : y_t \perp \perp z_t$. Table 3 presents the test results for the standardized residuals of the day-ahead and month-ahead models. In the day-ahead case, independence is rejected for the dependent variable y_t . This indicates that there are still unexploited dependencies. However, these may be nonlinear and potentially weak, which makes exploiting them difficult. The latter also holds for the coal price, the price of EUA and the seasonal component, although the p values are much higher for these quantities. For the month-ahead model, independence is not rejected for all variables. However,




FIGURE 8 Standardized residuals of the proposed day-ahead (top-left, green) and month-ahead (top-left, blue) model and corresponding densities (top-right) as well as sample autocorrelations (ACF) up to lag 40 of the day-ahead (green) and month-ahead (blue) standardized residuals and transformations thereof, 5% confidence intervals are indicated by grey dotted lines

Description	Variable	Statistic	p value
Day-ahead			
Day-ahead Price	y_t	20.147	0.010
Seasonal	$\sum_{s=t-A-2}^{t-A+2} y_s$	9.686	0.050
Coal price lagged	Coal_{t-1}	41.122	0.030
EUA price lagged	EUA_{t-1}	30.279	0.010
Power Price lagged	PWR_{t-1}	17.147	0.515
Monday Dummy	Monday	0.379	0.337
Standardized residuals lagged	z_{t-1}	1.240	0.485
Month-ahead			
Month-ahead price	y_t	10.027	0.056
Lagged month-ahead price	y_{t-1}	7.190	0.194
Lagged two-month-ahead price	$y_{t-1,MP2}$	6.940	0.178
First day of month dummy	First day of month	0.115	0.247
Days in month	D_t^M	0.999	0.280
Day since last close	D_t^{LTD}	1.281	0.146
Seasonal	\bar{y}_{t-1Y}^M	4.812	0.384
Lagged EUA price	EUA_{t-1}	14.787	0.182
Lagged oil price	Oil_{t-1}	12.370	0.311
Lagged temperature	T_{t-1}	8.322	0.509
Lagged standardized residual	z_{t-1}	1.008	0.663

TABLE 3 Test statistics and p values of the energy test of independence conducted to test whether z_t and the denoted listed are independent

Process	Parameter	Description	Estimate	\mathcal{H}_0	z-statistic
μ	λ	ETS smoothing	1.152	$\lambda = 1$	6.65
μ	ψ_0	Seasonal	-0.015	$\psi_0 = 0$	-1.29
μ	φ_1	AR1	-0.685	$\varphi_1 = 0$	-36.17
μ	θ_1	MA1	0.413	$\theta_1 = 0$	16.84
μ	ζ_0	Coal	0.019	$\zeta_0 = 0$	3.70
μ	ζ_1	EUA	0.078	$\zeta_1 = 0$	3.92
μ	ζ_2	Power	0.001	$\zeta_2 = 0$	1.72
σ	ω	Constant	0.025	$\omega = 0$	15.52
σ	α	AR1	0.291	$\alpha = 0$	38.98
σ	β	MA1	0.726	$\beta = 0$	128.60
σ	$\alpha + \beta$	Persistence	1.027	$\alpha + \beta = 1$	0.85
σ	γ	Leverage	-0.101	$\gamma = 0$	-9.38
σ	δ	Monday	1.342	$\delta = 1$	5.05
ν	ν	Nu	1.039	$\nu = 1$	1.05
τ	τ	Tau	6.425	$\tau > 4$	2.91

TABLE 4 Parameter estimates and test statistics of corresponding z-tests of the proposed day-ahead model along with significance levels indicated by colors

Note: $p \leq 0.10$, $p \leq 0.01$, $p \leq 0.001$.

the p value of the dependent variable is close to rejection, which may indicate some residual dependencies.

The maximum likelihood estimates themselves are asymptotically normal under some regularity conditions,

which often hold in practice. That is, we assume asymptotic normality and conduct z-tests to assess the parameter significance. We approximate the variance-covariance matrix by inverting the Hessian, which is calculated by

numerical differentiation using the symmetric difference quotient. The obtained Hessian will be equivalent to the Fisher information matrix, as all models were estimated by minimizing the negative log-likelihood function. Therefore, inverting it yields an estimate for the variance–covariance matrix. We approximated the variance–covariance matrix in every iteration step of the forecasting study and averaged afterward. This makes the estimated Hessian more robust, and the standard errors are not as low as they would be when only evaluating the last step of the expanding window study.

Table 4 summarizes the parameter estimates of the proposed day-ahead model along with their significance estimates. The smoothing parameter λ is significantly different from one, which indicates that the conditional mean process is not purely a random walk. The seasonality parameter is insignificant, although it reduces the out-of-sample CRPS (see Appendix A1). The coal and EUA parameters are significant, whereas the power parameter only reaches weak significance. The ARMA(1,1) process of the error-trend-seasonal (ETS) model is highly significant, indicating that an exponential smoothing model alone does not suffice in describing the process. All parameters of the TGARCH process are significant. The high test statistics of the AR and moving average (MA) terms suggest that a higher order TGARCH model might be even better. The skewness parameter ν is not significantly different from one. A value of one corresponds to asymmetric distribution. The estimated degrees of freedom τ are significantly higher than four. We can conclude that the fourth moment exists with high

certainty since the degrees of freedom function as a tail-index estimate for the T -distribution. This property is important because it renders inference on the TGARCH parameters, for which the fourth moment is needed, plausible.

The positive signs of coal and oil prices confirm previous findings in related research, as discussed in Section 1. The effect of the EUA was estimated to be significantly positive as well, which indicates that the day-ahead product is used to substitute coal for power production.

The α and β (the persistence parameters in the TGARCH process) sum up to 1.027. Conducting a z -test with $\mathcal{H}_0: \alpha + \beta = 1$ yields test statistic of 0.85. That is, $\alpha + \beta$ is not significantly different from 1. The latter is a characteristic property of Integrated GARCH (IGARCH) processes. Shocks have an infinite persistence in the (IGARCH) process. In consequence, long term forecasts of the conditional variance will not converge towards the unconditional variance, which is undefined in the (IGARCH) model. Therefore, forecasts will get less reliable as the forecasting horizon increases (Tsay, 2010). (IGARCH) behavior can be caused by unaccounted structural change in the conditional variance (Morana, 2002). However, whether those findings also apply to the proposed TGARCH model is questionable since it deviates from the plain GARCH model in various aspects. However, we did estimate the proposed model with $\alpha + \beta$ fixed to 1. This did not improve the forecasting performance (see Appendix A1). Therefore, further examining this issue remains a task for future scientific work.

TABLE 5 Parameter estimates and test statistics of corresponding z -tests of the proposed month-ahead model, along with significance levels indicated by colors

Process	Parameter	Description	Estimate	H0	z -statistic
μ	η	Risk premia	−0.353	$\eta = 0$	−3.36
μ	ψ_0	Seasonal	−0.005	$\psi_0 = 0$	−19.29
μ	φ_0	Constant	0.112	$\varphi_0 = 0$	24.78
μ	φ_1	AR1	0.997	$\varphi_1 = 1$	−11.38
μ	ζ_0	EUA	−0.005	$\zeta_0 = 0$	−6.45
μ	ζ_1	Oil	0.001	$\zeta_1 = 0$	16.19
μ	ζ_2	Temperature	0.001	$\zeta_2 = 0$	3.18
σ	ω	Constant	0.012	$\omega = 0$	13.23
σ	α	AR1	0.208	$\alpha = 0$	43.27
σ	β	MA1	0.834	$\beta = 0$	241.51
σ	$\alpha + \beta$	Persistence	1.043	$\alpha + \beta = 1$	2.59
σ	γ	Leverage	−0.085	$\gamma = 0$	−10.47
σ	δ	First day	1.403	$\delta = 1$	2.73
ν	ν	Nu	1.132	$\nu = 1$	3.38
τ	τ	Tau	6.858	$\tau > 4$	3.03

Note: $p \leq 0.10$, $p \leq 0.01$, $p \leq 0.001$.

Table 5 presents the parameter estimates and test statistics of the proposed month-ahead model. The risk parameter is negative as expected, and it is highly significant. The EUA parameter is significantly negative, which contrasts the positive estimate that was observed in the day-ahead model. This potentially indicates that the month-ahead product is not used for substitution but is rather influenced on its own. While this is a plausible explanation, one has to keep in mind that *day-ahead* EUA prices were used because other time series of EUA prices were not available. Analyzing these effects in more detail remains a topic for further research. However, while the sign of oil is plausible and matches the one from the day-ahead model, the sign of the temperature is not plausible. Lower temperatures should lead to a higher natural gas demand, and higher demand usually translated to higher prices if it is not fully anticipated upfront. The sign of the seasonal parameter is also unexpectedly negative. All parameters in the TGARCH process are significant. The δ estimates are significant around 1.4, indicating that the conditional variance is significantly increased on the first trading day of every month. The estimated degrees of freedom τ are larger than four as for the day-ahead model.

5.2 | Predictive performance

Table 6 shows the forecasting performance of the day-ahead and month-ahead models. The proposed models

outperform their closest benchmarks by $\approx 13\%$ and $\approx 9\%$ in terms of CRPS, respectively. The assumption of T -distributed errors is beneficial for all day-ahead models except the random forest. Interestingly, the normality assumption yields better performance for the benchmark models. The results further support the hypothesis that Oil Prices contain predictive power as the VAR model consistently outperforms the ARIMA counterpart. However, the results of the Random Forest are not in line with the results of Herrera et al. (2019) where the random forest substantially outperformed all other considered models.

However, the results in Table 6 do not reveal how these improvements are distributed over the predictive distribution. Figure 9 closes this gap. It shows the pinball loss of both proposed models and their benchmarks. The figure presents the level of the pinball loss (top), as well as the improvement in % relative to the best performing benchmark (bottom). Both proposed models outperform their benchmarks over the whole probability grid. Moreover, substantial improvements are realized in the distribution's tails, which are particularly relevant for risk management.

We conducted the DM-test to check whether the observed improvements in CRPS are significant. Figure 10 contains the results. The colors indicate the value of the test statistic, whereas the symbols note significance. Lower values (dark green) indicate that the model on the y-axis outperforms the model on the x-axis. All models outperform the Random Forest, while the

	Reference	MAE	RMSE	CRPS
Day-ahead				
Proposed model		0.3863	1.0843	0.2834
VAR (T)		0.4144	1.1598	0.3252
ARIMA(2,1,2) (T)		0.4215	1.2127	0.3324
VAR (normal)	(Geng et al., 2017)	0.4144	1.1598	0.3386
ARIMA(2,1,2) (normal)	(Siddiqui, 2019)	0.4215	1.2127	0.3466
Random forest (normal)	(Herrera et al., 2019)	0.5725	1.2548	0.4361
Random forest (T)		0.5725	1.2548	0.4523
Month-ahead				
Proposed model		0.3010	0.3995	0.2126
VAR (normal)	(Geng et al., 2017)	0.3184	0.4619	0.2336
ARIMA(2,1,2) (normal)	(Siddiqui, 2019)	0.3201	0.4635	0.2349
VAR (T)		0.3184	0.4619	0.2356
ARIMA(2,1,2) (T)		0.3201	0.4635	0.2371
Random forest (normal)	(Herrera et al., 2019)	0.4681	0.6167	0.3502
Random forest (T)		0.4681	0.6167	0.3603

TABLE 6 Predictive performance of the proposed models and benchmarks

Note: Lower values indicate better performance.




FIGURE 9 Pinball loss levels at all percentiles for the day-ahead (top-left) and month-ahead (top-right) models as well as the relative improvement compared to the best performing benchmark model - the VAR (*T*) model for day-ahead (bottom-left) and VAR (Normal) for month-ahead (bottom-right)

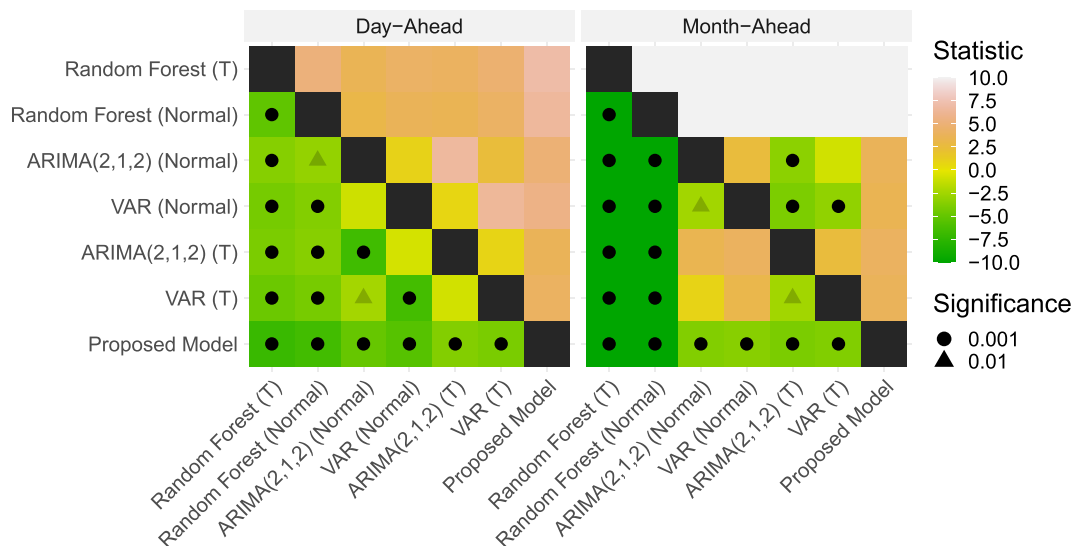


FIGURE 10 Results of DM-tests testing whether the *y*-axis model performs better than the *x*-axis model. The colors correspond to levels of the test statistics. Symbols denote selected significance levels

VAR outperforms most other models for both products. The proposed models significantly outperform all benchmarks, which makes them preferable.

We also tested if the observed differences in pinball loss (see Figure 9) are significant. Therefore, we used the DM-test with respect to the pinball loss of each percentile. That is, we can present the results concisely in Figure 11. It shows the *p* values (top) and test-statistics

(bottom). Lower values indicate that the proposed model outperforms the benchmark. Furthermore, all *p* values are below 5% for all benchmarks. This means that the proposed models significantly outperform the benchmarks at all percentiles.

Lastly, we constructed a probability integral transform (PIT) histogram for both proposed models to evaluate the calibration accuracy; see Figure 12. Each bin




FIGURE 11 P-values (top) and test statistics (bottom) of DM tests testing whether the predictions of the proposed models are better than those from the ARIMA and VAR models. Lower values indicate better performance of the proposed models




FIGURE 12 PIT histogram of the probabilistic out-of-sample predictions of both proposed models

represents a probability segment of the predictive distribution with a predicted probability of 5%. A perfectly calibrated model would produce a uniformly shaped histogram. That is, the figure indicates that both models are calibrated reasonably well while the month-ahead model has a slight tendency of underdispersion (Raftery et al., 2005).

To keep this paper concise, we decided to present the proposed model only instead of presenting each model considered during the model building process. We only discussed the proposed models in detail to keep this paper concise. However, we considered several variations of the two proposed models, either leaving out (–) or adding (+) model components. This is helpful to assess the relative importance of each component. We used the DM-test with respect to the CRPS to compare the predictive performance of these models. Figure 13 shows the results. The figure can be interpreted as follows. The main body shows the test-statistics of DM-tests testing

whether the y -axis model outperforms the x -axis model. Therefore, inspecting the last column of the day-ahead facet answers whether any model (denoted on the y -axis) outperforms the proposed model (the last entry on the x -axis). The latter is not the case, which means that neither adding nor removing model components yields a significantly better predictive performance. This also holds true for the proposed month-ahead model.

Some findings are worth highlighting here. First, the models assuming homoscedasticity are among the worst for both products. The day-ahead model profits in particular from including ARMA errors, coal, and the leverage effect in the sigma process. The month-ahead model profits from considering the rollover (First Day Mu) as well as including EUA and the leverage effect. Removing components not always produces a significant difference. For example, removing EUA prices from the proposed day-ahead model does not yield a significantly worse performance with respect to the DM-test. However, the EUA




FIGURE 13 Results of DM-tests with respect to CRPS, testing whether the y-axis model performs better than the x-axis model. Colors denote the level of the test statistics. Symbols indicate selected significance levels

coefficient itself is significant, as shown in Table 4 and including EUA prices reduces the out-of-sample CRPS. Therefore, we keep EUA prices included in the proposed model.

CRPS, RMSE, and MAE values of all models are summarized in Tables A1 and A2 in Appendix A. Results of DM-tests testing with respect to the absolute and squared error are also included in Appendix A (see Figures A1 and A2).

6 | CONCLUSION

In this paper, we examined the problem of forecasting European natural gas prices. We developed state-space time series models for probabilistic price forecasting of the day-ahead and month-ahead products. Both time series are heavy-tailed and feature conditional heteroscedasticity with leverage effects. In addition, the month-ahead time series contains risk premia effects. The model development process was guided by an extensive exploratory data analysis (see Section 2) and results from related research. The econometric models were estimated using maximum likelihood under skewed *T*-distribution. Out-of-sample forecasts obtained from an expanding window forecasting study were used to evaluate the forecasting performance using various

probabilistic performance measures, including the CRPS and the pinball loss.

Particularly, the day-ahead time series features pronounced heavy tails and benefits substantially from assuming a *t*-distribution. We modeled the conditional mean using a state-space model with ARMA errors, a seasonal component, and external regressors. We used a stepwise forward selection approach to determine the influence of external regressors. The results show a significant positive impact on the coal price. The latter confirms previous findings (Papież & Śmiech, 2011). Additionally, we find a positive influence of EUA and power peak prices on natural gas prices. This indicates a strong link between natural gas and electricity markets in Europe. Finally, the conditional standard deviation was modeled using a TGARCH process, extended by a level elevation on Mondays. As a result, the estimated volatility is increased by around 34% on Mondays. The proposed day-ahead model improves the CRPS by 12.9% compared with a VAR model with *t*-distributed errors, which is the best benchmark.

A major consideration for the month-ahead time series was the presence of risk premia and the rollover problem. The latter implies a sudden change in the delivery period on the first trading day of each month. This leads to pronounced price effects and, as a consequence, increased conditional volatility. The proposed model

covers these effects by using the last observed price of the Two-month-ahead product as a predictor for the first trading day and by elevating the conditional standard deviation on the first trading day. The latter increase is estimated to be 40% in the proposed month-ahead model. The observed positive changes at the beginning of the month are possibly due to risk premia. These risk premia should decrease over time because the uncertainty decreases as time gets closer to delivery. We find this highly significant pattern in the data estimating a total price reduction of 0.35£ per month. Including a seasonal component as well as EUA prices, oil prices, and the temperature reduced the CRPS. However, only the effect of EUA prices is significant. The impact is negative, which might indicate that the month-ahead product is not used for substituting coal, in contrast to the day-ahead model, which is positively influenced by EUA prices. The leverage component in the GARCH process is significant and negative in the proposed month-ahead model. Using the proposed month-ahead model compared to the VAR alternative with Gaussian errors leads to a 9% reduction in out of sample CRPS.

Some additional questions arise in light of the results. Further, analyzing the influence of EUA prices seems interesting for energy companies and governments. The main question to be answered here is why the impact of EUA prices on natural gas prices is positive for the day-ahead product while it is negative for the month-ahead product. Further investigation of the role of risk premia in natural gas markets is another attractive topic since the literature on natural gas price forecasting has not covered it yet.

DATA AVAILABILITY STATEMENT

Research data are not shared. The data that support the findings of this study are available from the corresponding author upon reasonable request.

ORCID

Jonathan Berrisch  <https://orcid.org/0000-0002-4944-9074>

Florian Ziel  <https://orcid.org/0000-0002-2974-2660>

REFERENCES

- Asche, F., Osmundsen, P., & Sandsmark, M. (2006). The UK market for natural gas, oil and electricity: Are the prices decoupled? *The Energy Journal*, 27(2), 27–40. <https://doi.org/10.5547/issn0195-6574-ej-vol27-no2-2>
- Berrisch, J. (2021). sstudentt 0.1.1: A python implementation of the skewed student-t distribution. <https://pypi.org/project/sstudentt/>, <https://github.com/BerriJ/sstudentt>
- Berrisch, J., & Ziel, F. (2021). CRPS Learning. *Journal of Econometrics*. forthcoming. <https://doi.org/10.1016/j.jeconom.2021.11.008>
- Bollerslev, T., Russell, J. R., & Watson, M. W. (2010). *Volatility and time series econometrics: essays in honor of Robert Engle*: Oxford University Press.
- Čeperić, E., Žiković, S., & Čeperić, V. (2017). Short-term forecasting of natural gas prices using machine learning and feature selection algorithms. *Energy*, 140, 893–900. <https://doi.org/10.1016/j.energy.2017.09.026>
- Chaton, C., Creti, A., & Villeneuve, B. (2008). Some economics of seasonal gas storage. *Energy Policy*, 36(11), 4235–4246.
- Chen, Y., Xu, X., & Koch, T. (2020). Day-ahead high-resolution forecasting of natural gas demand and supply in Germany with a hybrid model. *Applied Energy*, 262, 114486.
- Chkili, W., Hammoudeh, S., & Nguyen, D. K. (2014). Volatility forecasting and risk management for commodity markets in the presence of asymmetry and long memory. *Energy Economics*, 41, 1–18.
- De Livera, A. M., Hyndman, R. J., & Snyder, R. D. (2011). Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American statistical association*, 106(496), 1513–1527.
- Diebold, F. X., & Mariano, R. S. (2002). Comparing predictive accuracy. *Journal of Business & economic statistics*, 20(1), 134–144.
- Drachal, K. (2021). Forecasting selected energy commodities prices with Bayesian dynamic finite mixtures. *Energy Economics*, 99, 105283.
- Elgharabawy, A., & Sayed, A. (2020). A review on natural gas previous, current and forecasting prices and demand. *Journal of Petroleum and Mining Engineering*, 22(1), 61–64.
- Engle, R. F., & Patton, A. J. (2007). What good is a volatility model? *Forecasting volatility in the financial markets*: Elsevier, pp. 47–63.
- Fernández, C., & Steel, M. F. J. (1998). On Bayesian modeling of fat tails and skewness. *Journal of the American Statistical Association*, 93(441), 359–371.
- Ferrari, D., Ravazzolo, F., & Vespignani, J. (2021). Forecasting energy commodity prices: A large global dataset sparse approach. *Energy Economics*, 98, 105268.
- 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(3), 1038–1050.
- Gao, S., Hou, C., & Nguyen, B. H. (2021). Forecasting natural gas prices using highly flexible time-varying parameter models. *Economic Modelling*, 105, 105652.
- Geng, J.-B., Ji, Q., & Fan, Y. (2017). The relationship between regional natural gas markets and crude oil markets from a multi-scale nonlinear Granger causality perspective. *Energy Economics*, 67, 98–110. <https://doi.org/10.1016/j.eneco.2017.08.006>
- Gneiting, T. (2011). Making and evaluating point forecasts. *Journal of the American Statistical Association*, 106(494), 746–762.
- Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477), 359–378.
- Hailemariam, A., & Smyth, R. (2019). What drives volatility in natural gas prices?. *Energy Economics*, 80, 731–742.
- Hamie, H., Hoayek, A., & Auer, H. (2021). Modeling post-liberalized European gas market concentration—A game

- theory perspective. *Forecasting*, 3(1), 1–16. <https://doi.org/10.3390/forecast3010001>
- Herrera, G. P., Constantino, M., Tabak, B. M., Pistori, H., Su, J.-J., & Naranpanawa, A. (2019). Long-term forecast of energy commodities price using machine learning. *Energy*, 179, 214–221. <https://doi.org/10.1016/j.energy.2019.04.077>
- 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.
- Hyndman, P. R., Koehler, P. A., Ord, P. K., & Snyder, A. P. R. (2008). *Forecasting with exponential smoothing*, Springer Series in Statistics: Springer Berlin Heidelberg.
- Jianliang, W., Changran, L., & Meiyu, G. (2020). Daily natural gas price forecasting by a weighted hybrid data-driven model. *Journal of Petroleum Science and Engineering*, 192, 107240. <https://doi.org/10.1016/j.petrol.2020.107240>
- Karabiber, O. A., & Xydis, G. (2020). Forecasting day-ahead natural gas demand in Denmark. *Journal of Natural Gas Science and Engineering*, 76, 103193.
- Li, H., Chen, L., Wang, D., & Zhang, H. (2017). Analysis of the price correlation between the international natural gas and coal. *Energy Procedia*, 142, 3141–3146.
- Liang, C., Ma, F., Wang, L., & Zeng, Q. (2021). The information content of uncertainty indices for natural gas futures volatility forecasting. *Journal of Forecasting*, 40(7), 1310–1324.
- Livieris, I. E., Pintelas, E., Kiriakidou, N., & Stavroyiannis, S. (2020). An advanced deep learning model for short-term forecasting US natural gas price and movement. *IFIP International Conference on Artificial Intelligence Applications and Innovations*: Springer, pp. 165–176.
- Lu, X., Salovaara, J., & McElroy, M. B. (2012). Implications of the recent reductions in natural gas prices for emissions of CO₂ from the US power sector. *Environmental science & technology*, 46(5), 3014–3021.
- Lv, X., & Shan, X. (2013). Modeling natural gas market volatility using GARCH with different distributions. *Physica A: Statistical Mechanics and its Applications*, 392(22), 5685–5699.
- Mirantes, A. G., Población, J., & Serna, G. (2012). The stochastic seasonal behaviour of natural gas prices. *European Financial Management*, 18(3), 410–443.
- Morana, C. (2002). IGARCH effects: an interpretation. *Applied Economics Letters*, 9(11), 745–748.
- Mu, X. (2007). Weather, storage, and natural gas price dynamics: Fundamentals and volatility. *Energy Economics*, 29(1), 46–63.
- Nguyen, H. T., & Nabney, I. T. (2010). Short-term electricity demand and gas price forecasts using wavelet transforms and adaptive models. *Energy*, 35(9), 3674–3685.
- Nowotarski, J., & Weron, R. (2018). Recent advances in electricity price forecasting: A review of probabilistic forecasting. *Renewable and Sustainable Energy Reviews*, 81, 1548–1568.
- Papież, M., & Śmiech, S. (2011). The analysis of relations between primary fuel prices on the European market in the period 2001–2011. *Rynek Energii*, 5(96), 139–144.
- Petropoulos, F., Apiletti, D., Assimakopoulos, V., Babai, M. Z., Barrow, D. K., Taieb, S. B., Bergmeir, C., Bessa, R. J., Bijak, J., Boylan, J. E., Browell, J., Carnevale, C., Castle, J. L., Cirillo, P., Clements, M. P., Cordeiro, C., Oliveira, F. L. C., De Baets, S., Dokumentov, A., ..., & Ziel, F. (2022). Forecasting: theory and practice. *International Journal of Forecasting*. forthcoming. <https://doi.org/10.1016/j.ijforecast.2021.11.001>
- Pratson, L. F., Haerer, D., & Patiño-Echeverri, D. (2013). Fuel prices, emission standards, and generation costs for coal vs natural gas power plants. *Environmental science & technology*, 47(9), 4926–4933.
- Raftery, A. E., Gneiting, T., Balabdaoui, F., & Polakowski, M. (2005). Using Bayesian model averaging to calibrate forecast ensembles. *Monthly weather review*, 133(5), 1155–1174.
- Roncoroni, A., & Brik, R. I. (2017). Hedging size risk: Theory and application to the US gas market. *Energy Economics*, 64, 415–437. <https://doi.org/10.1016/j.eneco.2016.10.020>
- Safari, A., Das, N., Langhelle, O., Roy, J., & Assadi, M. (2019). Natural gas: A transition fuel for sustainable energy system transformation?. *Energy Science & Engineering*, 7(4), 1075–1094.
- Saghi, F., & Rezaee, M. J. (2021). An ensemble approach based on transformation functions for natural gas price forecasting considering optimal time delays. *PeerJ Computer Science*, 7, e409.
- Salehnia, N., Falahi, M. A., Seifi, A., & Adeli, M. H. M. (2013). Forecasting natural gas spot prices with nonlinear modeling using Gamma test analysis. *Journal of Natural Gas Science and Engineering*, 14, 238–249. <https://doi.org/10.1016/j.jngse.2013.07.002>
- Siddiqui, A. W. (2019). Predicting natural gas spot prices using artificial neural network. In 2019 2nd International Conference on Computer Applications & Information Security (ICCAIS), Riyadh, Saudi Arabia, pp. 1–6.
- Su, M., Zhang, Z., Zhu, Y., & Zha, D. (2019). Data-driven natural gas spot price forecasting with least squares regression boosting algorithm. *Energies*, 12(6), 1094.
- Su, M., Zhang, Z., Zhu, Y., Zha, D., & Wen, W. (2019). Data driven natural gas spot price prediction models using machine learning methods. *Energies*, 12(9), 1680.
- Székely, G. J., & Rizzo, M. L. (2013). Energy statistics: A class of statistics based on distances. *Journal of statistical planning and inference*, 143(8), 1249–1272.
- Székely, G. J., Rizzo, M. L., Bakirov, N. K., et al. (2007). Measuring and testing dependence by correlation of distances. *The Annals of Statistics*, 35(6), 2769–2794. <https://doi.org/10.1214/009053607000000505>
- Tamba, J. G., Essiane, N., Sapnken, E. F., Koffi, F. D., Nsouand, J. L., Soldo, B., & Njomo, D. (2018). Forecasting natural gas: A literature survey. *International Journal of Energy Economics and Policy*, 8(3), 216.
- Tang, Y., Wang, Q., Xu, W., Wang, M., & Wang, Z. (2019). Natural gas price prediction with big data, 2019 IEEE International Conference on Big Data (Big Data). Los Angeles, CA, USA: IEEE, pp. 5326–5330.
- Tsay, R. S. (2010). *Analysis of financial time series*: John Wiley & Sons.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., ..., van Mulbregt, P., & SciPy 1.0 Contributors (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17, 261–272.
- Wang, J., Cao, J., Yuan, S., & Cheng, M. (2021). Short-term forecasting of natural gas prices by using a novel hybrid method based

on a combination of the CEEMDAN-SE and the PSO-ALS-optimized GRU network. *Energy*, 233, 121082.

- Wang, T., Zhang, D., & Broadstock, D. C. (2019). Financialization, fundamentals, and the time-varying determinants of US natural gas prices. *Energy Economics*, 80, 707–719. <https://doi.org/10.1016/j.eneco.2019.01.026>
- Wang, Y., Liu, L., & Wu, C. (2020). Forecasting commodity prices out-of-sample: Can technical indicators help? *International Journal of Forecasting*, 36(2), 666–683.
- Wurtz, D., Chalabi, Y., & Luksan, L. (2006). Parameter Estimation of ARMA Models with GARCH/APARCH Errors an R and S Plus Software Implementation. *Journal of Statistical Software*, 55(2), 28–33.

AUTHOR BIOGRAPHIES

Jonathan Berrisch is a PhD candidate at the University of Duisburg-Essen. His research focuses on forecast combination, probabilistic forecasting, and forecasting in energy markets. In addition, he is passionate about open-source software, to which he contributes regularly.

Florian Ziel is Assistant Professor of Environmental Economics at the House of Energy Markets and

Finance at the University of Duisburg-Essen, Germany. He received his MSc in Statistics from University College Dublin (Ireland, 2012), his Diplom in mathematics from Dresden University of Technology (Germany, 2013), and his PhD on forecasting in energy markets from the European-University Viadrina in Frankfurt Oder (Germany, 2016). His research interests include data analytics with application to energy markets and systems. He is the first author of various peer-reviewed journal articles, most notably in top-tier IEEE Transactions on Power Systems, Applied Energy, Energy Economics, Renewable and Sustainable Energy Reviews, and International Journal of Forecasting.

How to cite this article: Berrisch, J., & Ziel, F. (2022). Distributional modeling and forecasting of natural gas prices. *Journal of Forecasting*, 41(6), 1065–1086. <https://doi.org/10.1002/for.2853>

APPENDIX A: RESULTS OF REDUCED MODELS

A.1 Day-ahead

TABLE A1 Predictive performance of the proposed day-ahead model and variations thereof sorted by the relative difference in CRPS (last column)

Num.	Changed component	MAE	RMSE	CRPS	CRPS Δ (%)
19	Proposed model	0.3863	1.0843	0.2834	0.0000
18	+ Temp (smoothed)	0.3866	1.0844	0.2834	0.0015
17	- Seasonal	0.3866	1.0846	0.2836	0.0495
16	+ Storage	0.3862	1.0840	0.2836	0.0520
15	+ Power base	0.3861	1.0837	0.2836	0.0528
14	IGARCH	0.3864	1.0843	0.2837	0.0871
13	+ Oil ahead 3	0.3869	1.0843	0.2837	0.0951
12	+ Oil ahead 1	0.3869	1.0847	0.2837	0.1077
11	+ Oil ahead 6	0.3870	1.0847	0.2838	0.1262
10	- Power	0.3866	1.0835	0.2838	0.1277
09	+ Oil ahead 9	0.3870	1.0842	0.2839	0.1549
08	+ Temperature	0.3872	1.0842	0.2841	0.2332
07	- Eua	0.3865	1.0853	0.2842	0.2857
06	- Monday elevation	0.3862	1.0843	0.2844	0.3417
05	- Arma errors	0.3868	1.0868	0.2845	0.3756
04	- Coal	0.3889	1.0849	0.2850	0.5677
03	- Exp. smoothing	0.3878	1.1010	0.2860	0.9132
02	- Leverage	0.3865	1.0843	0.2863	1.0169
01	Const. sigma	0.3885	1.0826	0.3053	7.7054




FIGURE A1 Results of the DM-test for the proposed day-ahead model and variations thereof with respect to absolute and squared error. The symbols indicate the 0.1, 0.05, 0.01, and 0.001 significance levels

A.2 Month-ahead

Num.	Changed component	MAE	RMSE	CRPS	CRPS Δ (%)
19	– Risk	0.3011	0.3997	0.2126	0.0238
18	Proposed model	0.3010	0.3995	0.2126	0.0000
17	– Sigma elevation	0.3011	0.3996	0.2126	–0.0032
16	– Temp. smooth	0.3009	0.3997	0.2127	–0.0287
15	– Temperature	0.3009	0.3999	0.2127	–0.0461
14	– Oil ahead 3	0.3009	0.4000	0.2128	–0.0836
13	– Seasonal	0.3010	0.3996	0.2129	–0.1213
12	IGARCH	0.3010	0.3993	0.2129	–0.1452
11	– Leverage	0.3011	0.3995	0.2132	–0.2744
10	+ Storage	0.3018	0.4006	0.2133	–0.3412
09	– Eua	0.3021	0.4010	0.2137	–0.4991
08	+ Power peak	0.3029	0.4011	0.2140	–0.6541
07	+ Power base	0.3036	0.4018	0.2146	–0.9087
06	+ Coal	0.3073	0.4029	0.2167	–1.9002
05	Const. sigma	0.3017	0.4009	0.2205	–3.7067
04	– First day Mu	0.3170	0.4617	0.2276	–7.0538
03	+ Oil ahead 9	0.3861	0.5133	0.2839	–33.5322
02	+ Oil ahead 6	0.5194	0.6788	0.3984	–87.3583
01	+ Oil ahead 1	0.6173	0.7443	0.4699	–120.9902

TABLE A2 Predictive performance of the proposed month-ahead model and variations thereof sorted by the relative difference in CRPS (last column)




FIGURE A2 Results of the DM-test for the proposed model and variations thereof with respect to absolute and squared error. The symbols indicate the 0.1, 0.05, 0.01, and 0.001 significance levels

DuEPublico

Duisburg-Essen Publications online

UNIVERSITÄT
D U I S B U R G
E S S E N

Offen im Denken

ub

universitäts
bibliothek

This text is made available via DuEPublico, the institutional repository of the University of Duisburg-Essen. This version may eventually differ from another version distributed by a commercial publisher.

DOI: 10.1002/for.2853

URN: urn:nbn:de:hbz:465-20231009-113048-7



This work may be used under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 License (CC BY-NC-ND 4.0).