

# *Using reanalysis data to quantify extreme wind power generation statistics: a 33 year case study in Great Britain*

Article

Published Version

Creative Commons: Attribution 3.0 (CC-BY)

Open Access

Cannon, D.J., Brayshaw, D.J. ORCID: <https://orcid.org/0000-0002-3927-4362>, Methven, J. ORCID: <https://orcid.org/0000-0002-7636-6872>, Coker, P.J. and Lenaghan, D. (2015) Using reanalysis data to quantify extreme wind power generation statistics: a 33 year case study in Great Britain. *Renewable Energy*, 75. pp. 767-778. ISSN 0960-1481 doi: 10.1016/j.renene.2014.10.024 Available at <https://centaur.reading.ac.uk/38448/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1016/j.renene.2014.10.024>

Publisher: Elsevier

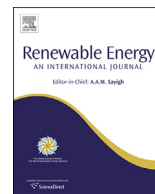
All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

[www.reading.ac.uk/centaur](http://www.reading.ac.uk/centaur)

## **CentAUR**

Central Archive at the University of Reading

Reading's research outputs online



# Using reanalysis data to quantify extreme wind power generation statistics: A 33 year case study in Great Britain



D.J. Cannon <sup>a,\*</sup>, D.J. Brayshaw <sup>a,b</sup>, J. Methven <sup>a</sup>, P.J. Coker <sup>c</sup>, D. Lenaghan <sup>d</sup>

<sup>a</sup> Department of Meteorology, University of Reading, UK

<sup>b</sup> National Centre for Atmospheric Sciences, University of Reading, UK

<sup>c</sup> School of Construction Management and Engineering, University of Reading, UK

<sup>d</sup> National Grid, UK

## ARTICLE INFO

### Article history:

Received 17 March 2014

Accepted 8 October 2014

Available online

### Keywords:

Wind power extremes

Persistent low wind

Wind power ramping

Inter-annual/seasonal variability

Climatology

Wind integration

## ABSTRACT

With a rapidly increasing fraction of electricity generation being sourced from wind, extreme wind power generation events such as prolonged periods of low (or high) generation and ramps in generation, are a growing concern for the efficient and secure operation of national power systems. As extreme events occur infrequently, long and reliable meteorological records are required to accurately estimate their characteristics.

Recent publications have begun to investigate the use of global meteorological “reanalysis” data sets for power system applications, many of which focus on long-term average statistics such as monthly-mean generation. Here we demonstrate that reanalysis data can also be used to estimate the frequency of relatively short-lived extreme events (including ramping on sub-daily time scales). Verification against 328 surface observation stations across the United Kingdom suggests that near-surface wind variability over spatiotemporal scales greater than around 300 km and 6 h can be faithfully reproduced using reanalysis, with no need for costly dynamical downscaling.

A case study is presented in which a state-of-the-art, 33 year reanalysis data set (MERRA, from NASA-GMAO), is used to construct an hourly time series of nationally-aggregated wind power generation in Great Britain (GB), assuming a fixed, modern distribution of wind farms. The resultant generation estimates are highly correlated with recorded data from National Grid in the recent period, both for instantaneous hourly values and for variability over time intervals greater than around 6 h. This 33 year time series is then used to quantify the frequency with which different extreme GB-wide wind power generation events occur, as well as their seasonal and inter-annual variability. Several novel insights into the nature of extreme wind power generation events are described, including (i) that the number of prolonged low or high generation events is well approximated by a Poisson-like random process, and (ii) whilst in general there is large seasonal variability, the magnitude of the most extreme ramps is similar in both summer and winter.

An up-to-date version of the GB case study data as well as the underlying model are freely available for download from our website: <http://www.met.reading.ac.uk/~energymet/data/Cannon2014/>.

© 2014 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/3.0/>).

## 1. Introduction

Due to the increasing market penetration of wind power, extreme wind power generation events (such as prolonged periods of low generation and ramps in generation) are of growing concern to policy makers and transmission system operators. Widespread

low (or high) power generation can persist because wind turbines are insensitive to changes in wind speed when it is low (and turbines produce little or no net power), or high (and turbines produce their rated maximum power). Such persistent events have important implications for electricity system capacity adequacy [1], as well as for wider energy system planning and strategic assessment purposes. In the near future, persistent (multi-day) low generation events will likely influence fuel reserve planning (especially for natural gas), whilst in the longer term, quantifying their frequency and severity will be essential for assessing the potential of innovative technologies such as bulk energy storage [2]. Ramps in

\* Corresponding author.

E-mail addresses: [d.j.cannon@reading.ac.uk](mailto:d.j.cannon@reading.ac.uk) (D.J. Cannon), [d.j.brayshaw@reading.ac.uk](mailto:d.j.brayshaw@reading.ac.uk) (D.J. Brayshaw), [j.methven@reading.ac.uk](mailto:j.methven@reading.ac.uk) (J. Methven), [p.j.coker@reading.ac.uk](mailto:p.j.coker@reading.ac.uk) (P.J. Coker), [david.lenaghan@nationalgrid.com](mailto:david.lenaghan@nationalgrid.com) (D. Lenaghan).

generation often occur at moderate wind speeds where turbine output ranges from zero to a rated maximum power. They also occur at extremely high wind speeds when turbines are shut down for safety, though this is much rarer [3,4]. Such ramps in generation provide challenges for transmission system operators, who schedule reserve holding in advance and require long term strategies for system balancing [5].

Assessing the frequency of extreme generation events directly from power system data is problematic as there is too little data available to determine representative return periods for events that recur infrequently [6]. This is because wind speeds vary on inter-annual and inter-decadal time scales [7,8]. In addition, the geographical distribution of wind farms is constantly changing. In Great Britain (GB), there has been a considerable shift towards wind farms located in the south and offshore. For this reason, weather events that occurred only a few years ago may not have the same impact on the current wind farm distribution as they did before. In response to these challenges, recent studies have estimated the statistical behaviour of the wind resource by inferring the long-term nationally-aggregated wind power output from surface-based wind speed observations. For example [3,9], estimated long-term mean generation statistics for the United Kingdom (UK) and GB respectively, including a brief analysis of low wind periods. Recent studies such as [4,10,11] have also used surface observations to estimate generation statistics.

As an alternative to surface-based observations, authors in academia [12–14], government [1] and industry [15] have begun investigating the potential usefulness of meteorological reanalysis data. Modern “reanalyses” are constructed using global numerical weather prediction models that assimilate observations from a wide variety of sources including land surface stations, buoys, radiosonde balloons, aircraft and satellites [16,17]. Reanalysis data is, by construction, coarsely resolved and so cannot represent small-scale wind fluctuations at a particular site [18]. Nevertheless, as will be shown in Section 2, good agreement with surface-based observations is found when considering variability over sufficiently large spatiotemporal scales.

For assessing wind power variability on a multi-hour, regionally-aggregated scale (as is the focus here), reanalyses may offer numerous advantages over surface-based observations. Firstly, wind observations are heavily influenced by their immediate locale (local topography, vegetation or buildings), and so may not accurately represent the conditions at nearby wind farms. In contrast, because reanalyses do not resolve these local features, they reproduce the large scale wind variability more faithfully. Secondly, changing measuring equipment and recording standards produce biases and discontinuities in the observational record. The impact of these biases on reanalysis data is reduced by the use of multiple observation sources, and by the consistent modelling (and data assimilation) methods used throughout [16,17]. Thirdly, there are few surface-based observations offshore, whereas reanalysis data has global coverage. Finally, modern reanalysis products estimate the wind at multiple vertical levels near the surface using atmospheric boundary layer parameterisations. Whilst still heavily idealised, their consideration of stability effects on the wind profile represents an improvement over the assumption of a neutrally-stratified boundary layer, which is implicit in most studies using surface-based wind observations [4,10,11].

### 1.1. Paper outline

This paper is divided into two main parts (Sections 2 and 3). Section 2 begins by investigating the accuracy with which data from the MERRA reanalysis [16] reproduces the observed variability in near-surface wind speed (Sections 2.1–2.2) and aggregated wind

power generation (Section 2.3) over different spatiotemporal scales. Statistics of long-term mean aggregated wind power and extreme events are then derived and compared to available power system data (Sections 2.4–2.5).

In Section 3, a 33 year climatology of GB-aggregated wind power generation from 1980 to 2012 is used to estimate the frequency of extreme events (persistent low or high generation and ramping), assuming the wind farm distribution of September 2012 (Section 3.1). The inter-annual and seasonal variability of the results is examined (Sections 3.2–3.3), as well as the sensitivity to changes in the assumed dependence of wind farm power generation on wind speed (herein, the “power curve”; Section 3.4).

Conclusions are presented in Section 4, where the potential impacts of the climatology for power system management are discussed.

## 2. Reanalysis verification

### 2.1. 10 m altitude wind speed comparisons

The degree to which wind speeds in MERRA reproduce surface-based, hourly, 10 m altitude UK wind observations from the MIDAS archive [19] will now be evaluated.<sup>1</sup> To facilitate a proper comparison, the gridded MERRA data was bi-linearly interpolated to obtain wind speeds at the co-ordinates of all 328 MIDAS stations. Overall, the MIDAS observations span 1980–2011, though no individual stations were operational for all 32 years.

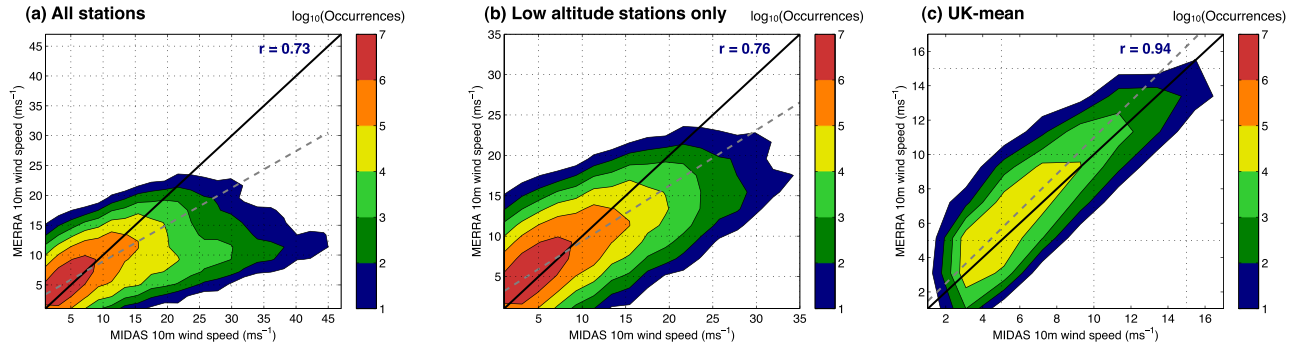
Fig. 1(a) shows a site by site comparison between the 10 m altitude wind speed records in MERRA ( $V$ ) and MIDAS ( $U$ ). As [14] similarly noted, whilst in most cases MERRA accurately reproduces the MIDAS wind speeds (the correlation coefficient is 0.73), there is a small systematic overestimation for around  $U < 6 \text{ ms}^{-1}$  and a large underestimation for around  $U > 20 \text{ ms}^{-1}$ . The worst underestimations are removed when stations above 300 m altitude are discounted (Fig. 1(b)). This is a result of the smoothed topography used in MERRA,<sup>2</sup> which leads to artificially low wind speeds for stations residing on the (unresolved) peaks [20]. The smoothed topography may similarly contribute to the small overestimation in wind speed for some low altitude stations.

Although MERRA cannot fully capture the observed MIDAS wind variability at individual locations, the mean wind speed (spatially averaged over all stations) is reproduced more accurately (Fig. 1(c)). The range of mean wind speeds is smaller than at individual sites, reflecting the reduced influence of extremely high winds which only simultaneously effect a small number of stations. The correlation coefficient between the mean wind speeds in MERRA and MIDAS is greatly increased (to 0.94), which is consistent with the “smoothing” commonly observed when averaging (or aggregating) over large numbers of stations [3,21]. This smoothing reduces the impact of small-scale wind variability, leaving the large-scale variability (well resolved by MERRA) dominant. The improved agreement in mean wind speed implies that MERRA should be considerably more successful in reproducing regionally-aggregated generation than that of an individual wind farm.

To evaluate the degree to which MERRA reproduces the temporal variability observed in MIDAS, the above analysis was repeated for the change in wind speed over different time spans. At individual locations, MERRA tends to underpredict the change in wind speed relative to MIDAS on short time spans ( $\Delta t = 3 \text{ hr}$ ,

<sup>1</sup> The MIDAS wind speed observations are not assimilated into MERRA.

<sup>2</sup> The smoothed topography in MERRA is a result of the coarse (approximately  $50 \text{ km} \times 50 \text{ km}$ ) horizontal grid used in the underlying numerical weather prediction model.



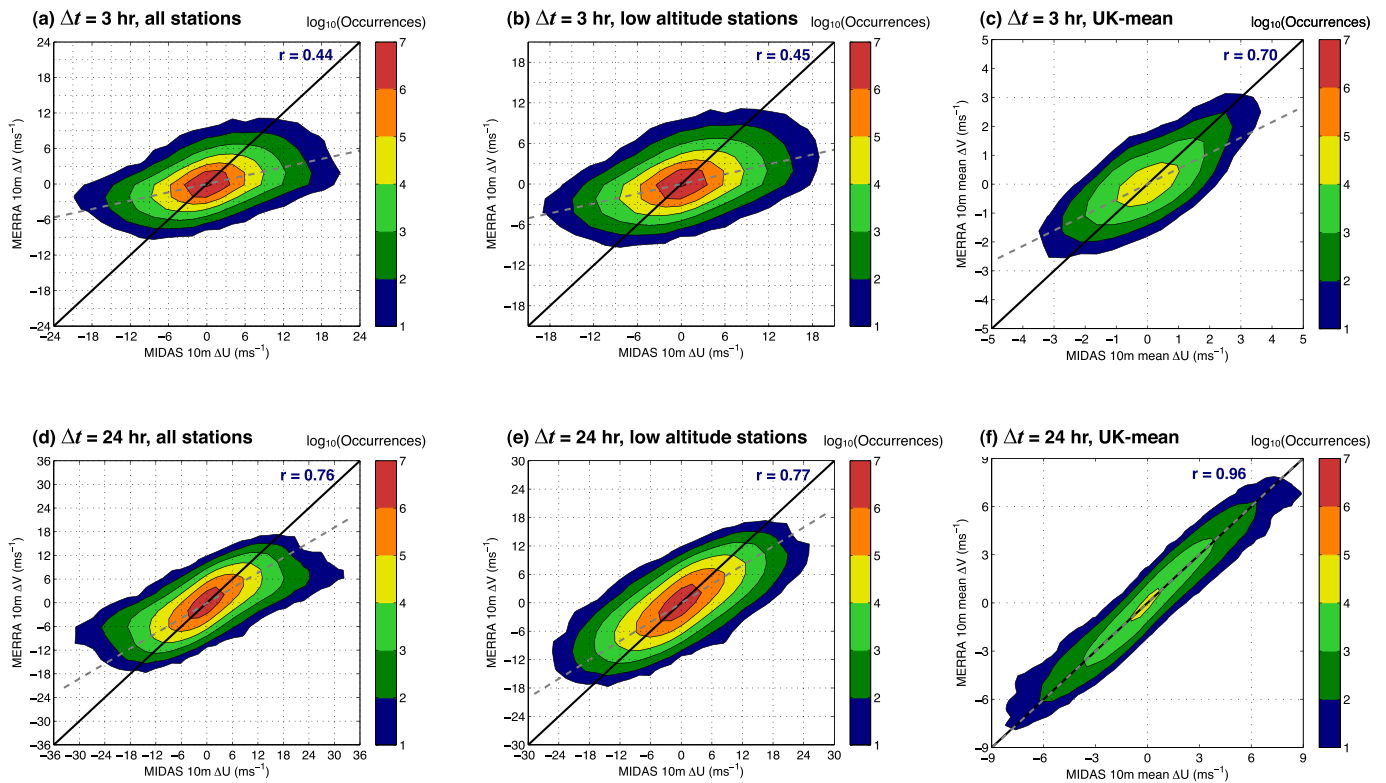
**Fig. 1.** Comparisons between the 10 m altitude wind speeds from MERRA and MIDAS between 1980 and 2011, for (a) all MIDAS station locations, (b) only stations below 300 m altitude, and (c) the mean wind speed over all MIDAS station locations. The number of occurrences within  $2.06 \times 2.06 \text{ m}^2 \text{ s}^{-2}$  bins is shown on a logarithmic scale ( $2.06 \text{ m}^2 \text{ s}^{-2}$  is four times the discretisation of the MIDAS wind speed data). The black solid line indicates a 1:1 agreement, whereas the dashed line shows a linear least squares fit to the data. The linear correlation coefficient is given by  $r$ .

Fig. 2(a)), but is more accurate over longer time spans ( $\Delta t = 24 \text{ hr}$ , Fig. 2(d)). The most extreme changes in wind speed are consistently underestimated, with the largest underestimations associated with high altitude stations (Fig. 2(b, e)). As before, the correlation coefficient increases markedly when considering the spatial mean over all stations (Fig. 2(d, f)).

This analysis shows that MERRA successfully reproduces the observed near-surface wind variability over large spatiotemporal scales, but less accurately reproduces localised wind variability (especially in regions of complex terrain) and changes in wind speed over short time spans. In Section 2.2, the precise spatio-temporal scales over which MERRA reproduces the observed variability are estimated.

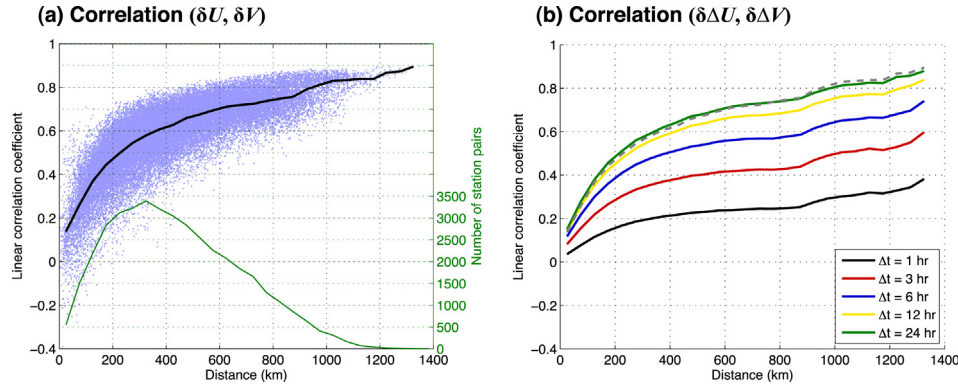
## 2.2. Estimating the spatiotemporal scales over which MERRA reproduces the observed wind variability

To estimate the spatial scales over which MERRA adequately captures the observed wind variability in MIDAS, the difference in wind speed between two stations ( $i$  and  $j$ ) in MERRA ( $\delta V = V_i - V_j$ ) and MIDAS ( $\delta U = U_i - U_j$ ) are compared. In Fig. 3(a), the correlation of  $\delta V$  and  $\delta U$  ( $r(\delta U, \delta V)$ ) is plotted as a function of the distance between the stations. Unsurprisingly, there are no station pairs for which  $\delta V$  and  $\delta U$  agree perfectly (i.e.,  $r(\delta U, \delta V) = 1$ ), however there is a clear improvement as the station separation increases. Taking the median  $r(\delta U, \delta V)$  as a function of distance,  $r(\delta U, \delta V) \rightarrow 0$  as the distance decreases to zero. In this extreme,  $\delta V \rightarrow 0$  as MERRA cannot



**Fig. 2.** Comparisons between the rate of change in 10 m altitude wind speed in MIDAS ( $\Delta U$ ) and MERRA ( $\Delta V$ ) over (a, b, c)  $\Delta t = 3 \text{ hr}$  and (d, e, f)  $\Delta t = 24 \text{ hr}$ . Panels show comparisons for (a, d) all MIDAS station locations, (b, e) stations lower than 300 m altitude, and (c, f) the mean wind speed over all MIDAS station locations. The number of occurrences within (a, b, d, e)  $2.06 \text{ ms}^{-1}$  by  $2.06 \text{ ms}^{-1}$  or (c, f)  $0.515 \text{ ms}^{-1}$  by  $0.515 \text{ ms}^{-1}$  bins are shown on a logarithmic scale ( $0.515 \text{ ms}^{-1}$  is the discretisation of the MIDAS wind speed data). The black solid 1:1 line indicates perfect agreement and the dashed line a linear least squares fit to the data (these lines overlap in (f)). The linear correlation coefficient is given by  $r$ .





**Fig. 3.** The linear correlation between MERRA and MIDAS for (a) the difference in wind speed between station pairs and (b) the difference in accelerations and decelerations in the wind between station pairs, as a function of the station separation. In (a), each dot represents a different station pair and the black line represents the median linear correlation coefficient for each 50 m increase in distance. Also shown is the number of station pairs within each 50 m distance bin (grey in print, green online). In (b), the median lines are shown for different time steps ( $\Delta t$ ). For comparison, the median line from panel (a) is reproduced in (b) as a dashed line. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

resolve the small scale variability affecting  $\delta U$ . As the station separation increases, the large scale atmospheric processes resolved by MERRA become important and  $r(\delta U, \delta V)$  increases rapidly. Although larger spatial scales generally yield higher correlations, the benefit of increasing distance slows markedly beyond around 300 km, where  $r(\delta U, \delta V) \approx 0.5$ . Averaging (or aggregating) over many stations on this spatial scale is therefore likely to produce a high correlation between the MERRA and MIDAS estimates.

To estimate the temporal scales over which MERRA accurately reproduces the observed variability, the above analysis is extended to compare the correlation between  $\delta(\Delta V) = \Delta V_i - \Delta V_j$  and  $\delta(\Delta U) = \Delta U_i - \Delta U_j$ . This tests the ability of MERRA to reproduce the observed spatial variability in accelerations (or decelerations) in wind speed, over varying time spans ( $\Delta t$ ). Fig. 3(b) shows the median dependence of  $r(\delta \Delta U, \delta \Delta V)$  as a function of distance, for varying  $\Delta t$  (for clarity, the individual station pairs are omitted, though they have similar distributions about the median as in Fig. 3(a)). As before, the median  $r(\delta \Delta U, \delta \Delta V) \rightarrow 0$  as the distance tends to zero regardless of  $\Delta t$ . The increase in  $r(\delta \Delta U, \delta \Delta V)$  with distance is however strongly dependent on  $\Delta t$ . Over short time spans,  $r(\delta \Delta U, \delta \Delta V)$  remains small for all station separations, whereas for large time spans, the increase is almost identical to that in Fig. 3(a).

In general, the degree to which small scale variability between stations is smoothed upon averaging or aggregating is dependent on the number of stations as well as their separation. The number of stations beyond which the benefit of extra smoothing is small was estimated at around 50 by Ref. [21], who studied the variability of wind power generation in Germany and Ireland. A similar figure was found here by analysing randomly-selected distributions of stations in MIDAS and MERRA (not shown). The number of MIDAS stations operational at any one time averages around 130 (approximately 40% of the total), and so is considerably higher than 50.

This analysis suggests that care should be taken when interpreting wind variability from MERRA on spatiotemporal scales below around 300 km and 6 h. In the following section, the MERRA data is used to construct a GB-aggregated wind power time series, which is evaluated against National Grid data.

### 2.3. GB-aggregated wind power

In this section, the accuracy with which MERRA can be used to reproduce the measured GB-aggregated hourly wind power from 2012 is determined, and understood in light of the results of Section 2.1. The wind farm distribution shown in Fig. 4(a) is used throughout this paper as it allows both for a comparison with the

2012 National Grid data and provides a contemporary distribution for the climatology presented in Section 3. For each wind farm location, a MERRA-derived power time series was derived by: (i) Bilinearly interpolating the horizontally gridded 2 m, 10 m and 50 m altitude winds to each location, (ii) vertically interpolating the winds to a representative turbine hub height (as estimated by National Grid for each wind farm), assuming a logarithmic change in wind speed with altitude,<sup>3</sup> (iii) applying an idealised power curve (as in Fig. 4(b)) to convert hub-height wind speed to wind farm capacity factor. The GB-aggregated capacity factor,

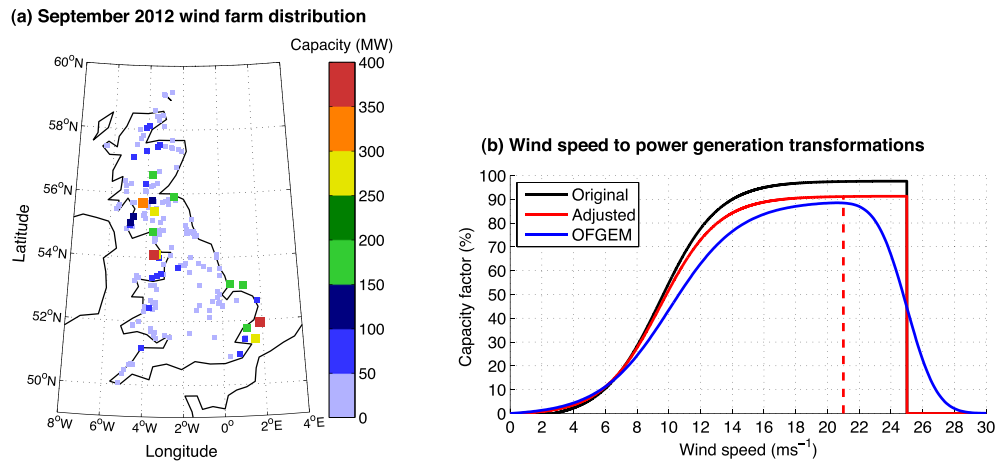
$$CF = \frac{100\%}{C} \sum_{i=1}^{188} c_i g_i(t), \quad (1)$$

is the power generated by each wind farm (the product of the local capacity factor,  $g_i(t)$ , and the wind farm capacity,  $c_i$ ) summed over all 188 wind farms in the distribution (Fig. 4(a)), and normalised by the total GB capacity ( $C = 7.0$  GW). A sensitivity test prior to publication was performed in which the distribution in Fig. 4(a) was replaced with one from April 2014. This showed the capacity factor time series to be only weakly sensitive to modest changes in the wind farm distribution (not shown).

Results will be presented here using both the “original” and “adjusted” power curves shown in Fig. 4(b) (the “OFGEM” curve will be used in Section 3.4). The original curve is based on the design performance of a Siemens 2.3 MW turbine, but has been modified by National Grid to reflect the average dependence of forecasted wind on measured generation (personal communication). The maximum output is (on average) less than 100% due to atmospheric phenomena such as turbine wakes, which decrease the wind speed within wind farms [22], as well as other phenomena such as turbine unavailability [23] and ageing [14]. For simplicity, wind farms are assumed to shut down above  $25 \text{ ms}^{-1}$  and return to full power at  $21 \text{ ms}^{-1}$  (typical values advised by National Grid).

Fig. 5(a) compares the MERRA-derived CF using the original power curve with National Grid generation data from 2012. This includes generation from all wind farms for which National Grid

<sup>3</sup> As the wind data is available at more than one vertical level, there is no need to assume a fixed roughness length in the vertical interpolation.



**Fig. 4.** (a) The wind farm distribution and capacity (colours) of September 2012 (data from National Grid). (b) A range of transformation functions used to convert hub-height wind speed to power output (termed “power curves”). The black curve is based on the design performance of a Siemens 2.3 MW turbine, but is modified to improve agreement between forecast wind speed data and measured power generation. The red curve has been adjusted to correct for small biases in the GB-aggregated power output found using the original curve (Fig. 5). The red dashed line indicates the wind speed at which wind farms come back online after they cut-out at  $25 \text{ ms}^{-1}$ . The blue curve is that assumed in Ref. [1]. The sensitivity of the results in Section 3 to the choice of power curve is discussed in Section 3.4. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

receive metering.<sup>4</sup> To facilitate a proper comparison, instances where wind farms were deliberately curtailed in response to transmission constraints are accounted for by adding the curtailed power back into the generation data.<sup>5</sup> Other human influences, such as turbine maintenance, remain. Even though, unlike the National Grid data, the MERRA-derived estimates assume a constant wind farm distribution that includes many unmetered wind farms, the two time series are highly correlated (with a correlation coefficient of 0.96), albeit with a small overestimation for high values. This high correlation can be understood given the results of Sections 2.1–2.2, which found that MERRA accurately represents wind variability on spatial scales greater than around 300 km (the mean capacity-weighted distance between wind farms is 328 km).

The adjusted curve in Fig. 4(b) is of the same form as the original curve, but is tuned to remove the systematic biases in Fig. 5(a). Fig. 5(b) shows a comparison between the MERRA-derived CF, using this adjusted curve, and that derived from the 2012 measured data. All MERRA-derived results from here on utilise this adjusted power curve.

From Sections 2.1–2.2, we expect MERRA to reproduce changes in CF over time intervals greater than around 6 h. Fig. 6 shows comparisons between the MERRA-derived  $\Delta\text{CF}$  (using the adjusted power curve) and equivalent National Grid values, for a range of  $\Delta t$ . At  $\Delta t = 3 \text{ hr}$ , the time series are reasonably well correlated (with a correlation coefficient,  $r = 0.77$ ), although the largest changes in CF are consistently underestimated. As  $\Delta t$  increases to 6 h and 12 h, the correlation increases ( $r = 0.86$  and  $0.93$  respectively) and the systematic underprediction in  $\Delta\text{CF}$  reduces considerably.

#### 2.4. Long-term mean statistics of GB-aggregated wind power

In this section, the MERRA-derived CF time series described in Section 2.3 is used to analyse the annual-mean CF and the frequency distribution of CF values. Fig. 7(a) shows a comparison between the MERRA-derived annual-mean CF from 1980 to 2012

and recent estimates from National Grid and the UK government (the Digest of UK Energy Statistics, herein DUKES [25]).<sup>6</sup> The MERRA-derived 33 year mean capacity factor is 32.5%; slightly above previous long term estimates ([3] suggested 30%). From the available National Grid and DUKES data, the variability in annual-mean CF is well reproduced by the MERRA-derived time series, including for the low generation year of 2010. The slight reduction in wind speed since the late 1980s is broadly consistent with the “stilling” observed in the UK [4] and more generally in the continental mid-latitudes [26]. This may be a result of inter-decadal variability associated with climate phenomena such as the North Atlantic Oscillation (NAO) which significantly influences European weather [8]. The large year-to-year variability is also correlated with inter-annual fluctuations in the NAO [7,27].

As shown in Fig. 7(b), the frequency distribution of hourly CF values is heavily skewed towards low values, with the most common CF around 4–13 % in 2012 and around 5–6 % for the 33 year MERRA-derived time series. The distribution closely matches that of the National Grid data (to within 15 h per unit CF on average). There are no occurrences above CF > 90 % in either the MERRA-derived or National Grid estimates. The cumulative frequency reveals the 33 year median CF = 26.4 %, which is significantly below the mean (32.5%) due to the positive skew in the frequency distribution. Percentiles from the cumulative distribution will be used in Sections 2.5 and 3 to define persistent low and high wind power events.

#### 2.5. Extreme wind power generation in 2012

As the central purpose of this paper addresses extreme wind power generation events (persistent low or high generation and ramping), we now evaluate the ability of the MERRA-derived power time series to reproduce the extremes of 2012.

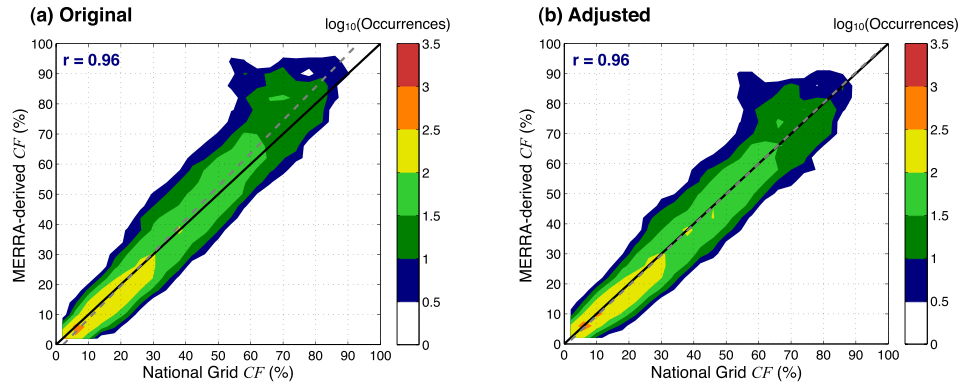
The number of persistent low generation events is presented in Fig. 8(a, d) as a function of both a threshold below which CF drops, and the length of time for which it persists below that threshold.<sup>7</sup> Events that persist beyond the beginning or end of the time series

<sup>4</sup> Typically this includes wind farms with capacity over 100 MW in England and Wales, over 30 MW in southern Scotland (Scottish Power’s transmission area), and over 10 MW in northern Scotland (Scottish Hydro’s transmission area) [24].

<sup>5</sup> On average, less than 0.1% of GB capacity was curtailed in 2012.

<sup>6</sup> The government estimates are for the whole of the UK.

<sup>7</sup> These are sometimes called “load-duration” curves.



**Fig. 5.** Comparisons between MERRA-derived and National Grid estimates of GB-aggregated wind power generation in 2012. The MERRA-derived capacity factor (CF) is calculated using (a) the original power curve, and (b) the adjusted power curve (Fig. 4(a)). The shading indicates the number of occurrences of CF within 4% by 4% bins, and is displayed on a logarithmic scale. The black solid line indicates a 1:1 agreement, whereas the dashed line shows a linear least squares fit to the data (these lines overlap in (b)). The linear correlation coefficient is given by  $r$ .

are immediately terminated. For example, there were 32 events for which  $CF \leq 10.3\%$  for at least 24 h according to the MERRA-derived estimates and 29 according to the National Grid data. Similarly, Fig. 8(b, e) shows the number of persistent high generation events as a function of both a threshold above which CF rises, and the time for which it persists above that threshold. The thresholds used correspond to percentiles of the cumulative frequency distribution in Fig. 7(b). The  $CF = 2.2\%$ ,  $6.3\%$  and  $10.3\%$  thresholds correspond to the 1st, 10th and 20th percentiles, whereas the  $CF = 55.3\%$ ,  $69.6\%$  and  $87.1\%$  thresholds correspond to the 80th, 90th and 99th percentiles.

For both persistent low and high generation events (Fig. 8(a, b)), there is good general agreement between the MERRA-derived and National Grid estimates. In most cases however, the number of short-lasting events is underestimated and the number of long-lasting events is overestimated. This is consistent with the observed underestimation of high frequency variability in the MERRA-derived time series (Fig. 6), which may otherwise break up persistent events into shorter segments. Fig. 8(d, e) shows the same plot but focuses on the rarest (and most extreme) persistent events. The MERRA-derived time series reproduces the most extreme events well in most (but not all) cases.

Fig. 8(c, f) shows the number of hours which preceded a ramp in CF of at least the given threshold magnitude, within different time windows. For example, there were 57 h in 2012 that preceded a ramp of at least  $\Delta CF = 50\%$  within 12 h according to the MERRA-derived estimates, and 55 according to the National Grid data. By definition, any ramp occurring within 12 h of a given hour must also have occurred within any time window greater than 12 h. There is generally good agreement between the MERRA-derived and measured ramps (Fig. 8(c)), albeit with a systematic underestimation of the number of hours preceding modest ramps. As before, this may be due to the lack of high frequency variability in MERRA, which may otherwise add to the maximum  $\Delta CF$ . This is also true for the rarest (and most extreme) ramps (Fig. 8(f)), though the underestimation reduces as the time window increases and the magnitude of high frequency variations becomes small relative to the size of the ramps.

This analysis demonstrates that, whilst imperfect, the frequency with which extreme wind power generation events occur in the MERRA-derived time series closely matches that from the National Grid data.

### 3. A 33 year climatology of extreme wind power generation in Great Britain

In this section, a 1980–2012 climatology of extreme wind power in GB is presented using the hourly time series of MERRA-derived

CF described in Section 2. The mean frequency (the number that occur in an average year) of different extreme events is presented in Section 3.1, after which the inter-annual and seasonal variability is discussed (Sections 3.2–3.3). Finally, the sensitivity of the results to the choice of power curve is analysed in Section 3.4.

#### 3.1. Mean frequency of extreme events

The mean frequency with which persistent low CF events occur is shown in Fig. 9(a) as a function of both the threshold below which CF drops and the time for which it persists below that threshold. The frequency reduces as the CF threshold is decreased or when the persistence time increases, as both provide a more stringent test for what constitutes a persistent low CF event. For example, there are on average 5.6 events per year where  $CF \leq 5\%$  for at least 24 h. Similarly, Fig. 9(b) shows the mean frequency with which persistent high CF events occur as a function of both the threshold above which CF rises and the time for which it persists above that threshold. In this case, the frequency reduces as the threshold CF is increased or when the persistence time increases. The dashed lines in Fig. 9(a, b) indicate the most persistent events in the 33 year time series, for each threshold CF.

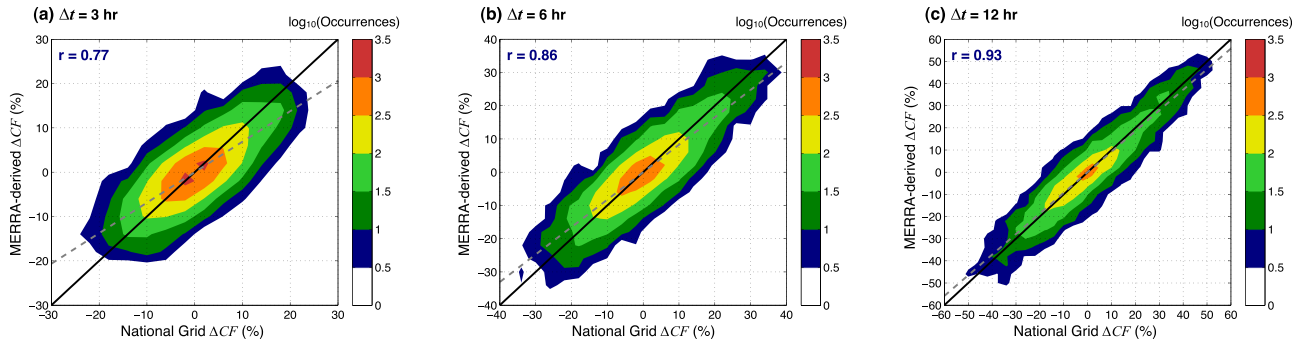
The mean frequency with which low or high generation events occur decreases approximately exponentially with increasing persistence, suggesting they can be approximated as a Poisson-like process where the mean frequency,

$$N = N_0 \exp\left(\frac{-t_p}{\lambda}\right), \quad (2)$$

where  $t_p$  is the persistence time,  $N_0$  is the mean frequency of events of any length (with  $t_p \geq 0$ ) and  $\lambda$  controls the rate at which  $N$  decreases with increasing  $t_p$ . Fig. 10(a, b) shows the mean frequency of persistent low and high generation events on a logarithmic scale, for a range of CF thresholds.

Both  $\lambda$  and  $N_0$  vary as a function of the threshold CF. To illustrate this, Fig. 10(c) shows the variation of  $\lambda$  with CF threshold. For all thresholds,  $\lambda$  was calculated via a linear regression of  $\log[N(t_p)]$ , for all points with  $N > 1 \text{ yr}^{-1}$  (so each  $N$  is based on more than 33 events). To properly compare the persistence of low and high generation events, the rate parameter is plotted not against the threshold CF itself, but against the corresponding percentile of the cumulative distribution in Fig. 7(b), from the most extreme percentile to the least. For the 20 most extreme percentiles, the values of  $\lambda$  are very similar for both low and high wind power generation events. For less extreme percentiles,  $\lambda$  is smaller for high





**Fig. 6.** Comparisons between MERRA-derived and National Grid estimates of the rate of change in GB-aggregated wind power generation in 2012. Changes over (a)  $\Delta t = 3$  hr, (b)  $\Delta t = 6$  hr and (c)  $\Delta t = 12$  hr are shown. The adjusted power curve is used in all panels. The shading indicates the number of occurrences of  $\Delta CF$  within 4% by 4% bins, and are shown on a logarithmic scale. The black solid line indicates a 1:1 agreement, whereas the dashed line shows a linear least squares fit to the data. The linear correlation coefficient is given by  $r$ .

generation events than for low generation events, implying that less extreme low generation events tend to persist longer than less extreme high generation events. This may be a consequence of atmospheric blocking, which is associated with low winds and can persist for weeks [28]. As shown by the alternative axes in Fig. 10(c), the percentiles of extremeness correspond to very different ranges of threshold CF for low and high generation events. This is a consequence of the CF frequency distribution being heavily skewed towards low values (Fig. 7(b)).

In Fig. 9(c), the mean frequency of hours for which there is a subsequent ramp in CF is shown as a function of a threshold  $\Delta CF$ , which the ramp surpasses, and the time window within which the ramp took place. Ramps become rarer as the threshold  $\Delta CF$  increases or as the time window decreases, as both modifications provide a more stringent test for what constitutes a ramp. The most extreme  $\Delta CF$  increases rapidly with the time window up to around 9–12 h, after which it plateaus. This corresponds to the transition time of a typical low pressure (cyclonic) weather system over the UK. As the time window increases to very large values (not shown), the most extreme ramp tends to the maximum permitted by the power curve ( $\Delta CF = 91.3\%$ ). Given the variability in CF over short time spans is likely to be underestimated (Section 2), the statistics for time windows less than around 6 h should be treated with caution.

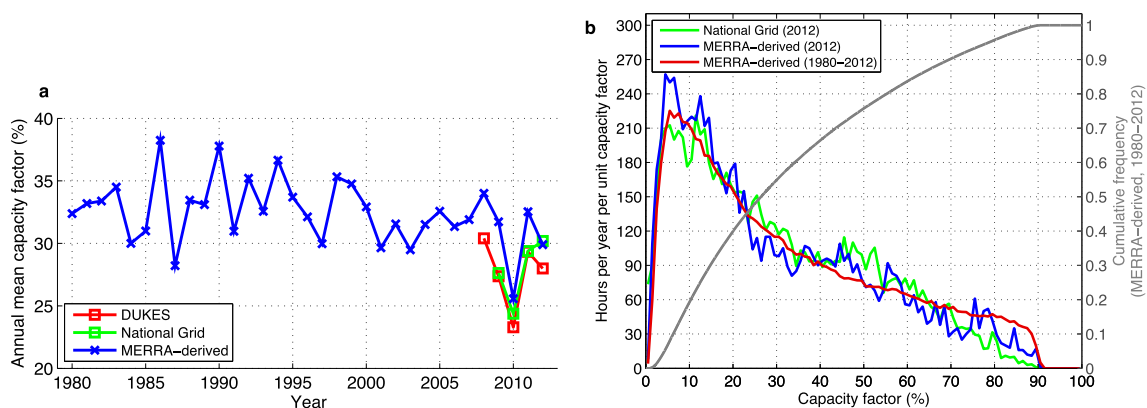
Unlike persistence events, ramps do not have beginning and end points defined by specific thresholds, and so are not counted independently. For example, a ramp may be counted multiple times if it corresponds to the largest  $\Delta CF$  within a given time window for more than 1 h in the time series. The number of hours for which

there is a subsequent ramp of at least a given threshold  $\Delta CF$  does not therefore decrease exponentially with increasing time window (not shown). For this reason, the analysis shown in Fig. 10 is not repeated for ramping events.

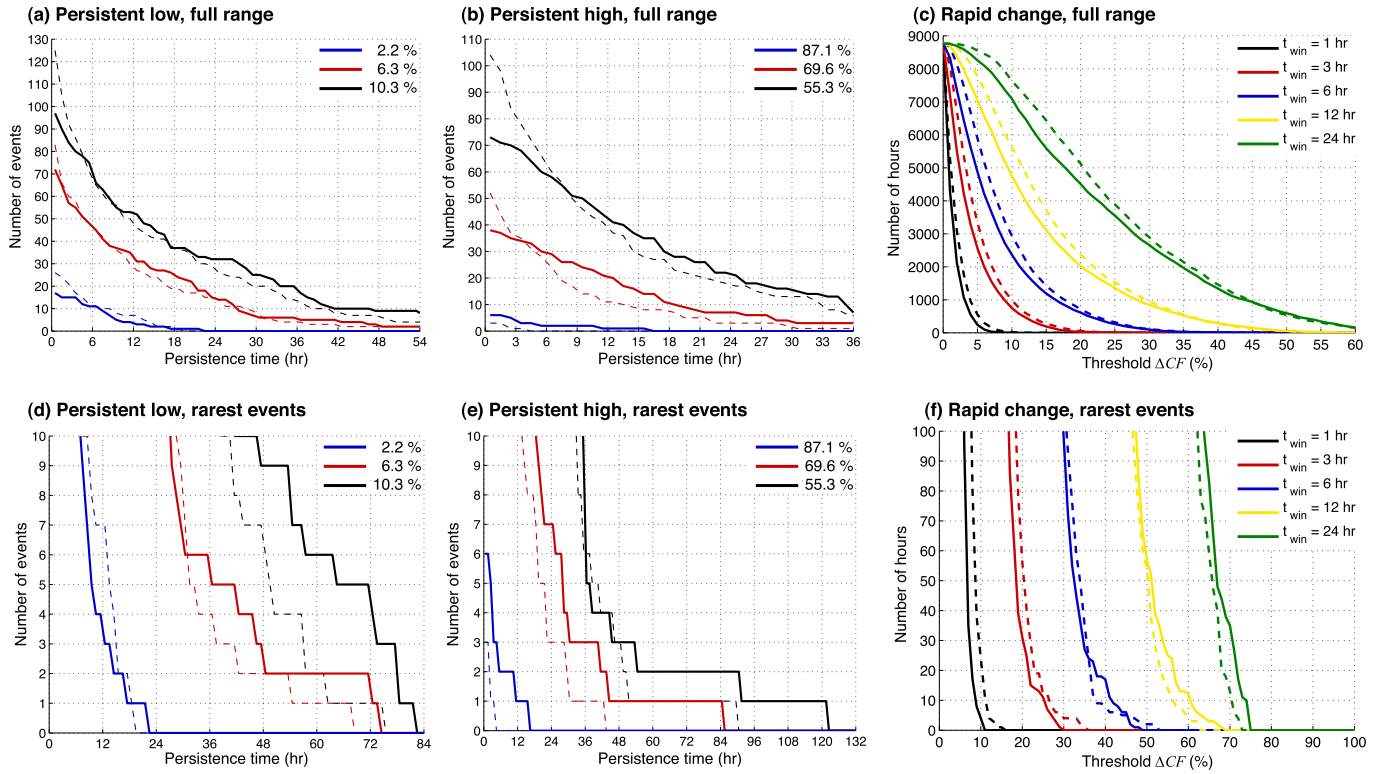
As a sensitivity test, the exclusion of high wind cut-out events was found to make little difference to the mean frequency (not shown). This is likely because they tend to be geographically isolated, and so have a small impact on GB-aggregated generation. In addition, similar results were found by analysing positive and negative ramps in isolation. Whilst similar qualitative trends were observed on smaller regional scales, the mean frequency of extreme events increased markedly as the smoothing effect of aggregation was reduced (not shown). Some differences between the regions of GB were noted, with a propensity for fewer low generation events, more high generation events and more ramps in more northerly regions.

### 3.2. Inter-annual variability

The results of Section 3.1 vary substantially from year to year. Fig. 11(a, b, d, e) shows the frequency of low and high generation events as a function of persistence time, for the CF thresholds introduced in Section 2.5. The frequency in a mean year  $\pm 1$  standard deviation is shown, as well as the highest and lowest number found in any one year. When the persistence time tends to zero, the number of events tends to the mean number of low CF events. As the persistence time increases, the number of events that persist at least that long reduces.



**Fig. 7.** A comparison of (a) different estimates of annual-mean capacity factor and (b) the frequency distribution of capacity factor values within the MERRA-derived and National Grid data. The cumulative frequency distribution obtained using the full (1980–2012) time series is shown in grey. Percentiles from this curve are used to define the thresholds for low and high generation events.



**Fig. 8.** Comparison between the MERRA-derived number of extreme generation events in 2012 (solid), and the number according to National Grid data (dashed). Shown are the number of (a) persistent low generation events and (b) persistent high generation events for three different CF thresholds (described in Section 2.4), and (c) the number of hours preceding a ramp in generation within the indicated time window ( $t_{win}$ ). Panels (d–f) are as in (a–c) but show only the rarest events.

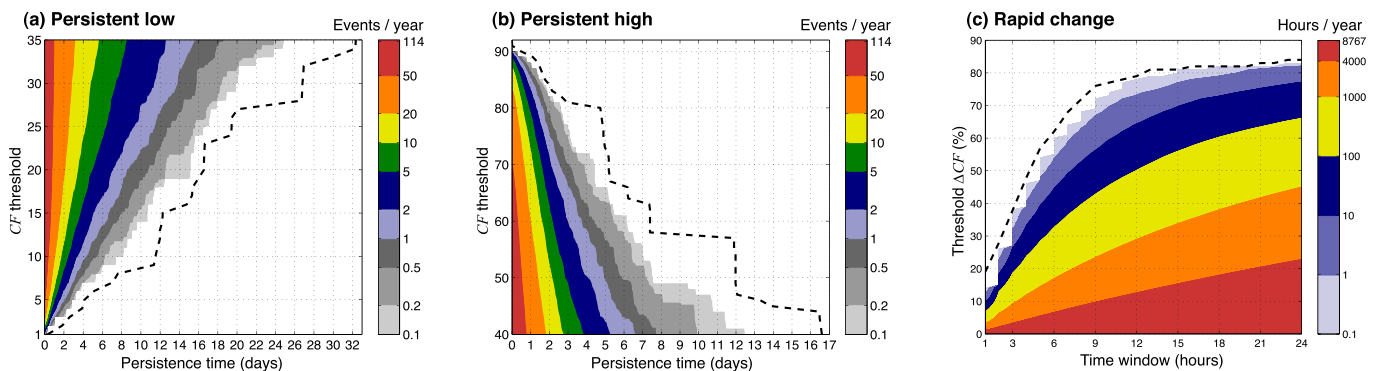
In Fig. 11(c, f), the mean number of hours preceding a ramp surpassing the  $\Delta CF$  threshold is shown, for three different time windows. As in the other panels, the frequency in a mean year  $\pm 1$  standard deviation is shown, as are the highest and lowest number that occurred in any one year. The number of hours tends to the mean number of hours per year (8767) as the threshold  $\Delta CF$  is reduced to zero, and reduces as the threshold  $\Delta CF$  increases.

For all types of extreme generation (low, high and ramping), there is large inter-annual variability. This is especially true for the most extreme events, for which some examples are presented in Table 1. For many extremes, the difference between the most and least active year exceeds the mean frequency.

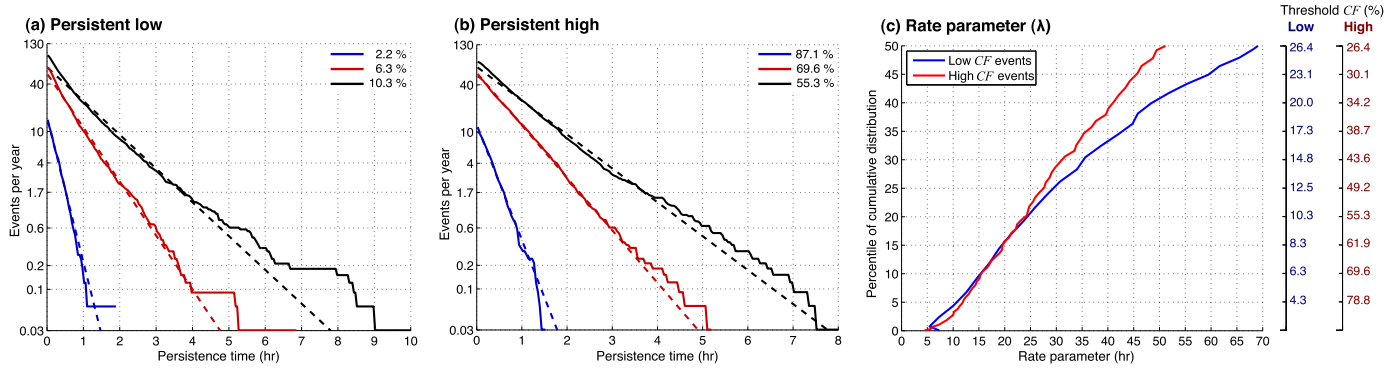
### 3.3. Seasonal variability

In addition to inter-annual variability, there is substantial seasonal variability in the frequency of extremes. Examples of specific event types are given in Table 2. As for inter-annual variability (Section 3.2), the range in the mean frequency from summer to winter can be larger than the frequency in a mean season (calculated as one quarter of the mean frequency).

In Fig. 12(a, d), the mean seasonal frequency of persistent low generation events is shown, as a function of persistence, for the  $CF \leq 6.3\%$  threshold. There is a clear propensity for both a greater number in summer than winter (with Spring and Autumn close to



**Fig. 9.** The mean frequency of extreme generation events derived from the MERRA reanalysis (1980–2012). (a) The frequency of persistent low generation events is expressed as a function of the threshold below which the capacity factor (CF) remains for at least the given persistence time. (b) The frequency of persistent high generation events is expressed as a function of the threshold above which CF remains for at least the given persistence time. (c) The frequency of hours for which there is a subsequent ramp in generation of at least  $\Delta CF$  within the given time window. As discussed in Section 2, the variability in CF over time windows less than around 6 h is likely to be underestimated. The dashed lines mark the most extreme events in the 33 year time series.



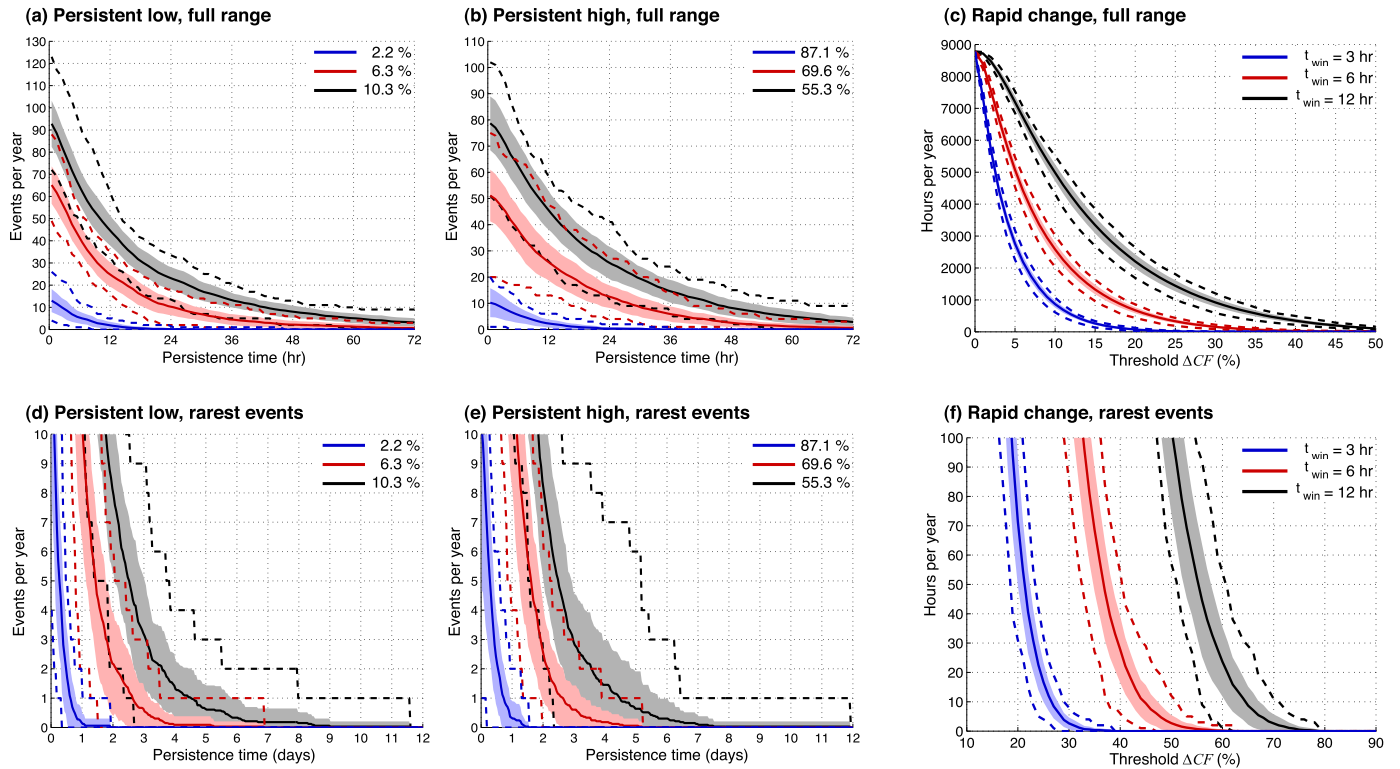
**Fig. 10.** The mean frequency of (a) low and (b) high wind power generation events (solid), for different CF thresholds. Also shown is a linear regression of  $\log[N(t_p)]$  (Eq. (2), dashed), which was fitted using all events for which  $N > 1 \text{ yr}^{-1}$ . (c) The rate parameter for low (blue) and high (red) generation events, which is a function of threshold CF. To compare the low and high generation events, the rate parameter is plotted as a function of a percentile representing the extremeness of the threshold CF. These are derived from the cumulative frequency distribution of CF values in Fig. 7(b). The threshold CF values corresponding to these percentiles are shown on the alternative y-axes (right). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

average) and a greater number of more persistent events. This occurs because the summer months are generally associated with lighter winds [29]. In addition, the longest lasting event with  $\text{CF} \leq 6.3\%$  in summer lasted 6.9 days, whereas in winter it lasted 3.5 days. This is consistent with known seasonal trends in the jet stream, which is often weaker in summer [30].

In Fig. 12(b, e), the mean seasonal frequency of persistent high generation events is shown as a function of persistence, for the  $\text{CF} \geq 69.6\%$  threshold. Mirroring the results for low generation

events, there are both a greater number of high CF events in winter than summer, and a greater number of more persistent events. The longest lasting event with  $\text{CF} \geq 69.6\%$  in summer lasted 1.6 days, but in winter lasted 5.2 days.

The mean frequency of ramps also varies seasonally. Fig. 12(c, f) shows, using a time window of  $t_{\text{win}} = 12 \text{ hr}$ , many more extreme ramps in winter than in summer. This is likely due to the increase in the number of cyclones impinging on the UK in winter, and thus goes hand in hand with an increase in the frequency of high



**Fig. 11.** The mean frequency of extreme generation events derived from the MERRA reanalysis (1980–2012). The frequency of (a) persistent low generation events and (b) persistent high generation events are expressed as a function of their persistence for three low CF thresholds that correspond to the 1st, 10th and 20th percentiles of the cumulative frequency distribution in Fig. 4(b). (c) The mean number of hours for which there is a subsequent ramp of at least  $\Delta\text{CF}$  within different time windows ( $t_{\text{win}}$ ). Panels (d–f) are as in (a–c) but show only the rarest events. All panels show the mean number (solid line) plus or minus one standard deviation (shaded), as well as the minimum and maximum numbers for any one year (dashed).

**Table 1**

The frequency of different extreme events, as derived from the MERRA reanalysis. For each event type, some example thresholds are shown alongside the corresponding mean frequency ( $\pm 1$  standard deviation). The minimum and maximum yearly totals are also shown.

Event type	Thresholds	Mean year	Min. year	Max. year
Persistent low	$CF \leq 6.3\%$ $t_p \geq 24$ hr	$10 \pm 3$ yr <sup>-1</sup>	2 yr <sup>-1</sup>	18 yr <sup>-1</sup>
Persistent high	$CF \geq 69.6\%$ $t_p \geq 24$ hr	$12 \pm 4$ yr <sup>-1</sup>	4 yr <sup>-1</sup>	27 yr <sup>-1</sup>
Extreme ramp	$\Delta CF \geq 50\%$ $t_{win} = 12$ hr	$103 \pm 29$ hr yr <sup>-1</sup>	57 hr yr <sup>-1</sup>	161 hr yr <sup>-1</sup>

generation events. Interestingly, there is little seasonal variability in the most extreme  $\Delta CF$ ; the largest ramp is  $\Delta CF = 74\%$  in summer and  $\Delta CF = 79\%$  in winter.

### 3.4. Sensitivity to changes in the power curve

As the underlying wind speeds do not change, using fixed CF or  $\Delta CF$  thresholds with different power curves modifies the percentile of extremeness to which the thresholds correspond. The results presented in this paper are thus sensitive to changes in the power curve. If these thresholds are instead set according to constant percentiles of extremeness (which vary along with the power curve), then the resultant frequency of extreme events is insensitive to changes in the power curve (not shown).

To illustrate the sensitivity of the results to changes in the power curve when constant CF or  $\Delta CF$  thresholds are used, Fig. 13 shows the mean frequency of the rarest extremes using the three wind farm power curves in Fig. 4(b). Whilst the frequency of low generation events is largely insensitive to the choice of power curve at the  $CF \leq 6.3\%$  or  $CF \leq 10.3\%$  thresholds (not shown), it is sensitive when  $CF \leq 2.2\%$  (Fig. 13(a)). Whilst  $CF \leq 2.2\%$  for at least 12 h around 2 yr<sup>-1</sup> using the adjusted curve, this is increased to 7 yr<sup>-1</sup> using the original curve and only 0.5 yr<sup>-1</sup> using the OFGEM curve. This is due to the subtle difference in which each curve begins generating. Whilst the OFGEM curve generates over 2.2% of capacity at a wind speed of just 2.8 ms<sup>-1</sup>, the adjusted curve requires at least 3.2 ms<sup>-1</sup>, and the original curve requires at least 4.0 ms<sup>-1</sup>.

The mean frequency of persistent high generation events is slightly sensitive to changes in the power curve at the  $CF \geq 55.3\%$  or  $CF \geq 69.6\%$  thresholds (not shown), and this sensitivity increases for the  $CF \geq 87.1\%$  threshold (Fig. 13(b)). Whilst  $CF \geq 87.1\%$  for at least 12 h around 2 yr<sup>-1</sup> using the adjusted curve, this is increased to 10 yr<sup>-1</sup> using the original curve and does not occur at all using the OFGEM curve. This sensitivity arises because the rated maximum CF for the OFGEM curve is only 88.5%, whereas the adjusted and original curves reach 91.3% and 97.6% respectively.

The mean frequency of ramps is also sensitive to changes in the power curve. The number of hours for which there is a subsequent  $\Delta CF \geq 50\%$  within 12 h is around 103 yr<sup>-1</sup> using the adjusted curve, but 177 yr<sup>-1</sup> using the original curve and just 49 yr<sup>-1</sup> using the

OFGEM curve. This sensitivity is due to the difference in slope of the power curves (Fig. 4(b)). The same change in wind speed can result in a larger ramp using the original curve, and a smaller ramp using the OFGEM curve.

These results demonstrate that whilst the statistics of extreme events are insensitive to the choice of power curve if the thresholds used to define the events correspond to the same climatological percentile of extremeness. However, for many practical applications, the thresholds are defined using a constant CF (or  $\Delta CF$ ) threshold. In such circumstances, whilst the general trends reported here are robust to changes in the power curve, the quantitative values can change markedly.

## 4. Conclusions

This paper examines the ability of a state-of-the-art global reanalysis data set (MERRA [16]) to accurately reproduce extreme wind power generation statistics, including for (i) persistent low generation, (ii) persistent high generation, and (iii) ramps in generation on sub-daily time scales. After extensive verification against 10 m altitude wind speed observations and measured nationally-aggregated generation (Section 2), a 33 year climatology of extreme wind power generation events is derived, assuming a fixed, modern wind farm distribution from Great Britain (GB; Section 3). An up-to-date version of the GB case study data as well as the underlying model is freely available for download at <http://www.met.reading.ac.uk/~energymet/data/Cannon2014/>.

MERRA is a coarse global atmospheric reanalysis (the horizontal grid size is around 50 km by 50 km) and is found to poorly reconstruct observed hourly variations in near surface wind speed at individual geographical locations. Nevertheless, it successfully captures the gross patterns of near surface wind variability at spatiotemporal scales greater than around 300 km and 6 h. To investigate wind power generation statistics, an hourly GB-aggregated time series is constructed by (i) spatially interpolating the MERRA wind speeds to the wind farm locations, (ii) extrapolating vertically assuming a logarithmic change between the available vertical levels to typical turbine hub heights, (iii) applying a simple transformation from wind speed to wind farm power generation, and (iv) aggregating over all (188) wind farms. The resultant hourly generation estimates are found to be highly correlated with GB-aggregated National Grid data for 2012, with a correlation coefficient of 0.96. This degree of correlation is similar to that obtained comparing longer-term averages nationally-aggregated generation (e.g., monthly averages [14]). The temporal variability is also well reproduced on time scales greater than around 6 h.

The frequency and severity of extreme generation events observed in 2012 is found to be well reproduced by the MERRA-derived time series. As such, it can be used to derive multi-decadal climatologies of extreme wind power production, assuming a modern wind farm distribution. As reanalysis data has global coverage, the GB case study presented here could be repeated for any distribution of wind farms (past, present or future), anywhere in the world. At 33 years, the MERRA-derived climatology for GB is considerably longer than direct generation records, which extend back only around 5–10 years and suffer from large inhomogeneities due to the rapidly changing wind farm distribution. This approach also avoids many known issues with 10 m wind mast observations, such as their sensitivity to local topographic effects and sparse offshore availability. It also avoids the computational expense of dynamical downscaling using high resolution meteorological models [31].

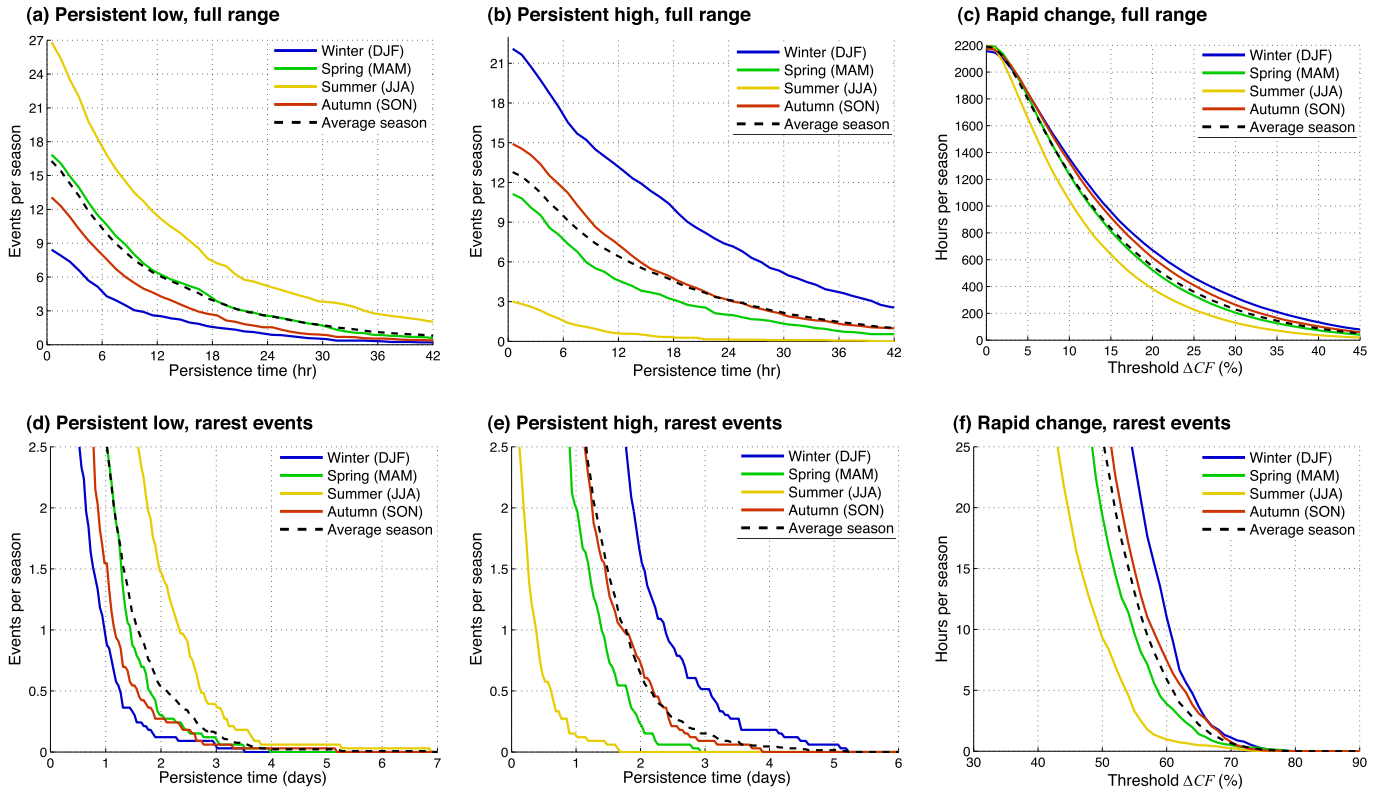
The 33 year mean capacity factor (CF) for GB is estimated at 32.5% (median 26.4%). This is slightly higher than previous long

**Table 2**

The mean seasonal frequency of different extreme events, as derived from MERRA. For each event type, some example thresholds are shown with the corresponding frequency of a mean (3 month) season, as well as the mean frequencies for events occurring only in summer (June to August) and winter (December to February).

Event type	Thresholds	Mean season	Mean summer	Mean winter
Persistent low	$CF \leq 6.3\%$ $t_p \geq 24$ hr	2.6	5.3	1.0
Persistent high	$CF \geq 69.6\%$ $t_p \geq 24$ hr	3.2	0.2	7.4
Extreme ramp	$\Delta CF \geq 50\%$ $t_{win} = 12$ hr	25.8	9.3	43.3



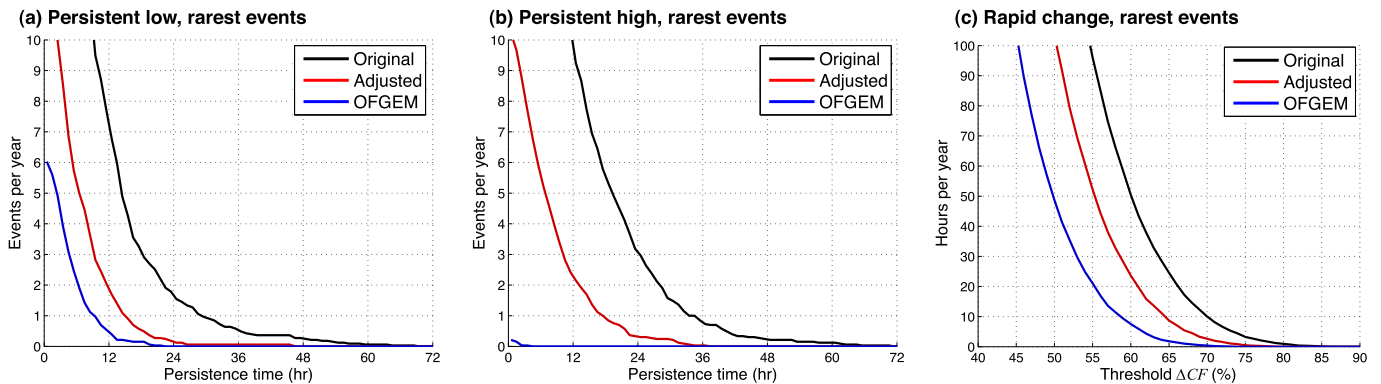


**Fig. 12.** The mean seasonal frequency of extreme generation events derived from the MERRA reanalysis (1980–2012), for four seasons and a mean season. Panels show (a) persistent low generation ( $CF \leq 6.3\%$ ), (b) persistent high generation ( $CF \geq 69.6\%$ ), and (c) ramps in generation within a 12 h time window. Panels (d–f) are as in (a–c) but show only the rarest events.

term estimates ([3] suggested 30%), which may be due to the different geographical distributions assumed (especially as Sinden did not include offshore sites, which tend to be windier). The annual-mean CF was found to range from 23.0% (in 2010) to 34.2% (in 1986). Such variability, if reflected on a site-by-site basis, would be highly relevant to the financing and operational revenue streams of wind farms as well as energy prices and trading.

The climatology is used to estimate the mean frequency of extremely persistent low and high wind power generation events across a wide range of thresholds. Moderately persistent low generation events (at least 2 days with  $CF \leq 5\%$ ) are found to occur 1.2 times  $yr^{-1}$ , whereas the lowest generation threshold for which there was a continuous 5 day lull in generation [2] was  $CF \leq 6\%$ . The

number of both low and high generation events decreases approximately exponentially with increasing persistence, implying they can be approximated as a Poisson-like process. This also demonstrates that there are no a priori meteorological or statistical reasons to focus on 5 day lulls specifically. These results were also found to contain large seasonal variations, with a tendency for more extended lulls in summer than winter. For example, whilst the most extreme 5 day lull occurred in summer (5 days with  $CF \leq 6\%$ ), in winter it occurred only at the  $CF \leq 9\%$  threshold. Extended periods of low generation (particularly in combination with low temperature and high electricity demand) are important for evaluating the capacity credit of wind power and, potentially, have ramifications for the security of supply in the presence of



**Fig. 13.** The mean frequency of the rarest extreme generation events from 1980 to 2012, as calculated using the three power curves in Fig. 4(b). The adjusted curve is that used to derive the 33 year climatology (Section 3). (a) Low generation events ( $CF \leq 2.2\%$ ), (b) high generation events ( $CF \geq 87.1\%$ ) and (c) ramps within a 12 h time window.



limited gas reserves. High generation events may become increasingly important as installed wind capacity increases against a relatively fixed transmission system, inducing deliberate curtailment to ensure local load balancing [32]. In future, we hope to link this research with reanalysis-based estimates of electricity demand (as outlined in Ref. [33]), thus enabling a more thorough investigation of the above power system impacts.

The results derived from MERRA for extreme ramps in generation must be treated with some caution given that generation variability over shorter time scales tends to be underestimated. This is clearly an area where dynamical downscaling can play a significant role (e.g., for evaluating reserve requirements on shorter time scales). Nevertheless, the MERRA-derived time series suggests that ramps of over 60% in GB-aggregated CF within 6 h are possible. Again, these statistics show large inter-annual and seasonal variability. Whilst large ramps are less common in summer than in winter, the size of the most extreme ramps is only slightly larger in winter. The degree to which extreme ramps are accurately predicted by operational weather forecast models is currently being investigated.

The statistics presented here were found to be quantitatively sensitive to the choice of wind farm power curve. The power curve used to derive the 33 year climatology was based on a single-turbine response curve and assumed a deterministic relationship between wind speed and generation. The sensitivity to changes in the power curve was found to be a consequence of the associated shift in the cumulative CF distribution. These sensitivities could be reduced using more accurate, farm-specific, power curves. The construction of these curves would benefit greatly from increased public access to farm-level generation data (site-specific, high frequency generation and turbine availability), as well as the adoption of a probabilistic, rather than deterministic transformation between wind speed and wind farm generation.

## Acknowledgements

The authors are grateful for the support of National Grid, who funded this research and provided data, and for the comments from two anonymous reviewers. The MERRA reanalysis data was provided by the Global Modeling and Assimilation Office (GMAO) at NASA Goddard Space Flight Center through the NASA GES DISC online archive. The surface-based wind speed observations were provided by the UK Meteorological Office via the British Atmospheric Data Centre (BADC). All plots were produced using MATLAB, and Fig. 4(a) was generated using the M\_MAP mapping package v1.4f (<http://www2.ocgy.ubc.ca/~rich/map.html>).

## References

- [1] OFGEM. Electricity capacity assessment report. Office of Gas and Electricity Markets; 2013.
- [2] Mackay DJC. Sustainable energy - without the hot air. UIT Cambridge; 2009.
- [3] Sinden G. Characteristics of the UK wind resource: long-term patterns and relationship to electricity demand. *Energy Policy* 2007;35:112–27.
- [4] Earl N, Dorling S, Hewston R, von Glasow R. 1980–2010 variability in u.k. surface wind climate. *J Clim* 2013;26:1172–91.
- [5] Gross R, Heptonstall P, Anderson D, Green T, Leach M, Skea J. The costs and impacts of intermittency. UK Energy Research Centre's Technology and Policy Assessment (TPA) function; 2006.
- [6] Zachary S, Dent C, Brayshaw DJ. Challenges in quantifying wind generation's contribution to securing peak demand. In: IEEE Power and Energy Society General Meeting, Detroit, MI; 2011. p. 1–8.
- [7] Brayshaw DJ, Troccoli A, Fordham R, Methven J. The impact of large scale atmospheric circulation patterns on wind power generation and its potential predictability: a case study over the UK. *Renew Energy* 2011;36:2087–96.
- [8] Bett P, Thornton HE, Clark RT. European wind variability over 140 yr. *Adv Sci Res* 2013;10:51–8. 12th EMS annual meeting and 9th European conference on applied climatology (ECAC) 2012.
- [9] Pöyry. Impact of intermittency: how wind variability could change the shape of the British and Irish electricity markets. Pöyry Energy Consulting; 2009. Available online at: [www.poyry.co.uk/news](http://www.poyry.co.uk/news).
- [10] Früh W-G. Long-term wind resource and uncertainty estimation using wind records from Scotland as example. *Renew Energy* 2013;50:1014–26.
- [11] Leahy PG, McKeogh EJ. Persistence of low wind speed conditions and implications for wind power variability. *Wind Energy* 2013;16:575–86.
- [12] Kiss P, Varga L, János IM. Comparison of wind power estimates from the ECMWF reanalyses with direct turbine measurements. *J Renew Sustain Energy* 2009;1:033105.
- [13] Kubik ML, Brayshaw DJ, Coker PJ, Barlow JF. Exploring the role of reanalysis data in simulating regional wind generation variability over Northern Ireland. *Renew Energy* 2013;57:558–61.
- [14] Staffell I, Green R. How does wind farm performance decline with age? *Renew Energy* 2014;66:775–86.
- [15] Brower MC, Barton MS, Lledó L, Dubois J. A study of wind speed variability using global reanalysis data. 2013. AWS Truepower technical report.
- [16] Rienecker MM, Suarez MJ, Gelaro R, Todling R, Bacmeister J, Liu E, et al. MERRA - NASA's modern-era retrospective analysis for research and applications. *J Clim* 2011;24:3624–48.
- [17] Dee DP, Uppala SM, Simmons AJ, Berrisford P, Poli P, Kobayashi S, et al. The ERA-interim reanalysis: configuration and performance of the data assimilation system. *Q J R Meteorol Soc* 2011;137:553–97.
- [18] Decker M, Brunke MA, Wang Z, Sakaguchi K, Zeng X, Bosilovich M. Evaluation of the reanalysis products from GSFC, NCEP, and ECMWF using flux tower observations. *J Clim* 2012;25:1916–44.
- [19] UK Meteorological Office. Met office integrated data archive system (MIDAS) land surface station data (1853–current). NCAS British Atmospheric Data Centre; 2013. Available from: <http://badc.nerc.ac.uk/>.
- [20] Howard T, Clark P. Correction and downscaling of NWP wind speed forecasts. *Meteorol Appl* 2007;14:105–16.
- [21] Hasche B. General statistics of geographically dispersed wind power. *Wind Energy* 2010;13:773–84.
- [22] Barthelmie RJ, Pryor SC, Frandsen ST, Hansen K, Schepers JG, Rados K, et al. Quantifying the impact of wind turbine wakes on power output at offshore wind farms. *J Atmos Ocean Tech* 2010;27:1302–17.
- [23] Conroy N, Deane JP, Gallachóir PPO. Wind turbine availability: should it be time or energy based? - a case study in Ireland. *Renew Energy* 2011;36:2967–71.
- [24] National Grid plc. Electricity ten year statement. 2013.
- [25] Department of Energy, Climate Change (DECC). Digest of UK energy statistics (DUKES). 2013. published 25th July.
- [26] Vautard R, Cattiaux J, Yiou P, Thépaut J-N, Clais P. Northern hemisphere atmospheric stilling partly attributed to an increase in surface roughness. *Nat Geosci* 2010;3:756–61.
- [27] Ely C, Brayshaw DJ, Methven JM, Cox J, Pearce O. Implications of the North Atlantic oscillation for a UK-Norway renewable power system. *Energy Policy* November 2013;62:1420–7. <http://dx.doi.org/10.1016/j.enpol.2013.06.037>.
- [28] Masato G, Hoskins BJ, Woolings TJ. Can the frequency of blocking be described by a red noise process? *J Atmos Sci* 2009;66:2143–9.
- [29] Holton JR, Hakim GJ. An introduction to dynamic meteorology. 5th ed. Elsevier Inc.; 2012.
- [30] Woolings T, Hannachi A, Hoskins B. Variability of the North Atlantic eddy-driven jet stream. *Q J R Meteorol Soc* 2010;136:856–68.
- [31] Hawkins S, Harrison G. A reanalysis of UK wind speeds using the WRF mesoscale model. In: European Wind Energy Conference and Exhibition (EWEC); 2010.
- [32] Quest H. Energy meteorology: stormy times for UK wind power [Master's thesis]. University of Reading; 2013.
- [33] Brayshaw DJ, Dent C, Zachary S. Wind generation's contribution to supporting peak electricity demand - meteorological insights. *Proc Institution Mech Eng Part O J Risk Reliab* 2012;226:44–50.