Forecasting the Distributions of Hourly Electricity Spot Prices

- Accounting for Cross Correlation Patterns and Non-Normality of Price Distributions

Arne Vogler
Co-Authors: Christian Pape, Oliver Woll and Christoph Weber

Rome, 4th November 2017

UNIVERSITÄT DUISBURG ESSEN
Open-Minded

The Rational behind Distribution Forecasts

- Despite being well established in other fields of time series analysis, distribution forecasting has only recently gained traction in electricity price forecasting (Nowotarski and Weron 2017).
- Yet, increased production of variable RES causes higher uncertainty.
- Thus, the usage of point forecasts only reduces the quality of decision making, due to the reduced amount of information provided.
- Forecasting the distribution of hourly prices is more appropriate for
  - the valuation of assets' flexibilities and optionality,
  - short-term decision making such as dispatch,
  - and providing further information about forecast quality.
- We propose an econometric-stochastic model that combines several established approaches to adequately capture distribution characteristics.

Agenda

Forecasting Approach	1
Evaluation of Forecasting Quality	2
Applications and Results	3
Conclusion	4

Estimation and Simulation Procedure (I)

Forecasting Approach

- Determine the main deterministic drivers and the residuals
  - Regression model to account for deterministic factors

$$x_{t,h} = \beta_{0,h} + \beta_{1,h}(L_{t,h} - S_{t,h}) + \beta_{2,h}(W_{t,h}) + \beta_{3,h}(C_{coal,t}) + \beta_{4,h}\,(C_{Gas,t}) + \varepsilon_{t,h}$$

  - β_h: Regression coefficients
  - $L_{t,h}$: Load
  - $S_{t,h}$: Solar
  - $W_{t,h}$: Wind
  - $C_{Coal,t}$ and $C_{Gas,t}$: typical variable costs of power plants incl. emission costs
  - $\varepsilon_{t,h}$: Residuals

Estimation and Simulation Procedure (II)

Forecasting Approach

- Map the empirical CDF of residuals onto a normal distribution

$$T_h\colon \varepsilon_{t,h} \mapsto \Phi^{-1}\left(C_h\left(\varepsilon_{t,h}\right)\right)$$

  – C_h: Empirical CDF of residuals in hour h
  – Φ: CDF of the standard normal distribution
  – Graphical representation corresponds to Q-Q-plot (Quantile Mapping)

- Factor Model
  – Normal Residuals

$$u_{t,h} = T_h\left(\varepsilon_{t,h}\right)$$

  – The common factors are constructed using a principal component analysis on the correlation matrix of the transformed factors.

Estimation and Simulation Procedure (III)

Forecasting Approach

- Factor time series are modelled with specifications from the ARMA-GARCH class
  - ARMA (1,1)

$$f_{t,i} = \alpha_{1,i}\ f_{t-1.i} + \alpha_{2,i}\ w_{t-1,i} + w_{t,i}$$

    - α: Coefficients of ARMA part
    - $f_{t,i}$: Factors i from different time steps t
    - $w_{t,i}$: Innovation with $w_{t,i} \sim N(\mu, \sigma)$
  - GARCH (1,1)

$$\sigma_{t,i}^2 = \gamma_{0,i} + \gamma_{1,i}\ \sigma_{t-1.i}^2 + \gamma_{2,i}\ w_{t-1,i}^2$$

    - γ: Coefficients of GARCH part
    - $\mathrm{w}_{\mathrm{t}-1,\mathrm{i}}^2$: Lagged Squared Innovation
    - $\sigma_{\mathrm{t}-1.\mathrm{i}}^2$: Lagged Variance

Estimation and Simulation Procedure (IV)

UNIVERSITÄT DUISBURG ESSEN

Open-Minded

Forecasting Approach

- Monte Carlo Simulation
  - We reverse the estimation procedure using Monte Carlo simulations to derive prediction samples.
  - We essentially characterize the distribution of the price of hour h at day t as the empirical cumulative distribution function of a Monte Carlo prediction sample.

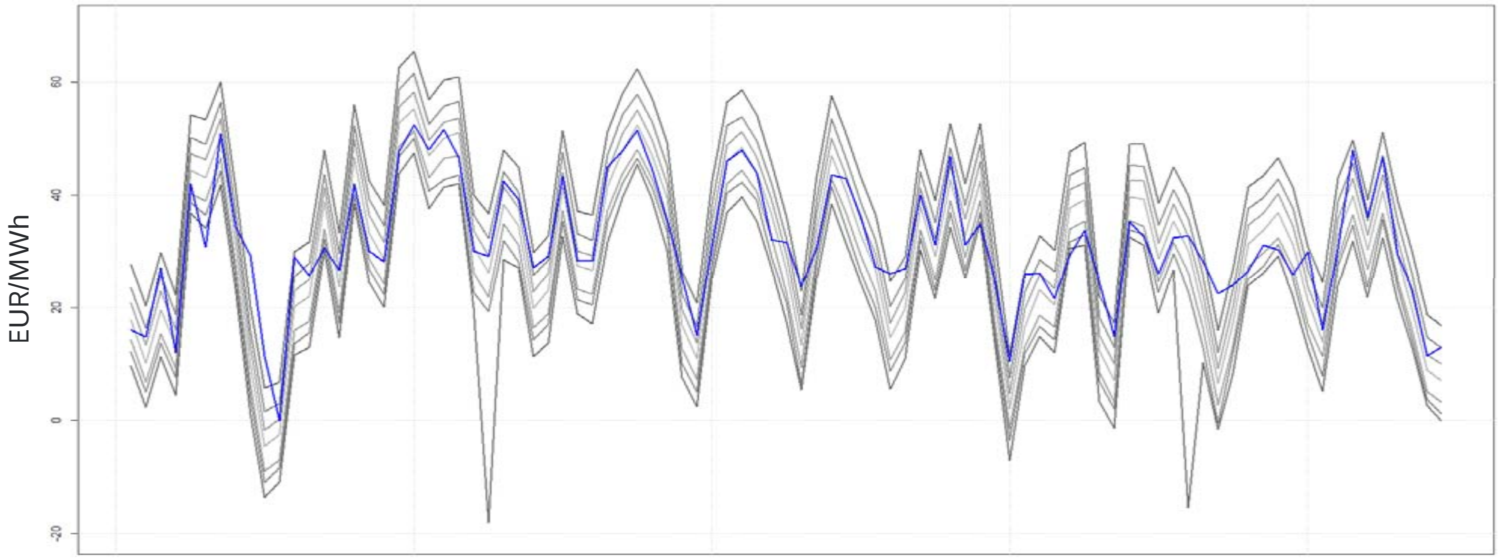

Predicted Quantiles versus Realizations
(Price of Hour 15 from 1st February to 30th March 2015)

The Evaluation Framework (I)

Evaluation of Forecast Quality

- Let $\{G_{Y_t|I_t}\}_{t=1}^{T}$ denote the sequence of conditional distributions, where $G_{Y_t|I_t} = \Pr(Y_t \leq y|I_t)$ denotes the conditional distribution of a random variable Y_t given some information set I_t.
- The corresponding sequence of one-step-ahead distribution forecasts is denoted by $\{F_{Y_t|I_t}\}_{t=1}^{T}$, where F may be of parametric or non-parametric form.
- We wish to test for calibration; that is, whether the distribution forecast sequence corresponds to the underlying distribution sequence.
- Hypotheses:

$$H_o: \{G_{Y_t|I_t}\}_{t=1}^{T} = \{F_{Y_t|I_t}\}_{t=1}^{T}$$

$$H_A: \{G_{Y_t|I_t}\}_{t=1}^{T} \neq \{F_{Y_t|I_t}\}_{t=1}^{T}$$

The Evaluation Framework (II)

Evaluation of Forecast Quality

- The evaluation of distribution forecasts rests on the probability integral transform (PIT), also know as Rosenblatt transformation (1952).
  - Under the null hypothesis the sequence of probability integral transforms (PITs), $\{F_{Y_t|\mathrm{I}_t}(Y_t)\}_{t=1}^{T}$, is uniformly distributed on $[0,1]$ and independent, given that I_t contains all relevant information.
  - The PIT sequence of the distribution forecasts, $\{G_{Y_t|\mathrm{I}_t}(Y_t)\}_{t=1}^{T}$, can be used to assess calibration.

- The graphical evaluation framework
  - The classic econometric testing framework rests on a graphical analysis of these PIT values.
  - Histogram and sample autocorrelation function

The Evaluation Framework (III)

Evaluation of Forecast Quality

- Evaluation and formal tests
  - Corradi and Swanson (2006) show that the PITs are still uniformly distributed on $[0,1]$, while the independence result cannot be upheld under the null hypothesis when I_t does not contain all relevant information.
  - Consequently, formal tests for uniformity of the PITs have to account for potential autocorrelation and classic distribution tests, which rely on i.i.d. observations, cannot be applied.
  - Knüppel (2015) proposes a test that is robust to autocorrelation and for which standard critical values can be used.

→ Uniformity constitutes a necessary but not sufficient condition for an ideal distribution forecast. We thus require the PIT values to be at least uniformly distributed.

→ Any dependence patterns may shed light on the characteristics of the information set underpinning our specifications.

Application (I)

Application and Results

- We test our econometric-stochastic approach against German day-ahead prices for 2015.
- We consider 6 different specifications in total for either the residual or the common factors.

	AR(1)	*AR(2)*	*ARMA(1,1)*	*GARCH(1,1)*	*PCA*
Model 1	x				
Model 2			x	x	
Model 3		x			
Model 4	x				x
Model 5			x	x	x
Model 6		x			x

- We calculate daily out-of-sample day-ahead forecasts using a rolling window of length 730; thus running 8760 Monte Carlo price simulation for each specification.

Application (II)

- We find little variation in point forecasting performance across the considered specifications as indicated by the point forecast error measures for individual hours.

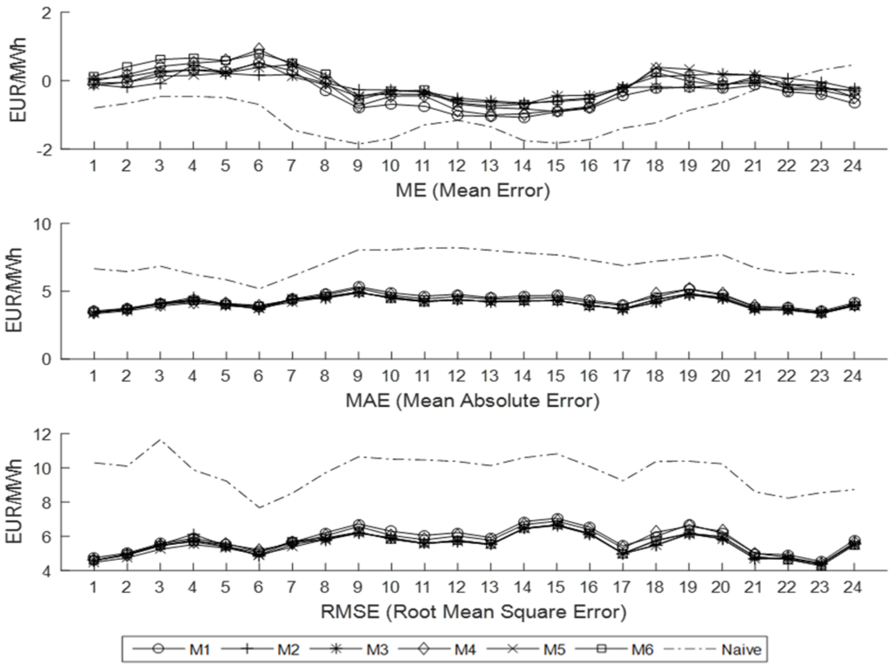

Application (III)

Application and Results

- Yet, we find strong variation in probabilistic forecasting performance as indicated by the histograms of PIT values across hours.

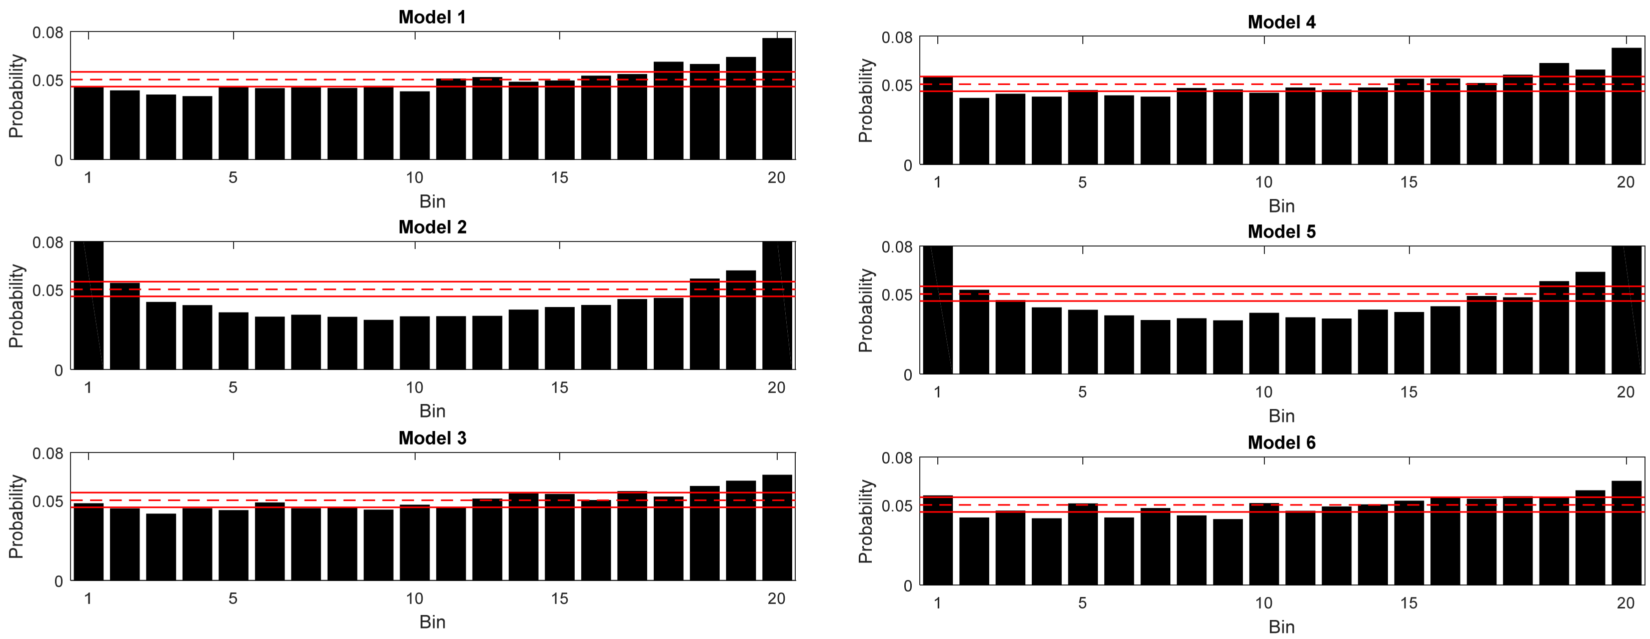

Application (IV)

Application and Results

- We fail to reject the null hypothesis of calibration for 22 hours of 2015 under the preferred specification (Model 3).

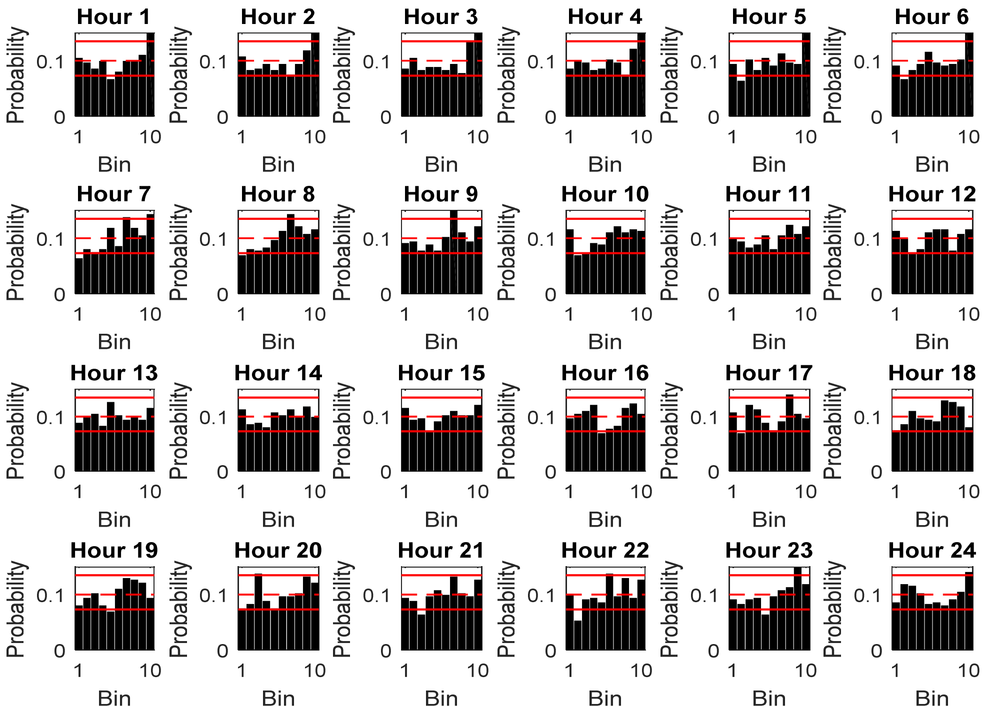

Application (IV)

Application and Results

- The formal calibration tests, due to Knüppel (2015), confirms the results of the preceding graphical analysis.

Hour	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*11*	*12*
Model 1	0	0	0	0	0	0	0	0	1	1	1	1
Model 2	0	0	0	0	0	0	0	0	0	0	0	0
Model 3	1	1	0	1	1	0	1	1	1	1	1	1
Model 4	0	0	0	0	0	0	0	0	1	1	1	1
Model 5	0	0	0	0	0	0	0	0	0	0	0	0
Model 6	1	0	0	0	1	1	1	1	1	1	1	1
Hour	*13*	*14*	*15*	*16*	*17*	*18*	*19*	*20*	*21*	*22*	*23*	*24*
Model 1	1	1	1	1	1	1	1	1	1	1	1	1
Model 2	0	0	0	0	0	0	0	0	0	0	0	1
Model 3	1	1	1	1	1	1	1	1	1	1	1	1
Model 4	1	1	1	1	1	1	1	1	1	1	1	1
Model 5	0	0	0	0	0	1	0	0	0	0	0	0
Model 6	1	1	1	1	1	0	1	1	1	1	1	1

Failure to reject the null hypothesis of calibration of the distribution forecast sequence is indicated by 1.

Conclusion

- The econometric-stochastic approach is able to capture the main characteristics of daily hourly prices in Germany and delivers calibrated distribution forecasts.

- A few comments on model particularities are warranted
  - Factor models adequately address cross correlations and ensure smooth price paths
  - Time-varying volatility seems to be less important for price processes of individual hours, as GARCH specifications do not improve results
  - The conditional distributions are correctly specified with respect to the considered information set; yet, dynamic misspecification seems to be present.
  - Since the latter can only be addressed by identifying and using the “relevant” information set, which is difficult in purely empirical applications, the extension of the evaluation framework by other criteria seems warranted in such settings.

Thank you for your attention!

Authors/ Speakers

Address

Backup: Q-Q-Plot

Forecasting Approach

- Quantile Mapping

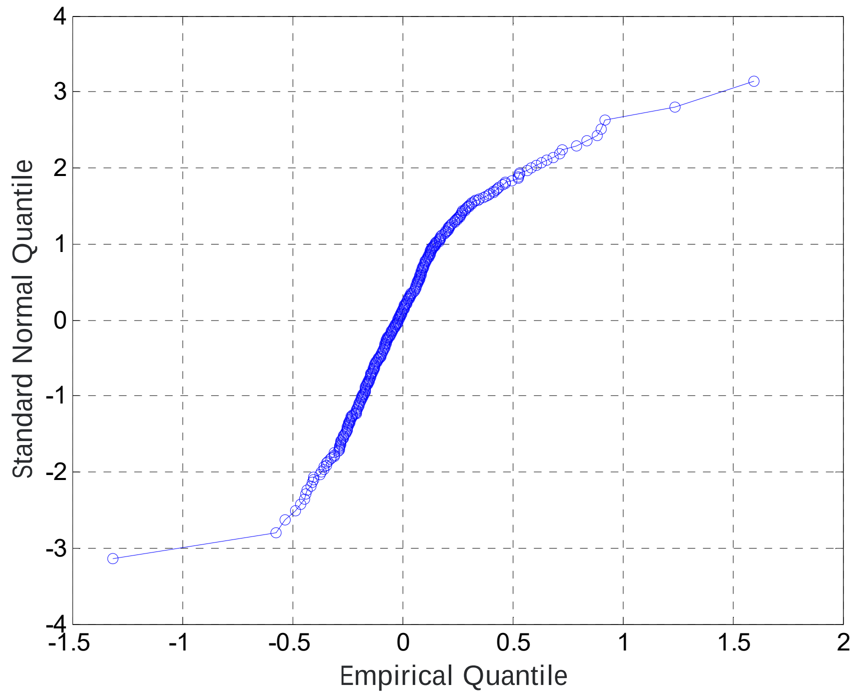

Backup: Orthogonal Factor Model

- The Orthogonal Factor Model (Johnson and Wichern (2002))

$$X = \mu + LF + \varepsilon$$

implies a specific covariance structure

$$\Sigma = LL' + \Psi, \qquad \Psi = Cov(\varepsilon)$$

which can be used to solve for factor loadings L and common factors F by spectral decomposition.

$$\Sigma = \underbrace{\left[\sqrt{\lambda_1}e_1, \dots, \sqrt{\lambda_p}e_p\right]}_{L}\left[\sqrt{\lambda_1}e_1, \dots, \sqrt{\lambda_p}e_p\right]'$$