Use of meteorological data for improved estimation of risk in capacity adequacy studies

Citation for published version:
Wilson, AL, Zachary, S & Dent, CJ 2018, Use of meteorological data for improved estimation of risk in capacity adequacy studies. in *2018 IEEE International Conference on Probabilistic Methods Applied to Power Systems (PMAPS).* Institute of Electrical and Electronics Engineers (IEEE). https://doi.org/10.1109/PMAPS.2018.8440216

Digital Object Identifier (DOI):
10.1109/PMAPS.2018.8440216

Published In:
2018 IEEE International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)

Use of meteorological data for improved estimation of risk in capacity adequacy studies

Amy L. Wilson
School of Mathematics
University of Edinburgh
Edinburgh, UK
Amy.L.Wilson@ed.ac.uk

Stan Zachary
School of Mathematics
University of Edinburgh
Edinburgh, UK
s.zachary@gmail.com

Chris J. Dent
School of Mathematics
University of Edinburgh
Edinburgh, UK
Chris.Dent@ed.ac.uk

***Abstract*—This paper investigates the use of additional meteorological (temperature, wind) data, to obtain more reliable estimates of risk metrics for capacity adequacy in systems with significant wind generation. The advantage of this approach, which requires the careful verification of a number of statistical conditions (well satisfied in practice), is that a considerably longer dataset may be used to estimate the distribution of wind generation, and so also that of demand-net-of-wind, thereby overcoming the otherwise serious problem of a lack of data in the critical region of high demand and low wind. A further advantage is that a block-bootstrapping procedure may be used to assess the remaining uncertainty associated with wind generation, and the effect of that uncertainty on desired risk metrics. This is in contrast to the hindcast approach, where no such uncertainty estimation is possible and where there may be serious concerns about the calculated values of the risk metrics.**

***Index Terms*—Capacity adequacy, reliability**

I. Introduction

The planning of a secure electricity supply requires the use of metrics such as Loss of Load Expectation (LoLE) and Expected Energy Unserved (EEU) to assess shortfall risks. It is important that methods used to estimate these risk metrics are accurate and robust, and that uncertainty in their estimated values is properly considered.

In such analyses the future period (year or peak season) under study is usually divided into n time intervals. Then $\text{LoLE} = nP(Z \leq 0)$ and $\text{EEU} = n\mathbb{E}[\max(-Z, 0)]$ where the "time-collapsed" random variable Z is such that $P(Z \leq z) = \frac{1}{n}\sum_{t=1}^{n} P(Z_t \leq z)$ and where Z_t is the excess of supply over demand in time interval t. In Great Britain (GB), where solar power does not contribute at times of peak demand, we have $Z = X + W - D$ where the similarly time-collapsed random variables X, W and D represent respectively conventional generation, wind generation and demand. (For more details on background see [1] and references therein.)

Conventional generation X is usually and reasonably modelled as probabilistically independent of the pair (D, W) and its distribution uncontroversially evaluated using the availability probabilities of individual generators. This distribution needs to be convoluted with (the negative of) that of the random variable $D - W$ representing demand-net-of-wind. Estimation of the distribution of the latter is more difficult. It is usually made from observed paired (demand, wind speed) traces for historical years, observations being made at hourly or half-hourly intervals. For each such historical year the demand trace is "forward-mapped" (typically by appropriate rescaling) to provide a corresponding demand trace for the future year under study; the wind speed trace—usually of observations at each point on a geographical grid—is similarly forward-mapped to a total wind generation trace for the future year under study by considering the likely locations, capacities and power curves of wind farms in that future year. Thus each year of historical (demand, wind) data yields a corresponding trace of paired "observations" (d_t, w_t), $t = 1, \ldots, n$, of total demand and wind generation for the future year of interest. These paired "observations" may then be used to estimate the predicted distribution of demand-net-of-wind, and hence ultimately the values of risk metrics as described above. Separate estimates of risk metric values for the future year of interest may be made on the basis of each year of historical data; alternatively the historical data may be pooled to obtain overall estimates [2].

The distribution of demand-net-of-wind $D - W$ is therefore to be estimated from one or more time series of n paired "observations" (d_t, w_t) of (D, W). A major difficulty here is that it is only the extreme right tail of this distribution—that corresponding to particularly high demand coupled with the simultaneous occurrence of low or almost no wind—which makes any significant contribution to the left tail of the distribution of the random variable $Z = X + W - D$, and so to the values of risk metrics such as LoLE and EEU. Further it is precisely in this region of high demand and low wind that there is very little data upon which to base the required estimation. The problem is compounded by the fact that, as demand patterns gradually change through time, there are typically relatively few years of historical data (d_t, w_t) which can be used as described above.

The *hindcast* approach simply estimates the tail of the distribution of $D - W$ by the empirical distribution of the observations $d_t - w_t$, of which there are typically very few in the critical region. An alternative [3] is to estimate first the distribution of demand, either by its empirical distribution or via some form of smoothing in its right tail, and to then

This research was supported by EPSRC grants EP/K03832X/1 and EP/N030028/1 and National Grid

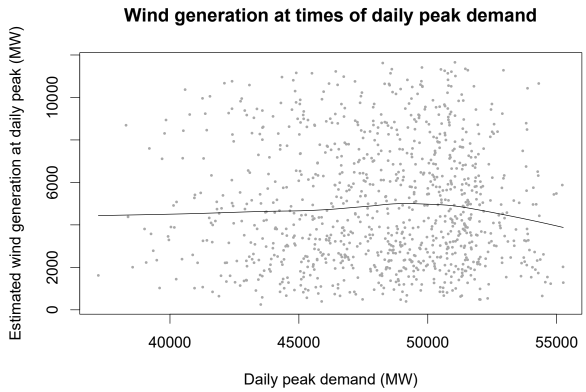

Fig. 1. Plot of wind generation against daily peak demand with smoothed lowess curve showing estimated mean wind generation conditional on demand.

form some reasonable estimate of the distribution of wind generation *conditional* on demand such that this distribution varies smoothly as demand is varied. This allows observations of wind when demand is *not* at its most extreme to contribute also to the estimation of the conditional distribution of wind when demand *is* extreme. The simplest version of such an approach is given by assuming wind generation W to be *independent* of demand D, thereby allowing the distribution of W, including its critical left tail, to be estimated from all the observations w_t of W. However, the independence assumption requires checking and is not always sufficiently well satisfied. Figure 1 gives a scatter plot of wind generation against daily peak demand throughout the peak season over seven years in GB. As can be seen by the smoothed curve overlaid on the plot, there is some evidence that mean wind generation drops at high demands.

This paper explores an alternative approach to overcoming the lack of paired observations (d_t, w_t) of (D, W) in the critical region of high demand and low wind. As before, the distribution of demand D is estimated from the forward-mapped observations d_t of that variable, corresponding to the limited number of years of relevant historical data. However, the observations d_t are now supplemented with concurrent observations te_t of the further meteorological variable *temperature*, denoted by TE. *In the probabilistic arguments below, everything is conditioned on the time t—i.e. the time of year and time of day—at which observation is made.* Given this time, the variables TE and D are generally closely statistically associated. (For example, we shall show this is so in GB.) If this statistical association were perfect (i.e. if, given time as above, TE and D were deterministically related), then the conditional distribution of wind W, given time t and demand D, would be the same as that of wind W, given time t and the corresponding value of temperature TE. This conditional distribution could then be estimated from a meteorological dataset consisting of concurrent observations of temperature TE and wind W obtained over a considerably greater number of years than the necessarily limited number of years used to obtain the observations (d_t, w_t) of (D, W). The result would be a considerable increase in the reliability of the estimated conditional distribution of wind W.

In reality the above statistical association, while typically strong, is not perfect. Thus, given the time t at which observation is made, and the observed value of temperature TE at that time, there is some (typically small) residual variation in demand D. Suppose, however, that there is satisfied the further condition that, again given the time t and corresponding temperature TE, demand D and wind W are approximately independent—something which we again show to be typically the case. Then, in the procedure described above in which at each time t the observation d_t of demand D is supplemented with the concurrent observation te_t of temperature TE, the conditional distribution of wind W given both D and TE is just the conditional distribution of W given TE alone and may again be estimated from a considerably longer meteorological dataset. We may thus proceed as before with the caveat that what is now being estimated is the conditional distribution of wind W given both D and TE. Future values of temperature TE given time t are not precisely known and the effect of this additional conditioning is to introduce a little additional noise into the estimation of the distribution of W given D (which is what is really required), partially offsetting the benefits gained from the use of an extended dataset as above. Provided that, as asserted above, the statistical association between temperature TE and demand D, given time t, is close, then the effect is small. However, it is necessary to check that this is indeed the case.

Thus the method of proceeding is as follows. Historical data consisting of concurrent observations of (demand, wind, temperature) are mapped forward as previously described to produce concurrent "observations" of these variables appropriate to the future year under study. The forward-mapped data are first used to check, given knowledge of the time t at which observation is made, both the closeness of the statistical association of TE and D, and the conditional independence of D and W given TE. Then, conditional on the (D, TE) data derived from any given historical year, an estimate is made of the distribution of $D - W$ as follows: at each time t the forward-mapped "observation" (d_t, te_t) of (D, TE) is supplemented by a sufficiently large number N of simulations drawn from the conditional distribution of W given the observed temp te_t and time t (see Section IV). This conditional distribution is well estimated from an auxiliary *wind generation model*, utilising a considerably longer meteorological time series of concurrent temperature and wind observations. For each such time t, there is thus obtained a set of N simulations of demand-net-of-wind $D - W$ appropriate to the future year under study. Required risk metrics such as LoLE and EEU may now be estimated conditioned on the use of each such historical year of (D, TE) data. Alternatively these may be aggregated over such years to give unconditional estimates.

Further, the use of a block-bootstrapping technique for assessing the uncertainty in the fitting of the wind generation model (see Section V) may be used to assess the corresponding uncertainty in the risk metrics of interest.

II. Data

The "future" year under study is taken to be winter 2014–15. The risk of a shortage of capacity is negligible in summer in GB, so only the peak winter season is considered. This peak season is set to 21 weeks starting from the first Sunday in October. Two datasets described in [4] and [5] are used for the analysis:

- A primary dataset consisting of hourly historical observations of aggregate GB demand for the seven winter seasons from 2007–08 to 2013–2014 inclusive. Observations are available for earlier years but these are not thought to be representative of current demand profiles. These historical observations are scaled to correspond to demand conditions in the 2013–14 winter season. A lowess curve is used to smooth the 90% quantiles of demand for each winter season, giving a scaling factor for each year. Observations in each winter season are then multiplied by the ratio of the scaling factor in 2013–14 to the scaling factor in the given year. This scaling is designed to preserve variation in demand between years due to different weather conditions, but to eliminate any trend in demand through time that might be caused by changes in underlying demand patterns. The fitted lowess curve models the variation in the 90% quantile of demand through time due to changes in these underlying demand patterns. This variation, unlike fluctuations in weather, is expected to be smooth. Rather than smoothing the mean, the 90% quantile was used because only observations of high demand are of relevance for assessing the risk of shortage of capacity. These re-scaled demand observations are referred to in the text as 'forward-mapped' demand. Note that the study year is 2014–15 so an assumption is made that demand conditions in 2013–14 are similar to demand conditions in winter 2014–2015. In practice, this assumption has no impact on the methodology as demand can be scaled to any required level. These forward-mapped demand observations are paired with concurrent historical population-weighted average hourly air temperature measurements for GB.
- A further dataset consisting of hourly "observations" of aggregate GB wind generation, appropriate to the study year and represented by wind capacity factors. These observations are obtained using 23.5 years of historical wind speed measurements (mid 1991–2014), converted to total GB wind generation by aggregating over the locations of the installed wind generation fleet as at January 2015. Wind capacity factors are then obtained by dividing the total GB generation by the total installed capacity. These wind generation observations are paired with concurrent historical population-weighted average hourly air temperature measurements for GB. This "long" dataset is used to estimate the conditional distribution of wind given temperature and time.

III. Investigation of the relationship between demand, wind and temperature

This section uses the forward-mapped historical (demand, wind, temperature) data to investigate whether the required conditions of the Introduction are satisfied, i.e. whether, conditional on time, demand D and temperature TE are reasonably closely statistically associated, and whether, conditional on time and on temperature TE, demand D and wind W are approximately independent. The analysis is based on the primary dataset described in Section II consisting of the available seven years of such historical data.

We use a regression model for the daily peak demand PD in the peak season. The wind generation available on the GB system is such that any shortfall in total generation is likely to occur at (or close to) times of peak demand, so that it is primarily at these times that the statistical assumptions of the present approach require to be tested.

The regression model fitted omits in the first instance any wind term—the effect of adding this term is tested subsequently to check the second of the above conditions. The model includes as explanatory variables the day of the week, a temperature variable TE (defined more precisely below) and Fourier terms [6] to capture the annual variation in peak demand (see [7] for further discussion of these variables). Letting pd_t denote the observed daily peak demand (in MW) for the day t, the model may be written

$$\begin{aligned} pd_t =& \beta_0 + \sum_{j=1}^{2} \left(\beta_1^j \sin\left(\frac{2j\pi t}{365.25}\right) + \beta_2^j \cos\left(\frac{2j\pi t}{365.25}\right) \right) \\ &+ \beta_3 te_t + \beta_4(dow_t) + \epsilon_t, \end{aligned} \quad (1)$$

where dow_t is a categorical variable mapping the day t to the corresponding day of the week, te_t is the observed value of temperature TE at the time of daily peak demand on day t, and ϵ_t is an error term with zero mean and constant variance; the constants $\beta_0, \beta_1^j, \beta_2^j, \beta_3$ and the effects $\beta_4(\cdot)$ corresponding to the levels of the categorical variable dow_t are to be estimated. In each year, 13 days over the Christmas period are omitted from the analysis. As demand is lower than average on holidays, this has no effect on the conclusions. The "observed" temperature te_t on day t is the value, at the time of the daily peak demand, of a smoothed hourly temperature te_h defined in each hour h as

$$\begin{aligned} te_h &= \frac{1}{2}(te_{h-24} + to_h), \text{where} \\ to_h &= \frac{1}{4}(ta_h + ta_{h-1} + ta_{h-2} + ta_{h-3}), \end{aligned}$$

and ta_h is the air temperature in hour h in degrees Celsius. This smoothed temperature variable, frequently employed by National Grid, is designed to smooth the measurements in each hour to reflect the lagged dependence of demand on temperature. For possible alternatives see [7].

Least squares is used to fit the model (1) to the primary dataset of seven years of concurrent demand and temperature data. All terms in the model are found to be necessary,

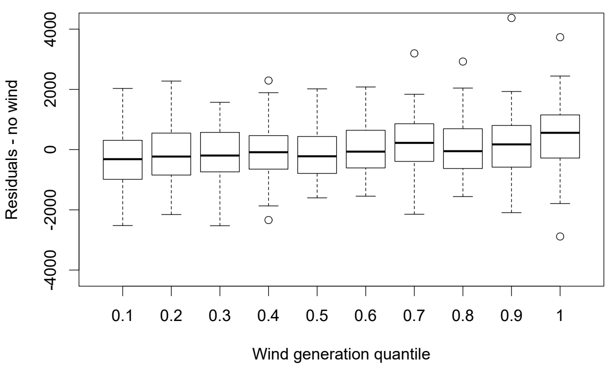

Fig. 2. Boxplot of residuals of model (1) against wind capacity factor w_t split by quantile.

as measured by the coefficient of determination, which for the full model is 0.925. This coefficient of determination is sufficiently close to one as to imply that model (1) is very successful in explaining the variation in peak demand. Thus there is sufficient evidence to support the first condition of the Introduction, that temperature and time capture most of the variation in demand.

To test the second condition of the Introduction, i.e. whether, conditional on time and on temperature TE, demand D and wind W are approximately independent, we now investigate whether the inclusion of wind generation (represented by wind capacity factors) in model (1) provides significant extra explanation of the variation in demand. Let w_t be the observed wind generation capacity factor at the time of peak demand on day t. Figure 2 is a boxplot of the residuals of model (1) against the further observations w_t. The values of w_t are split into ten equally sized bins corresponding the quantiles of their empirical distribution. (For example, the first component of the plot on the left hand side shows the variation in the residuals of the model for the lowest 10% of the observed capacity factors w_t.) As can be seen, any dependence of the residuals of model (1) on the observed capacity factors w_t is weak, and in particular there is no observable relationship in the more critical region corresponding to lower values of wind generation.

As a further test of the conditional independence of demand D and wind W, given time t and temperature TE, a wind capacity factor term $\beta_5 w_t$ may be added to model (1). The addition of this extra term has very little effect on the estimated coefficients $\beta_0, \beta_1^j, \beta_2^j, \beta_3$ and the effects $\beta_4(\cdot)$. While the new coefficient β_5 corresponding to the capacity factors is formally statistically significant under the assumption of independence of the model residuals, nevertheless, given temperature and time, wind generation makes almost no further contribution to the prediction of demand, only increasing the coefficient of determination for the model fit from 0.925 to 0.929.

Thus, from the above analyses, we may conclude that the required conditional independence of demand D and wind W, given time and temperature TE, is also sufficiently well satisfied.

IV. Wind Generation Model

This section develops the model for wind generation (represented by its capacity factor) conditional on temperature and time, where the latter variable includes time of day, time of year and year itself. Wind capacity factors may be converted back to wind generation by multiplying by assumed installed wind capacity for the future year under study. The distribution of future demand-net-of-wind, conditional on demand and temperature data for any given historical year, is then obtained as described in the Introduction. The wind generation model is fitted using concurrent wind and temperature data for 23.5 peak seasons, as described in Section II, and uses data from all hours of the day rather than just the hour of peak demand. This "long" dataset is sufficiently extensive as to permit reasonably reliable inference of the distribution of wind generation (represented by the corresponding wind capacity factor) given temperature and time.

Let the paired (wind capacity factor, temperature) observation for hour t in year i be (w_{it}, te_{it}). Since the w_{it} are constrained to lie between 0 and 1, so as to better fit a suitable model we make the logistic transformation $\tilde{w}_{it} = \log(w_{it}/(1-w_{it}))$ [8]. A regression model is then used for the dependence of transformed wind capacity factor observations on temperature TE and time. For the dependence of wind on temperature a linear term is here sufficient. For the dependence of wind on time it is necessary to include a categorical variable hod_t mapping each time t to the corresponding hour of the day and sufficient Fourier terms to capture slowly varying annual dependence; a year effect y_i for each year i is also required. Thus, the model fitted is given by

$$\begin{aligned}\tilde{w}_{it} &= \beta_0 + \beta_1 te_{it} + \beta_2(hod_t) + y_i \\ &\quad + \sum_{j=1}^{5}\left(\beta_3^j \sin\left(\frac{2j\pi t}{8766}\right) + \beta_4^j \cos\left(\frac{2j\pi t}{8766}\right)\right) + \epsilon_{it}, \quad (2)\end{aligned}$$

where the constants $\beta_0, \beta_1, \beta_3^j, \beta_4^j$, the hour effects $\beta_2(\cdot)$ and the year effects y_i are all to be estimated. For this standard least squares regression proves sufficient. The pattern of annual variation of wind capacity is complex but fairly distinct, and we use a total of five Fourier terms to model the annual variations in capacity factor; the inclusion of further terms runs a distinct risk of overfitting [6].

As measured by the standard assumption of independence of residuals, the inclusion of each of the terms in the model (2) is statistically significant. In particular the estimate of the temperature coefficient β_1 is 1.68×10^{-1} with a standard error of 1.51×10^{-3} and so the dependence on temperature, if modest, is nevertheless highly statistically significant. The effects for the individual years are also required. Their omission from the model reduces the coefficient of determination from 0.178 to 0.157 (recall that there is much residual variation in wind generation not explicable by either temperature or time). As the model is fitted on a logistic scale and is required as a first step in the estimation of the entire distribution of wind generation, as described in detail below, we do not report the numerical values of all the fitted coefficients. Further there is in reality considerable dependence in the sequence of model residuals, and this increases the statistical uncertainty in the model fit. Section V describes how to use block bootstrapping to properly account for this.

Any given historical year generates a forward-mapped (demand, temperature, time) trace for the future year of interest. For each time period t during that year, given the forward-mapped "observation" pair (d_t, te_t), a sufficiently large number N of simulations of wind generation are obtained by the use of the fitted model (2) in which the residuals are randomly sampled from the *empirical* distribution of the entire set of residuals from the fitted model. (Note the need to reverse here the logistic transformation of the wind capacity factors.) The year effect in model (2) is taken to be that appropriate to the historical year on which conditioning is taking place (an alternative would be to treat the year effect as a random effect when fitting the model and to sample from the fitted distribution of that effect). Thus, for the given time period t within the future year under study and conditional on (d_t, te_t), there is obtained by simulation an estimate of the (marginal) distribution of demand-net-of-wind. Hence, conditional on the entire (demand, temperature) trace for the given historical year, point estimates of LoLE and EEU for the future year of interest are obtained as described in the Introduction.

V. Bootstrapping for wind uncertainty

While the above estimates of wind generation are conditional on the (forward-mapped) demand and temperature data for a given historical year, it is important to quantify the uncertainty in these estimates arising from the use of necessarily finite data to fit the wind generation model (2). A considerable difficulty here is that the successive residuals within this model are not independent. Therefore, in order to quantify uncertainty in the coefficients of model (2), we use a block-bootstrapping approach (see [9] or [10] for details). The temporally ordered sequence of residuals from the fitted model (2), which is considered to form a stationary time series, is divided into successive non-overlapping blocks of equal length. A total of B new sequences of residuals, each of the same length as the original, are formed by resampling with replacement from these blocks. Each such *bootstrap* sequence of residuals is then added back to the fitted model (2) to form a new dataset for which the model (2) is *refitted* to give a *bootstrap* estimate of its coefficients. The length of the blocks is chosen sufficiently large that the residuals within distinct blocks are reasonably independent of each other. However, within each block the sequential dependence structure of the residuals is preserved. Thus the empirical distribution of the processes of residuals over a sufficient number B of such bootstrap samples mimics reasonably closely the distribution of the stochastic process of residuals corresponding to the original model (2). Consequently the empirical distribution of the B bootstrap estimates of the coefficients of that model probabilistically quantifies the uncertainty in the original estimated coefficients. Further quantities which depend on the fitted model (2), such as LoLE and EEU, may similarly be calculated for each of the B bootstrap estimates of the coefficients of model (2), so that the empirical distributions of the resulting bootstrap estimates of LoLE, EEU, etc, again probabilistically quantify the uncertainty—due to wind variation—in the original estimates. Probability intervals (confidence intervals) for the estimated quantities are given by the corresponding quantiles of the empirical distribution of the bootstrap estimates. (Note also that, in simulating the conditional distribution of wind, given time and temperature, as described above, we may continue to use random draws from the set of residuals from the original fit of the model (2): the uncertainty in the *distribution* of wind generation is almost completely captured by the uncertainty in the fitted coefficients for the model (2).)

In order to determine an appropriate block length, it is necessary to examine the sequential structure of the residuals $\hat{\epsilon}_{it}$ from the original fit of the model (2). Figure 3 shows both their partial autocorrelation function and autocorrelation function [11]. These show that there is statistically significant correlation out to time lags of around 160 hours, even after accounting for correlation at shorter time lags. A reasonable block length is therefore judged to be one week (or 168 hours).

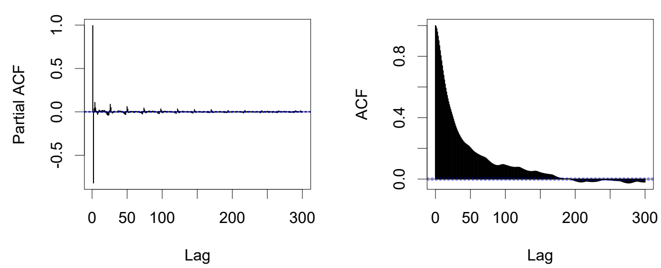

Fig. 3. Partial autocorrelation function (left) and autocorrelation function (right) of residuals $\hat{\epsilon}_{it}$ of model (2).

Thus, separately for each year of historical data (forward-mapped to the future year under study), and conditional on the (demand, temperature) data for that historical year, estimates may be made of LoLE and EEU for the future year of interest as described above, in particular using simulation to estimate the conditional distribution of wind given time and temperature. The fit of the wind generation model (2) describing this conditional distribution is based on the "long" 23.5 year set of (temperature, wind) data and is independent of the chosen historical year of (demand, temperature) data. It follows straightforwardly that pooled estimates of LoLE and EEU based on all seven years of historical data are simply the means of the estimates based on the individual years. In all cases bootstrap estimation may then be used to quantify the uncertainty in estimates. Again the bootstrap procedure relates to uncertainty in the model (2), so that the same bootstrap estimates of the coefficients of this model may be used in every case.

VI. Results

For the given data (as described in Section II) the methodology of the present paper, including the procedure for simulating wind generation described in Section IV, is used to calculate point estimates of the LoLE and EEU for the future year of interest. (For the wind generation model sufficient simulation is used to ensure convergence.) Separate estimates of these risk metrics are obtained based on each of the seven years of historical (demand, temperature) data, forward-mapped as described. A 90% probability interval—accounting for uncertainty arising from use of a finite dataset to estimate wind generation—is estimated for each LoLE and EEU using bootstrapping as described in Section V.

To calculate the risk metrics it is necessary to assume a distribution of conventional generation. We let the random variable X represent the amount of conventional generation available at any point in time and follow [3] in using a two-state model and assuming that X is independent of time, wind generation and demand. Data obtained from National Grid give a future scenario of installed conventional power plant capacities and their corresponding availability probabilities. To protect the sensitive nature of these data, a random error term is added to the installed capacities. The results presented here should therefore be thought of as an exemplar rather than a precise representation of the GB system. A distribution for conventional generation, traditionally represented as a capacity outage probability table [12], is estimated by convoluting the conventional plant capacities and availabilities.

Table I gives estimates of LoLE and EEU conditional on the forward-mapped (demand, temperature) profiles of 2007–08 to 2013–14. Bootstrapped 90% probability intervals are given in parentheses. These are based on 1,000 bootstrap estimates of the coefficients of model (2). As can be seen, there is substantial uncertainty in the risk metrics, highlighting the importance of estimating this wind uncertainty—something which is not possible with the hindcast approach. The variation in LoLE and EEU corresponding to the use of different years of historical (demand, temperature) data is considerable, and can be used to assess the effect that different (demand, temperature) profiles have on the LoLE and EEU estimates.

Hindcast estimates of LoLE and EEU are also shown in Table I. These are generally lower than those obtained using model (2). This difference may reflect variation in the relationship between wind and temperature through time. By using model (2) to simulate wind generation, the relationships between wind, time and temperature are smoothed over 23.5 years of data. In contrast, each hindcast estimate uses only one year of wind data (and hence frequently does not lie within the given 90% probability interval). In particular, for each hindcast estimate, LoLE and EEU are almost entirely determined by a very small number of data points, and it is entirely plausible that in a given year of historical data, there are simply no data points with very high demand and almost no wind, despite the fact that this situation is probabilistically quite possible and, when it occurs, makes a major contribution to the above risk metrics. Thus there is a serious risk that the hindcast approach may significantly underestimate both LoLE and EEU. Further investigation is required here.

Year	LoLE Simulation	LoLE Hindcast	EEU Simulation	EEU Hindcast
07–08	2.99 (2.27,3.83)	2.64	3337 (2474,4337)	2512
08–09	3.13 (2.44,3.91)	1.76	3112 (2403,3940)	1588
09–10	3.80 (3.00,4.71)	1.57	3970 (3083,4980)	1477
10–11	15.91 (13.40,18.65)	9.64	21085 (17413,25189)	12049
11–12	1.18 (0.89,1.49)	0.86	1034 (774,1327)	785
12–13	8.44 (6.59,10.56)	7.84	9578 (7343,12134)	9315
13–14	0.55 (0.39,0.75)	0.31	437 (301,597)	219
Mean	5.14 (4.57,5.78)	3.52	6079 (5315,6938)	3992

TABLE I
LoLE and EEU estimates with 90% probability interval, conditional on the forward-mapped (demand, temperature) trace in the given historical year. Hindcast estimates are given for comparison.

VII. Conclusion

This paper has investigated the use of additional meteorological data, in particular the use of (temperature, wind) data, to obtain more reliable estimates of risk metrics for capacity adequacy. The advantage of this approach, which requires the careful verification of a number of statistical conditions (which appear to be well satisfied in practice), is that a considerably longer dataset may be used to estimate the distribution of wind generation, and so also that of demand-net-of-wind, thereby overcoming the serious problem of a lack of data in the critical region of high demand and low wind. A further advantage is that a block-bootstrapping procedure may be used to assess the remaining uncertainty associated with wind generation, and the effect of that uncertainty on the desired risk metrics. In contrast, for the hindcast approach, no such uncertainty estimation is possible and it is furthermore the case that the hindcast approach may occasionally seriously underestimate the values of these risk metrics.

References

[1] C.J. Dent *et al*, *Capacity Value of Solar Power*, PMAPS 2016.

[2] *National Grid EMR Electricity Capacity Report*, 2017.

[3] C.J. Dent and S. Zachary, *Estimation of Joint Distribution of Demand and Available Renewables for Generation Adequacy Assessment*, Preprint available at http://arxiv.org/abs/1412.1786.

[4] I. Staffell and S. Pfenninger, *Using bias-corrected reanalysis to simulate current and future wind power output*, Energy, 114, 2016, p1224–1239.

[5] I. Staffell and S. Pfenninger, *The increasing impact of weather on electricity supply and demand*, Energy, 145, 2018, p65–78.

[6] A.C. Harvey, *Forecasting, structural time series models and the Kalman filter*, 1990.

[7] R. Hyndman and S. Fan, *Density forecasting for long term peak electricity demand*, IEEE Transactions on Power Systems, 25(2), 2010, p1142–1153.

[8] M.C.M. Troffaes, E. Williams and C.J. Dent, *Data analysis and robust modelling of the impact of renewable generation on long term security of supply and demand*, IEEE PESGM, 2015.

[9] A. Davison and D.V. Hinkley, *Bootstrap methods and their application*, 1997.

[10] D. Politis, *The impact of bootstrap methods on time series analysis*, Statistical Science, 18(2), p219–230.

[11] C. Chatfield, *The analysis of time series: an introduction*, 2016.

[12] R. Billinton and R. N. Allan *Reliability evaluation of large electric power systems*, Springer, 2012.