Quantifying Volatility Reduction in German Day-ahead Spot Market in the Period 2006 through 2016

Abdolrahman Khoshrou
Intelligent and Autonomous Systems Group
Centrum Wiskunde & Informatica
The Netherlands
Email: a.khoshrou@cwi.nl

Eric J. Pauwels
Intelligent and Autonomous Systems Group
Centrum Wiskunde & Informatica
The Netherlands
Email: Eric.Pauwels@cwi.nl

Abstract—In Europe, Germany is taking the lead in the switch from the conventional to renewable energy. This poses new challenges as wind and solar energy are fundamentally intermittent, weather-dependent and less predictable. It is therefore of considerable interest to investigate the evolution of price volatility in this post-transition era. There are a number of reasons, however, that makes the practical studies difficult. For instance, EPEX prices can be zero or negative. Consequently, the standard approach in financial time series analysis to switch to logarithmic measures is inapplicable. Furthermore, in contrast to the stock market prices which are only available for trading days, EPEX prices cover the whole year, including weekends and holidays. Accordingly, there is a lot of underlying variability in the data which has nothing to do with volatility, but simply reflects diurnal activity patterns. An important distinction of the present work is the application of matrix decomposition techniques, namely the singular value decomposition (SVD), for defining an alternative notion of volatility. This approach is systematically more robust toward outliers and also the diurnal patterns. Our observations show that the day-ahead market is becoming less volatile in recent years.

Keywords—*Day-ahead price, Electricity market, Energy switch, Singular value decomposition, Time-series analysis, Renewable energy sources.*

I. INTRODUCTION

Renewable Energy Sources (RES) are assuming an increasingly pre-eminent role in the German electricity production. Significant cost reductions, on the one hand, and tremendous technology advances and reliability improvements, on the other hand, have primed a growing interest in green electricity. Germany is pursuing an ambitious goal, a switch from fossil fuel to renewables, "Energiewende" (energy transition). By 2050, the emission of greenhouse gases is planned to reduce by 80-95% (see [1]). To achieve this, energy consumption is to be reduced by 50% and at least 80% of electricity is to come from renewables. In line with that, Germany has substantially expanded its RES-capacity, in particular wind and solar [2]. Consequently, the need for accurate predictions for the quantity of green electricity which is going to be fed into the grid at any given moment is becoming increasingly important. Moreover, energy transition could put pressure on the energy sector in terms of flexibility in managing power demand and supply. It is therefore to be expected that this evolution will have a significant impact on German electricity prices.

This paper provides an overview of the price volatility of the German day-ahead EPEX market. If price volatility increases, it can cause additional risks for suppliers and consumers on the electricity market. Negatively impact the reliability of the power grid is also possible. Due to the integration of the European grids the problems will not be limited to the German grid. Increasing instability in the German grid means also a higher instability in the neighboring grids. In the present work, we have opted to focus on this market as it represents an important and growing segment where market mechanisms are clearly visible. In particular, we focus on the following question: How can the evolution of the price volatility of the day-ahead market over the past eleven years (i.e., 2006-2016) be quantified?

The rest of this paper is organized as follows: Section II contains background studies and literature review. Section III focuses on the data description for day-ahead market; also explains the source and a brief summary of the day-ahead market mechanism. The evolution of a new type of volatility measure for time series data is discussed in Section IV. The general reduction in the volatility of price data is argued in Section V.

II. BACKGROUND AND LITERATURE REVIEW

Recently, the impact of variable generation on the electricity market has attracted a lot of attention. We briefly highlight a number of important contributions which are related to the topics discussed in this paper. Denny et al. [3] study the functionality of the increased interconnection between Great Britain and Ireland to facilitate the integration of the wind farms into the power system. This work suggests a reduction in average price and its volatility in Ireland to be an outcome of the increased interconnection. With the growing contribution of intermittent energy sources, transmission grid extensions and increasing the cross-border interconnection capacities seem inevitable. Schaber et al. [4] examine the viability of this approach and its effects, based on projected wind and solar data until 2020; they also conclude that expanding the grid is, indeed, helpful in coping with externalities which come with

the deployment of RES. The relation of substantial expansion of photovoltaic (PV) installations in Germany and Italy with daytime peak price fall in these countries is discussed in [5]. In [6] the influence of the RES on the German day-ahead market is studied; the authors also have considered priority that the German government assigns to the green electricity over fossil fuels in case of adequate supply. This paper also argues that the emergence of negative prices in the German day-ahead market is the result of the integration of RES.

In the present paper, our main focus is the evolution of the price volatility. Therefore, a novel approach to quantify the volatility of the German day-ahead market is proposed.

III. DATA

The European Power Exchange (EPEX SPOT SE) operates on the Central Western European (CWE) spot market, i.e. Swiss, French, German and Austrian short-term electricity markets. Striving for the creation of a single integrated electricity market, EPEX SPOT functions as an organized wholesale market place for trading large quantities of electricity between the market members. These members are mostly non-final consumers and big players in the energy sector such as utilities and aggregators, industrial producers, the Transmission System Operators (TSOs), banks, financial service providers and energy trading entities that are working within the energy sector on a daily basis. In fact, this company offers its clients the technologies, electronic trading systems and platform to operate their orders based on reference prices.

a) Day-ahead Auction Spot Market: The day-ahead market is an exchange for short-term electricity contracts. The trading in this market is driven by its participants. A buyer, typically a utility, needs to assess how much energy (MWh) it will need to fulfill its customers requirements for the following day, and how much its purchase price is going to be (Euro/MWh) hour by hour. The seller, for example the owner of a wind farm, also submits the quantity he is prepared to deliver the next day and the price level on hourly basis. The deadline for the members to submit the price and the quantity for which they seek to make an agreement is 12:00 CET. These "bids" are fed into a complex algorithm to calculate the *clearing* price. From 00:00 CET the next day, the sellers deliver the power at the contracted rate.

b) Day-Ahead Spot Prices (in Euro/MWh): This is the price (for each time-slot of the next day) as set by the spot market. As previously mentioned, on the day-ahead market the hourly price of the traded quantity (in Euro/MWh) is set a day earlier. Fig. 1 illustrates an overview of the hourly values of the price on the day-ahead market in Germany and Austria from 2006 until 2016 (data source: [7]).

IV. VOLATILITY OF THE DAY-AHEAD PRICE

A. Introduction

It is reasonable to question whether the intermittency of renewables would render the price more erratic. This can happen due to technical problems in the grid, e.g., congestion; or simply due to poor performance of the day-ahead energy forecasting models. For that, we are interested in the volatility of the day-ahead price. Loosely speaking, volatility
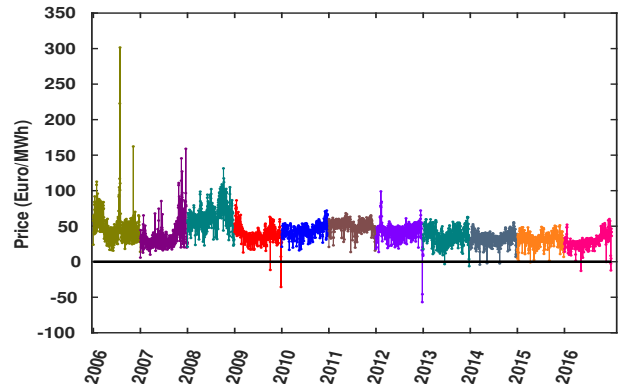

Fig. 1. German day-ahead spot prices from 2006 through 2016 (daily averages); an overall downward trend in the day-ahead price is notable.

refers to the random fluctuations of a time series about its expected value. There are various methods to define and quantify volatility, from applied models like Garman/Klass to coefficient of variation and formal Stochastic Volatility models such as GARCH, Heston models and the like, see, e.g., [8], [9]. However, there are at least two reasons why it is problematic to blindly transfer standard fintech methodology to the current setting:

- Since EPEX prices can be zero or negative, the standard approach to switch to logarithmic measures is not applicable.

- More importantly, while the stock market prices are only available on trading days; EPEX prices cover the whole year, including weekends and holidays. As a consequence, there is a lot of underlying variability in the data which has nothing to do with volatility, but simply reflects the diurnal activity patterns.

For these reasons, we will first pre-proccess the raw data to eliminate the underlying patterns; and subsequently focus on quantifying the volatility of the resulting *residuals*. The nature of this pre-processing is discussed in the following subsection.

B. Extracting underlying trends

Singular value decomposition (SVD) is a popular matrix decomposition technique; which we use to extract the underlying daily and seasonal patterns. SVD is applied extensively in matrix computations, but can also be put to good use in the study of time series which have exogenously induced periods. This is often the case in economics time series, where the variables of interest show daily or annual periods. In the case at hand, it is clear that prices would show daily patterns that might change relatively slowly over the year. To apply SVD, we rewrite the time series as a matrix where each column records the 24 hourly values for a particular day [10]. In this way, a time series representing a typical year is recast as a matrix A of size 24×365. If the results for every day were identical, then all the columns would be identical and the matrix would have a rank equal to one ($rank(A) = 1$). Put differently, the matrix could be written as the product of a 24×1 column and a 1×365 row. In practice, however, the values for subsequent days will tend to be slightly different
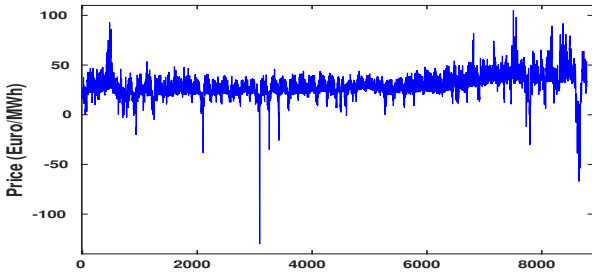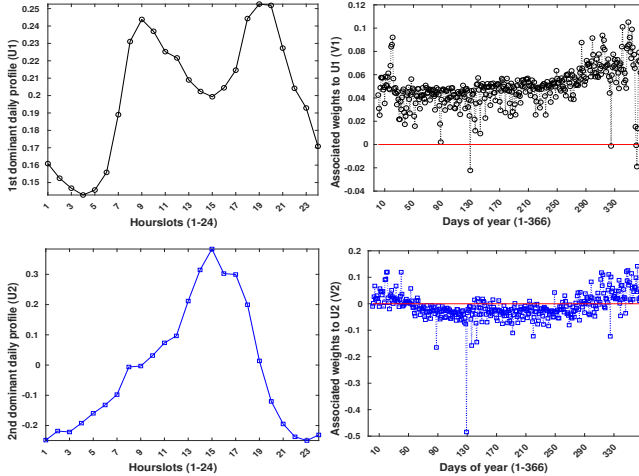
Fig. 2. Raw price data for 2016

Fig. 3. SVD-based rank-2 approximation: Left column: first two columns of U-matrix representing a weighted averaged profile for each day (top) and a first order correction (again one value for each hour over a 24 hour period). Right column: corresponding amplitudes (V columns).
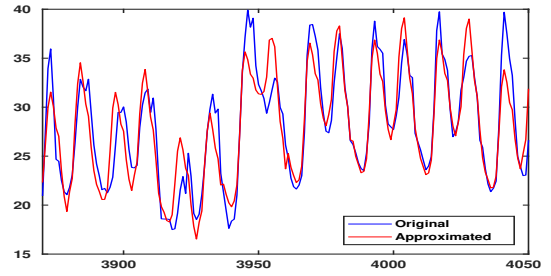

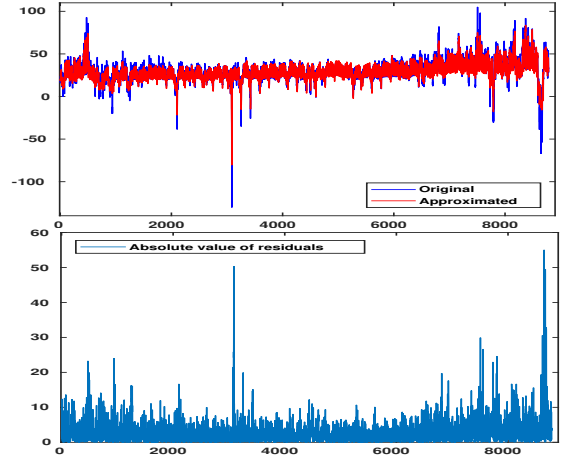Fig. 4. Detail of rank-2 approximation (red) superimposed on original data (blue).

Fig. 5. Top: Original and rank-2 approximation of the price data for 2016. Bottom: Absolute values of the residuals after rank-2 approximation. Residuals as a means to measure volatility.

and therefore A has a rank that exceeds one. Nevertheless, SVD assures us that A can still be expanded as a sum of rank 1 matrices (i.e. a column times a row) where each next contribution in the sum is less important. In mathematical parlance, any given matrix $A_{h \times d}$ can be written as:

$$A = \sum_{k=1}^{r} \sigma_k U_k V_k^T \quad \text{or more succinctly} \quad A = USV^T \quad (1)$$

where $U \in \mathcal{O}(h)$ and $V \in \mathcal{O}(d)$ are matrices with orthonormal columns, with U_k and V_k denoting the k^{th} column of U and V, respectively; and S is an $h \times d$ matrix for which the only strictly positive elements σ_k (so-called singular values) are situated on the main diagonal [11]. The importance of this decomposition lies in the fact that truncating the expansion in the right hand side of (1) after the p^{th} term yields the best approximation of the original matrix A by a matrix of (lower) rank $p \leq r$ [12].

$$A_p = \arg \min_{rank(R)=p} \|A - R\| \quad (2)$$

where the norm $\| \cdot \|$ can be either the Frobenius or spectral L_2 norm. We illustrate the process for the price data of the year 2016. The raw hourly values are shown in Fig. 2. After recasting this time series into a 24×366 matrix (2016 was a leap year!), we can compute the first two terms in the SVD expansion. The results are illustrated in Fig. 3. The left panels depict the first two columns of the U matrix. The first one (top left) resembles an appropriately weighted daily average

(averaged over the year). The price peaks in the morning and late afternoon are clearly visible. The second column of U (depicted in the bottom left panel) provides a correction that needs to be added to the average in order to get a better approximation. The panels on the right-hand side represent the first two columns of the V matrix and specify an amplitude coefficient for every day. To obtain an approximation for the actual data on day d, one must multiply each profile on the left by the d-th amplitude coefficient on the right and of course the corresponding singular value; then add them all up. To get some idea of what a rank-2 approximation looks like, Fig. 4 shows a detail of the approximation (in red) superimposed on the actual data (blue). Fig. 5 contains a more general case of the reconstruction of price data in 2016, along with the absolute value of the corresponding residuals.

C. Quantifying volatility

In the current section, we focus on the residuals of the price after compensating for daily patterns using a rank-2 approximation; it is done to quantify the volatility of the data over the years. The choice to focus on rank-2 approximation is relatively arbitrary and unimportant. Fig. 6 illustrates that the singular values of the price data throughout different years follow the same pattern. Furthermore, as can be seen in the top panel of Fig. 7, the absolute value of the residuals adhere remarkably well to an exponential distribution (with mean 2.97). It is almost only the top 1% that is substantially higher in value than expected. Moreover, the lower panel of Fig. 7
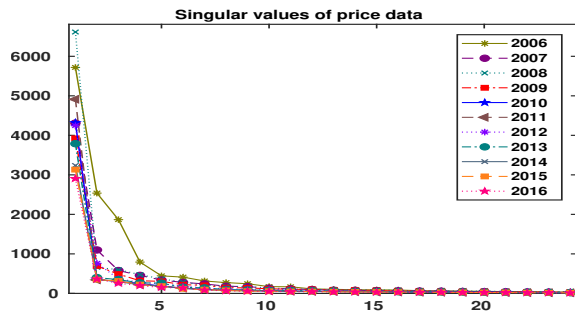
Fig. 6.   An overview of the evolution of the singular value of the price data in each year.
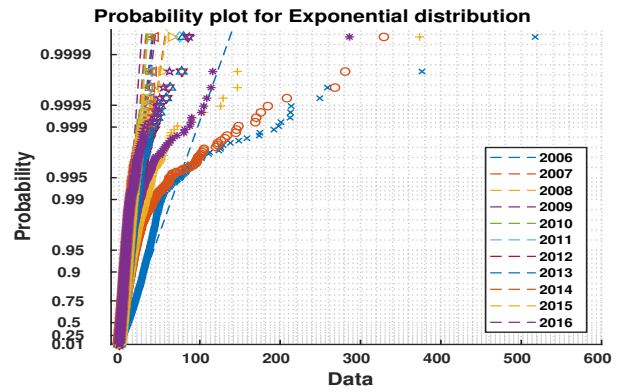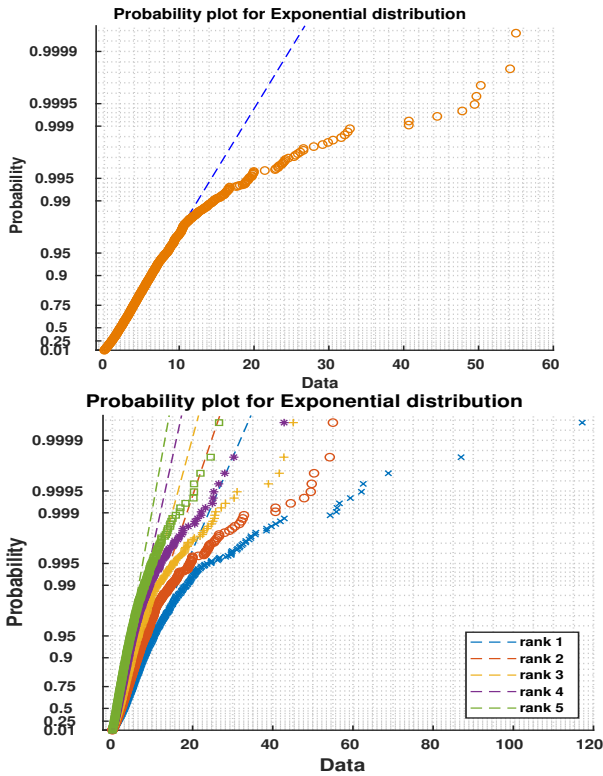

Fig. 7.   Top: The residuals of the rank-2 approximation of the day-ahead market prices for year 2016: almost 99% of residuals adhere to exponential distribution. Bottom: Higher rank approximations yield similar results.

Fig. 8.   Exponential distribution of the residuals of the rank-2 approximations for different years.

Fig. 9.   Evolution of the price volatility on the German day-ahead market in the period 2006 through 2016 . The individual data points record the estimated exponential parameter (mean) based on all but the 1% highest values of the residuals. The regression line has a slope equal to -0.58 which is highly significant (95% confidence interval: -0.89 : -0:26). See main text for more details.

highlights the fact that the higher rank approximations yield similar results. Fig. 8 confirms that this pattern is also seen throughout the period (2006-2016); we focus on in this paper. Based on this observation, we propose the following approach to quantify the evolution in (annual) volatility in the years 2006 through 2016:

1) We remove the influence of daily and seasonal variations by fitting a rank-2 approximation, and extracting the residual of the actual data with respect to this approximation. As volatility is influenced by both positive and negative fluctuations, we focus on the absolute values of these residuals.

2) For every year we fit the lowest 99% of the residual values with an exponential and compute the corresponding parameter (i.e. mean of the exponential).
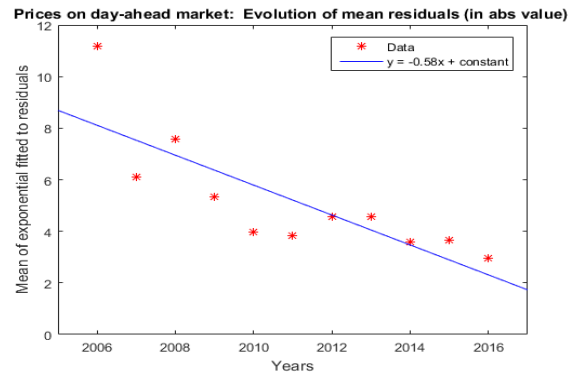
This value corresponds to the size of the residuals.

3) Typically, the top 1% of the observed distribution is much larger than expected (based on the bulk of the distribution). We characterize these values by computing the median value this top 1% segment separately.

The results are shown in Figs. 9 and 10. The former figure shows a robust estimate for the mean of exponential distribution for each year. The estimate is based on the lowest 99% of the residual (absolute) values and is therefore robust with respect to the 1% extremely large values. The 99% vs. 1% is dictated by the exponential prob-plot in Figs. 7 and 8 which show a clear divergence at the 99% mark. To quantify this decreasing trend, we have also computed the regression line, which yields a statistically significant downward slope equal to -0.58 (with 95% confidence interval: -0.89:-0.26). The quality of this regression can be further improved by fitting a *power law*, but at this point we simply want to illustrate the significant downward trend. The evolution of the top 1% is shown in Fig. 10 where these data are represented by their median value. The picture we get from this indicates that whereas extreme residuals were not uncommon prior to 2010, these values fell significantly and have been roughly constant during the last five years. The message from both figures combined is that
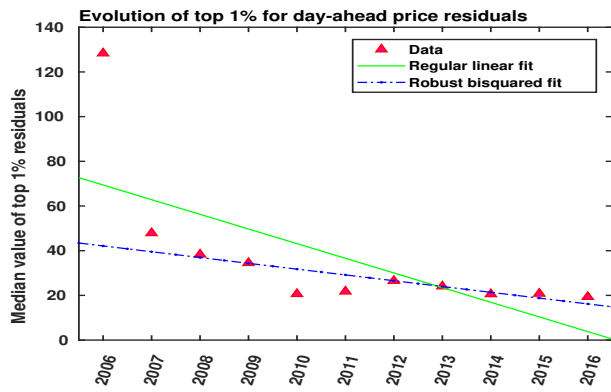
Fig. 10. Evolution of top 1% residuals over the years 2006-2016. Each data point represents the median value of the top one percent residuals (in absolute value). As such, these values characterize the extreme deviations in day ahead prices. The plot clearly shows that these extreme values decreased significantly before 2010, and then stayed approximately constant.

Fig. 11. Price data for 2016: Results of the permutation test. The histogram of the L-values for 1000 random permutations of the actual data. Obviously, the actually observed value (indicated in red) is significantly larger than values we would expect for data sets without structure (with a p-value $< 10^{-3}$). This confirms our observation that volatility shows a seasonal pattern.

volatility has decreased significantly over the past decade.

   a) Volatility tends to be higher in winter: Fig. 5 suggests that volatility tends to be lower in summer (middle part of graph) than winter (extremal parts of the graph). To demonstrate that this is indeed the case we use a measure based on the angular momentum. More precisely, if the residual for hour slot h is given by $R(h)$ and the distance between the hour slot h and the central hour slot $h_m = n/2 = 4392$ equals $|h - h_m|$ the observed angular momentum is defined a:

$$L_{obs} = \sum_{h=1}^{n} R(h)(h - h_m)^2 \qquad (3)$$

where n is equal to the total number of hour slots. If we re-scale the values of the hour slot in such a way that $h_m = 0$ and $-1 \leq h \leq 1$ (divided by 1000, for ease of comparison), we obtain $L_{obs} = 10.676$. A high value for L_{obs} refutes the assumption that the residuals are uniformly distributed throughout the year, and favours an interpretation in which residuals (and hence volatility) is higher in the winter season. To judge the significance of this result, we use a permutation test. The rationale is straightforward: if the residuals are uniformly distributed over the hour slots, then a random permutation of the values should not result in a significantly different value for L_{obs}. The results for 2016 are depicted in Fig. 11. It can be shown that all the other years produced similar results.

V. Conclusions

   The growing share of renewables (especially wind and solar) in the German energy mix appears to result in a higher degree of volatility of the electricity prices; as RES are intermittent and less predictable. Any errors in day-ahead projections for these sources can lead to higher variability in the market. In this paper, we looked at the electricity prices on the German day-ahead market in the period 2006 through 2016. The occurrence of zero or even negative prices makes it difficult to simply use the standard volatility measures which are common in fintech and stock markets. Therefore, we focused on residuals obtained after removing the underlying diurnal or seasonal pa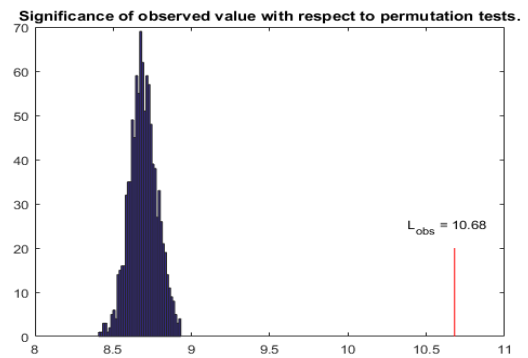tterns. This was accomplished by applying singular value decomposition (SVD) to the time series to get a low rank approximation of the raw data. Our results show that price volatility has significantly decreased during the past eleven years. The most likely cause for this is the improved accuracy of the forecasting tools for renewables in recent years.

References

[1] "Fourth energy transition monitoring report," https://www.bmwi.de/BMWi/Redaktion/PDF/V/vierter-monitoring-bericht-energie-der-zukunft-englische-kurzfassung,property=pdf,bereich=bmwi2012,sprache=de,rwb=true.pdf.

[2] "Information portal renewable energy," http://www.erneuerbare-energien.de/EE/Navigation/DE/Service/.

[3] E. Denny, A. Tuohy, P. Meibom, A. Keane, D. Flynn, A. Mullane, and M. Omalley, "The impact of increased interconnection on electricity systems with large penetrations of wind generation: A case study of ireland and great britain," *Energy Policy*, vol. 38, no. 11, pp. 6946–6954, 2010.

[4] K. Schaber, F. Steinke, and T. Hamacher, "Transmission grid extensions for the integration of variable renewable energies in europe: Who benefits where?" *Energy Policy*, vol. 43, pp. 123–135, 2012.

[5] K. Barnham, K. Knorr, and M. Mazzer, "Benefits of photovoltaic power in supplying national electricity demand," *Energy Policy*, vol. 54, pp. 385–390, 2013.

[6] N. Adaduldah, A. Dorsman, G. J. Franx, and P. Pottuijt, "The influence of renewables on the german day ahead electricity prices," in *Perspectives on Energy Risk*.   Springer, 2014, pp. 165–182.

[7] "Epexspot, day-ahead auction," https://www.epexspot.com/en/product-info/auction/germany-austria.

[8] L. C. G. Rogers and S. E. Satchell, "Estimating variance from high, low and closing prices," *The Annals of Applied Probability*, pp. 504–512, 1991.

[9] R. T. Baillie, C.-F. Chung, and M. A. Tieslau, "Analysing inflation by the fractionally integrated arfima–garch model," *Journal of applied econometrics*, pp. 23–40, 1996.

[10] A. Khoshrou, A. B. Dorsman, and E. J. Pauwels, "Svd-based visualisation and approximation for time series data in smart energy systems," in *2017 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*, Sept 2017, pp. 1–6.

[11] K. Baker, "Singular value decomposition tutorial," *The Ohio State University*, vol. 24, 2005.

[12] G. H. Golub and C. Reinsch, "Singular value decomposition and least squares solutions," *Numerische mathematik*, vol. 14, no. 5, pp. 403–420, 1970.