

How does wind farm performance decline with age?☆



Iain Staffell*, Richard Green

Imperial College Business School, Imperial College London, London SW7 2AZ, UK

ARTICLE INFO

Article history:

Received 17 May 2013

Accepted 25 October 2013

Keywords:

Wind farm
Load factor
Degradation
Ageing
Reanalysis
Levelised cost

ABSTRACT

Ageing is a fact of life. Just as with conventional forms of power generation, the energy produced by a wind farm gradually decreases over its lifetime, perhaps due to falling availability, aerodynamic performance or conversion efficiency. Understanding these factors is however complicated by the highly variable availability of the wind.

This paper reveals the rate of ageing of a national fleet of wind turbines using free public data for the actual and theoretical ideal load factors from the UK's 282 wind farms. Actual load factors are recorded monthly for the period of 2002–2012, covering 1686 farm-years of operation. Ideal load factors are derived from a high resolution wind resource assessment made using NASA data to estimate the hourly wind speed at the location and hub height of each wind farm, accounting for the particular models of turbine installed.

By accounting for individual site conditions we confirm that load factors do decline with age, at a similar rate to other rotating machinery. Wind turbines are found to lose $1.6 \pm 0.2\%$ of their output per year, with average load factors declining from 28.5% when new to 21% at age 19. This trend is consistent for different generations of turbine design and individual wind farms. This level of degradation reduces a wind farm's output by 12% over a twenty year lifetime, increasing the levelised cost of electricity by 9%.

© 2014 The Authors. Published by Elsevier Ltd. All rights reserved.

1. Introduction

Ageing is a fact of life. Its effects are inevitable for all kinds of machinery, reducing the efficiency, output and availability of steam and gas turbines, solar PV modules, batteries and automobiles alike. Previous work on wind turbines has considered the reliability of individual components and the effect of ageing on availability, but any impact on the energy production of turbines or farms has not been widely reported.

If load factors (also known as capacity factors) decrease significantly with age, wind farms will produce a lower cumulative lifetime output, increasing the levelised cost of electricity from the plants. If the rate of degradation were too great, it could become worthwhile to prematurely replace the turbines with new models, implying that the economic life of the turbine was shorter than its technical life, further increasing its cost.

This could have significant policy implications for the desirability of investing in wind power, as argued in a recent report by

Hughes for the Renewable Energy Foundation (REF) [1]. That report suggested that the load factors of wind farms in the UK have declined by 5–13% per year, normalising for month-by-month variations in wind speeds. These findings could represent a significant hurdle for the wind industry, but they require replication.

Several factors can confound the relationship between age and observed output in a fleet of wind farms, given that a turbine's output is dependent on wind speeds at its site and the efficiency with which it captures the energy in that wind. For example, if wind speeds have fallen slightly over time, farms would have lower load factors in recent months, when they were at their oldest, giving a spurious correlation between age and poor performance. If improvements in design increase a turbine's output relative to capacity (its power coefficient) then newer turbines (of the improved design) will have higher load factors than old turbines, so that turbine output appears to decline with age, when really it improves with newer generations. On the other hand, if the best (windiest) sites were occupied first, then old farms could have higher load factors than new ones built on inferior sites, so that turbines would appear to improve with age.

This paper uses public domain data to infer the hour-by-hour wind speeds at the site of every wind farm in the UK, and the power curve for each farm's model of turbine to estimate the output that they would ideally produce. This technique corrects for the

☆ This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

* Corresponding author.

E-mail addresses: staffell@gmail.com, i.staffell@imperial.ac.uk (I. Staffell).

confounding factors (wind patterns, turbine model and site quality), and validates well for farms that report their half-hourly output to National Grid. Simulated ideal outputs are compared with actual monthly load factors from a large portion of the UK's fleet over the last decade (282 wind farms, 4.5 GW, 53 TWh), yielding the normalised performance of each wind farm accounting for its wind resource availability, and a set of weather-corrected load factors which reveal the effects of ageing. We measure the level of age-related degradation at the national level, accounting for the vintage of turbine and local site conditions at each wind farm. We test different generations of technology and individual wind farms to confirm that specific units experienced similar declines in performance. We find the ageing effect to be present, but much smaller than predicted by Hughes, in line with experience of other rotating machinery. The specific causes of this performance loss and their relative contribution are not considered in this paper, although an overview of potential reasons is given in the discussion and conclusions.

Due to the amount of data and processing required for this study we provide online supplementary material which documents our sources and their validation in greater depth, along with downloadable datasets of UK wind farms and their energy output histories.

2. Previous studies

All machinery experiences an unrecoverable loss in performance over time. Gas turbine efficiency suffers an unrecoverable decline of 0.3–0.6% per year despite regular washing and component replacement, or by 0.75–2.25% without [2]. Similarly, the output of solar photovoltaic panels declines by 0.5% per year on average [3]. This loss in performance is not routinely accounted for in studies of the levelised cost of electricity (LCOE) of wind power. Recent studies by Mott MacDonald, Parsons Brinckerhoff and Arup accounted for the efficiency of conventional plants falling by 0.15–0.55% per year, but omitted any such factor for wind turbines [4–6].

Previous studies of wind turbines have focussed on availability and reliability [7–9]. There appear to be no long term fleet-level studies into loss of output from wind farms in open literature. Regardless of technology, quantifying performance degradation is difficult because consistent and validated field data is hard to obtain [2]. The recent study by Hughes [1] is therefore significant, in that we believe it is the first to attempt to estimate the rate of decline in wind farm load factors on a national scale.

Hughes analysed over 10 years of operating data from the British and Danish fleets of turbines, finding rates of performance degradation that are much higher than for other technologies, and which vary remarkably between the UK and Denmark, and between onshore and offshore turbines. This was based on econometric analysis of monthly load factors, using a regression which corrected for the quality of each wind farm's location, the monthly variation in national wind conditions, and the age of each farm. Hughes argues (and shows mathematically) that accounting for monthly wind conditions with a set of 'fixed effects' determined by the regression is econometrically superior to using a measure of average wind speeds across the country, since site-specific conditions differ from the national average and the output of wind turbines depends non-linearly on the wind speed at every moment in time, which is very poorly captured by its average over a month.

We therefore use wind speed data with high temporal and spatial resolution, and measure the performance of wind farms by estimating their theoretical potential output over the course of a month and comparing this with the actual reported load factors. While we believe we are the first researchers to assess wind farm performance with this kind of *ex-post* data, a number of papers present techniques to estimate output levels from time series of wind data.

Many studies have used hourly wind speed data recorded by met masts; for example investigations into wind variability by Pöyry [10] and SKM [11], and estimates of future national output by Green et al. [12,13] and Sturt and Strbac [14]. Hourly met mast speeds have been directly compared to metered wind farm load factors in Northern Spain [15] and Scotland [16], showing that accurate estimates can be made for monthly energy generation, but not for hourly power outputs.

More recent studies use reanalyses as a source of wind speed data: atmospheric boundary layer models which process physical observations from met masts and other sources into a coherent and spatially complete dataset, and are widely used to produce wind atlases. Kiss et al. [17] were first to compare the European ERA-40 reanalysis to nacelle measurements of wind speed and power output at two turbines in Hungary, finding "surprisingly good" agreement. Hawkins et al. [18] were able to replicate UK monthly load factors using a custom reanalysis model, while Kubik et al. [19] compared the global NASA reanalysis to half-hourly farm output in Northern Ireland, finding it to be more accurate than met mast data. The first practical application appears to have been made by Ofgem to estimate the equivalent firm capacity of the UK's wind fleet during winter peaks in demand [20]. Both Hawkins and Ofgem noted that the reanalysis outputs need to be scaled down by a constant factor (29% and 20% respectively) in order to match actual production in the UK, a finding which we elaborate upon in this paper.

3. Data sources

Predicting a given wind farm's output is far from being a new science: on-site monitoring of conditions using wind turbine SCADA systems is commonplace; and software tools such as WaSP or consultancies such as GL Garrad Hassan are widely used in the field. Data is not made publicly available, and these services come at a price of several thousand Euros.

On the other hand, national average data cannot reveal what is happening at individual wind farms. We therefore employ farm-specific data for output and site-specific data for wind speeds, taken from free and publicly accessible datasets. The primary data used in our main analysis are described in this section, and additional data used for validation are described in Section 4. Further information on our data is given in the [Supplementary Material](#).

3.1. Ofgem/REF output data

All wind farms enrolled in the UK government's incentive scheme, the Renewables Obligation, publish their monthly outputs (in MWh) in the Ofgem Renewables and CHP Register.¹ Hughes extracted and cleaned this data, cross-linking outputs with details about each wind farm (its capacity and date of commissioning), and ensured that each wind farm contained only the same model and vintage of turbine [1]. This cleaned dataset was published on the internet by the Renewable Energy Foundation (REF) [21]. We are very grateful to Prof. Hughes and the REF for making this rich data source available to the community.

We further validated this dataset, corrected the commissioning date for 15 wind farms (which were incorrectly reported by Ofgem), integrated further meta-data for each farm (the geographical location, wind turbine model and hub height), and extended the time-series by 8 months, adding data from April to December 2012.

Our modified dataset is provided as [Supplementary Material](#) to this paper. It contains 1687 farm-years of load factor data, covering onshore turbines built from 1991 onwards and spanning 11 years of

¹ www.renewablesandchp.ofgem.gov.uk.

operation. The study was restricted to onshore wind farms, as only a small amount of data was available for the UK's 20 offshore farms, and these would need to be considered separately as they face a very different operating environment and maintenance issues to onshore farms. Table 1 and Fig. 1 provide a selection of summary statistics for the data.

Figs. 2 and 3 show what has happened to the load factor of these wind farms as they get older. Fig. 2 simply plots the distribution of all observed load factors against age, showing a steady decline of -0.44 ± 0.04 absolute percentage points per year past age one ($1.69 \pm 0.17\%$ loss per year relative to the UK mean load factor).

As explained in the introduction this is not necessarily due to ageing if newer turbine models are more efficient at extracting energy from the wind. The oldest turbines in the sample (aged 15–19 years) were built in the early- to mid-1990s; typically 300–500 kW two or three bladed machines on 25–50 m towers. These will clearly be outperformed by the latest generation of 2–3 MW turbines which are no older than 5 years.

To control for technology effects, Fig. 3a charts the individual histories of the 53 farms which have more than ten years of data, using a 12-month moving average to smooth out seasonal variations in the wind. Load factors tend to rise during the first year of operation while turbines are still being commissioned until the farm achieves full operation and teething problems are ironed out. Fig. 3b summarises the annual degradation rates, estimated by running individual linear regressions on each farm's unsmoothed load factors, excluding the first year. The areas with the darkest shading are derived from farms which are old enough to give ten years of data; lighter areas add the more recent farms which have fewer observations.

The degradation rates of individual farms are predominantly bunched around 0 to -1 percentage points of absolute load factor per year, but several outliers make the overall distribution fat-tailed. A Cauchy (or Lorentz) distribution therefore provides a better fit than a normal distribution, centred on -0.48 with a half width at half maximum of 0.36. The distribution becomes wider for farms with fewer observations; and would stretch all the way from -25 to $+15$ points per year if farms with less than five years' data were included. With an absolute degradation rate of -0.48 ± 0.36 points per year, a typical wind farm loses $1.81 \pm 1.32\%$ of its output per year on average. The range of degradation rates for individual farms is clearly much greater than uncertainty on the average rate for the dataset as a whole.

Neither Fig. 2 nor Fig. 3 corrects for weather effects. To extract the true rate of degradation in the UK's wind farms, we account for variations in the weather over the last decade using a detailed wind resource assessment, and the rate of technology improvement by modelling the specific turbines installed at every wind farm.

3.2. NASA wind speed data

A database of wind speeds for the British Isles was created using NAA's MERRA dataset: a historical reanalysis of global atmospheric observations assimilated and processed using the Goddard Earth Observing System (GEOS-5) [22,23]. Wind observations with "fairly complete global coverage" are taken from weather stations, balloons, aircraft, ships, buoys and satellites, and processed by the model to give data with hourly resolution on a $\frac{1}{2}^\circ$ latitude and $\frac{2}{3}^\circ$ longitude grid (approx. 55 by 44 km), at heights of 2 and 10 m above the surface displacement height (d , the point at which a logarithmic wind profile would tend to zero) and at 50 m above ground.

We acquired data for the UK, Ireland and surrounding waters ($-15\frac{1}{3}^\circ$ to 10° E, 46.5° to 65.5° N) from 1993 to 2012, giving a database with 1.06 billion observations ($175,320$ temporal \times 1521 geographic \times three speed variables plus displacement height). This database was then processed using R [24] as follows:

Table 1

Summary statistics for the extended UK wind output dataset.

Number of observations	20,243
Temporal resolution	Monthly
Number of wind farms	282 onshore
Farm capacity	0.5–322 MW
(median)	6.5 MW
(mean)	13.5 MW
Total capacity	4.4 GW
Construction dates	1991–2012
Data recording period	2002–2012

- For each hour, the east (u) and north (v) components of wind speed were extracted at all three heights;
- A nonparametric polynomial surface was fitted to each set of spatially gridded observations using a 2-dimensional LOESS regression. This allowed wind speeds at any coordinates to be locally interpolated using the nearest twelve observations, as in Fig. 4;
- The magnitude of the wind speed vector (w) was calculated at each location from $w = \sqrt{u^2 + v^2}$;
- Wind speed was extrapolated from the three observation heights to the hub height of each farm (averaging 60 ± 14 m) using the log law described in the next section.

4. Data processing and validation

The aim of our method was to use NASA modelled wind speed data as a predictor for the monthly energy output of wind farms. Two forms of validation were performed to give confidence that this technique can estimate load factors for a given wind farm that are broadly representative and free of seasonal or inter-annual bias.

The first test was whether interpolated values from the NASA GEOS-5 model accurately represent actual wind speeds measured at a particular location; and the second was that these speeds, when extrapolated to hub height and transformed using a representative power curve, accurately match the actual metered output from a wind farm.

The data processing and validation was split into five stages:

- The NASA speeds at 10 m were compared to ground-based observations from the Met Office to validate our use of the GEOS-5 model (reported in our [Supplementary Material](#));
- Wind speeds at each farm's location were extrapolated from 50 m to the hub height of that farm to account for wind shear;
- These extrapolated speeds were transformed into estimated ideal load factors using the power curve of the installed turbine model;
- The ideal hourly load factors were compared to the half-hourly metered output data for farms where this was available, to check that the preceding methods were robust;
- Ideal load factors were aggregated for each farm monthly over the period of 2002–12, for comparison against the Ofgem/REF dataset.

The following sections briefly summarise our methods and findings. An extended validation section which covers each topic in greater depth is provided as supplementary material.

4.1. Extrapolating wind speeds to turbine hub height

Extrapolating wind speeds from the height of measurement stations to the much higher hub height of wind turbines is "probably one of the most critical uncertainty factors affecting the wind power assessment at a site" [25]. The change in horizontal wind speeds with height (known as wind shear or the wind profile) is

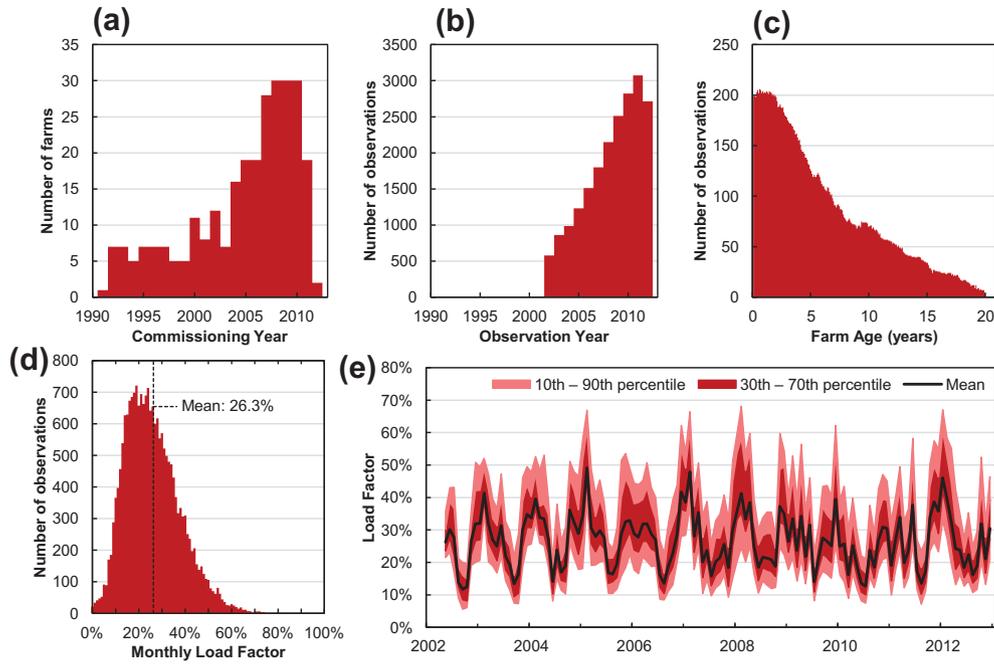


Fig. 1. Histograms summarising the vintage of UK wind farms in our dataset (a), the distribution of monthly observations from this fleet (b–c), and the distribution of load factors for individual farms over the last eleven years (d–e).

generated by friction from the earth's surface, and so is highly dependent on the specific site conditions: the surface roughness of the terrain, air temperature, season, atmospheric stratification, and the wind speed itself.

Many simplifications for extrapolating wind speeds exist, most notably the empirically derived power law and the theoretically derived log law (Eq. (1)) [25,26]

$$w_x = A \log\left(\frac{h_x - d}{z_0}\right) \quad (1)$$

The simplified log law assumes that wind speed, w , is related to the logarithm of its height, h , under the assumption of neutral atmospheric stability. The logarithmic wind profile tends to zero at a height of the roughness length, z_0 , plus the surface displacement height, d ,

and is scaled by a constant, A , which equals the friction velocity (u^*) divided by 0.4 (von Karman's constant). The log law provides robust extrapolations of wind speeds over a range of wind speeds, locations and altitudes [25,27]. The latter is especially important as turbine hub heights range from 25 to 100 m, and the Hellman exponent used in the power law decreases non-uniformly with height.

The NASA data give simultaneous wind speeds at three heights, along with d , allowing the coefficients for wind shear to be calculated for the specific site and time of the observation. Eq. (1) was linearised to Eq. (2), allowing the coefficients A and z_0 to be estimated by a least-squares regression of the three NASA observations. The coefficient values were independently estimated for each site and time period with no smoothing or prior values, but were found to be temporally and spatially stable.

$$w_x = A \log(h_x - d) - A \log(z_0) \quad (2)$$

At the average UK hub height of 60 m, wind speeds are 6–9% larger than at 50 m and 32–41% larger than at 10 m. The NASA data greatly reduces the extrapolation uncertainty compared with using 10 m met mast data.

4.2. Converting wind speed to power

The power curve for an 'ideal' wind farm was applied to the speed data at each site, estimating the potential output from a farm with perfect availability, perfect calibration, and no site-related performance loss (e.g. turbulence from surrounding geography or wake effects from other turbines). We consider this to be the maximum possible attainment and name it the ideal yield.

The meta-data that we integrated into the Ofgem/REF dataset gave the model of turbine used at each of the 282 wind farms. We compiled the power curves for 50 of these turbine models, accounting for 92% of the wind farms in the UK. For the remaining farms, the best match was found from the known curves, based on the installed turbine's capacity and power density (peak power divided by swept area). [Supplementary Table 1](#) gives details of all the turbines considered.

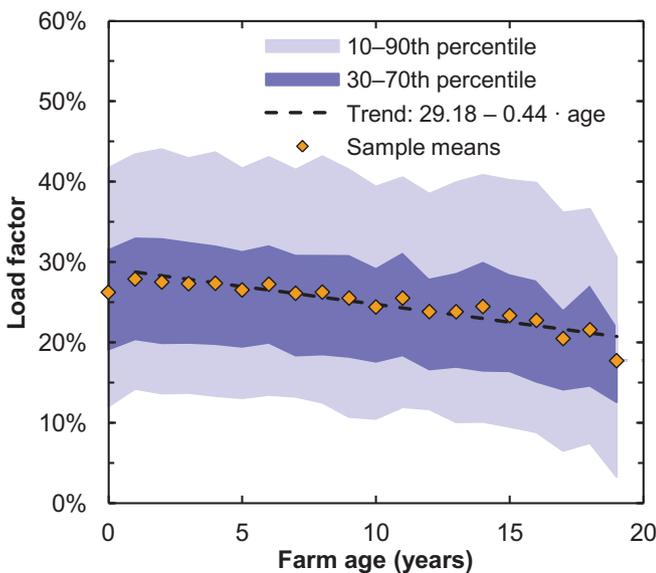


Fig. 2. The mean and distribution of load factors across all farms as they age.

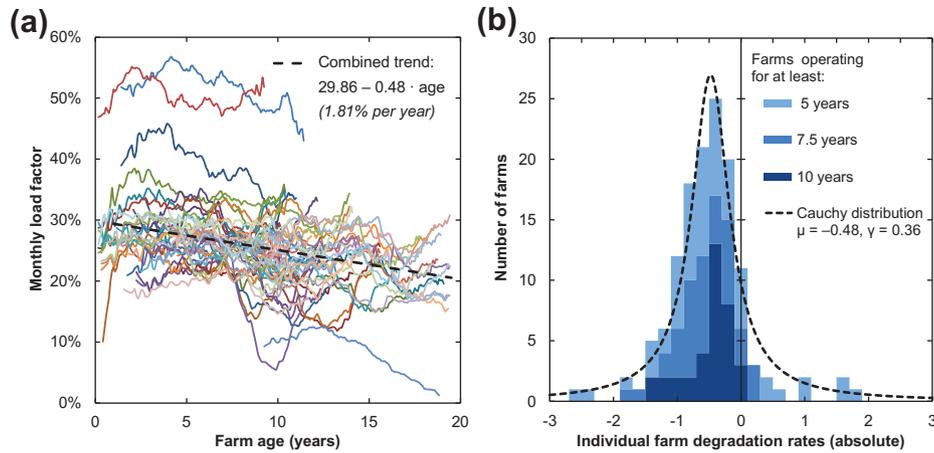


Fig. 3. Lines tracing the output of individual farms as they age (a), and a histogram summarising the linear trends for individual farms, giving the distribution of annual degradation rates (b).

The power curve for a single turbine is shown in Fig. 5 with the aggregate power curve for a typical farm of these turbines. The multi-turbine power curve accounts for the fact that wind speeds at the location of each individual turbine within a farm will vary according to a normal distribution. Following the practice of [10,20,28], the turbine power curve for each farm was convoluted by a normal distribution. The standard deviations were determined by the estimated geographic area that each farm covers [28], based on observations that UK farms occupy 100 m² of land per kW capacity [29]. The mean for each farm was chosen to normalise the total energy production to that of the individual turbine curve.

4.3. Comparing predicted and metered output

The Ofgem/REF output data has monthly resolution, which is too coarse to see whether the structure of hourly wind variations

accurately replicates output from operating turbines. Our second measure of wind farm energy production is the half-hourly metered output from all transmission-connected generators published on Elexon’s TIBCO relay service.² Data from 2005 to 2012 covers 25 TWh of output from the 47 wind farms highlighted in Fig. 6a.

Fig. 6b and c compare the hourly metered output from Black Law wind farm with NASA wind speeds interpolated at its location at 80 m above ground, and the ideal energy yields derived from these speeds.

Each point in Fig. 6b represents one hour’s operation to give the empirical power curve for that farm. This empirical curve follows the features of the farm-aggregated Siemens SWT-2.3–82 power curve (solid line), albeit shifted to the right (as NASA speeds are the theoretical maximum, and actual speeds will be lower due to local site features) and downwards (due to downtime and sub-optimal turbine calibration), with substantial scatter (as local conditions like turbulence vary over time).

Fig. 6c shows 1000 h of metered output together with the simulated ideal output, which was scaled by a constant factor of 0.698 for reasons explained in the next section. The simulated ideal load factors for other farms compare similarly well to their metered outputs, with plots given in Supplementary Fig. 9.

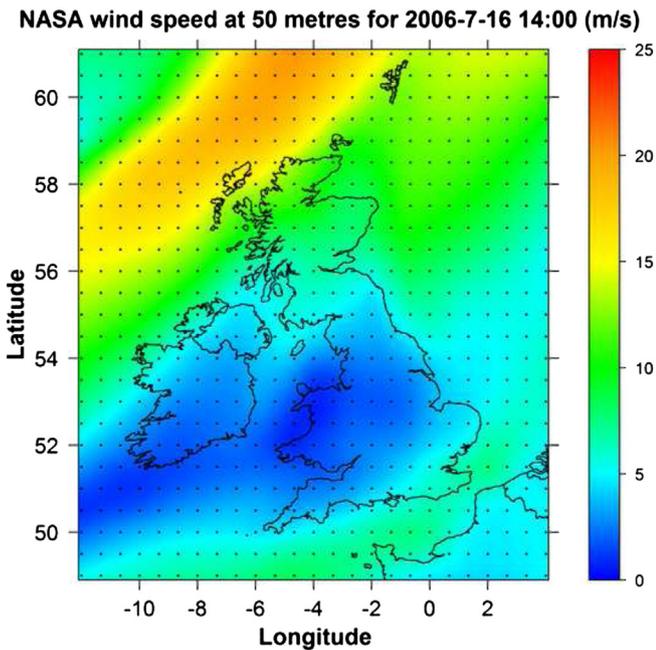


Fig. 4. Colour map showing the interpolated wind speeds for a given hour over the British Isles, highlighting the locations of MERRA grid points. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

5. Results

5.1. Simulated wind speeds

The average monthly wind speed at UK onshore wind farms was estimated to be 7.5 ± 1.5 m/s at the location and hub height of each farm, which average 62 m above ground. This average speed has experienced a slight decline over the last 12 years, although the trend is not statistically significant (-0.23 ± 0.37 m/s per decade). Fig. 7 shows the correlation between the simulated monthly NASA wind speeds, averaged overall operating sites, and the national average load factors reported to Ofgem. The correlation is very high as the simulation is able to represent the actual sites of generation and the evolution of this site population over time, reaffirming the strong linear relationship between speed and output at monthly resolution. The correlation between monthly average speed and reported load factor is also high for individual wind farms (averaging 0.84), as shown in Supplementary Fig. 12.

² www.bmreports.com.

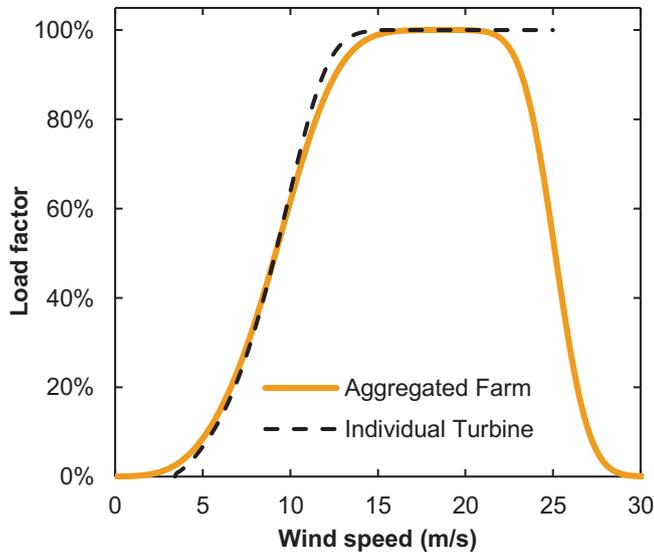


Fig. 5. The power curve for a single Vestas V80 2 MW wind turbine and for a 50 MW farm of these turbines.

If performance declines with wind farm age, we would expect the reported load factor to fall relative to the simulated wind speed over the sample period. This is in fact the case although it is hard to see in Fig. 7, because so many new farms have been built that the average age of the UK fleet has only risen only from 4.8 to 7.2 years between 2002 and 2012. The degradation rates derived below imply that this would reduce the fleet's average load factor by only 1 percentage point.

5.2. Ideal load factor and performance ratio

The ideal load factors derived from these wind speeds ignore a number of factors that will reduce the actual output attained by a wind farm at the given location and hub height. Three are well-understood:

1. *Machine availability*: analysis of national fleets suggests 4–7% downtime for farms and the electrical infrastructure they rely upon [9,30], which translates to an 11% reduction in energy output as turbines on average fail in windier than average conditions [31];
2. *Operating efficiency*: sub-optimal control systems, misaligned components and electrical losses within the farm are found to reduce output by 2% in well-performing field installations relative to the turbine's supplied power curve [30];
3. *Wake effects*: wind farms suffer from power loss as interactions between neighbouring turbines increase turbulence and reduce wind speeds; for relatively small (up to 20 turbine) onshore farms estimates are in the region of 5–15% [25,32–35];

and two are less well understood:

4. *Turbine ageing*: based on the findings from Fig. 2 (and presented later in this paper), energy output from the UK's fleet is 7.5% lower than it would be for a fleet of the same turbines as-new, due to their average age being 5.9 years across the sample period³;

³ In practice ageing is not a separate issue from availability and efficiency, as these likely fall over a turbine's lifetime from the as-new values listed in points 1 and 2, producing the ageing effect that we observe.

5. *Site conditions*: imperfections in a turbine's surroundings are not considered in our model; for example: turbulence intensity, terrain slope, blockage effects, blade fouling (by dirt, ice, insects, etc.), or masking by surrounding terrain. These impacts are highly site specific and hard to quantify with a single factor, with the only source we found estimating that they reduce output by 2–5%, plus 1% per 3% increase in turbulence intensity [36].

Combining the first four terms, we could expect the ratio of observed to ideal load factors to be $0.89 \cdot 0.98 \cdot 0.90 \cdot 0.925 = 0.725$. We call this metric the Performance Ratio (PR), which is analogous to availability, except it deals with output rather than uptime. Fig. 8 plots the relationship between actual and ideal load factors, showing that the performance ratio is unbiased across the range of simulated wind conditions.

Based on the simulated wind speeds and the model of turbine installed at each farm, the ideal load factor of UK onshore wind farms should average 38.4%, whereas the mean observed load factor for these farms from 2002 to 12 has been 26.3%. The average performance ratio of the farms is therefore $68 \pm 19\%$, confirming previous work from Ref. [18] which found that a scale factor of 0.69 gave good correlation between load factors derived from a custom reanalysis and the Ofgem ROC data.

This result does not imply that UK wind farms produce only two-thirds of what they ought to, for the ideal yield represents a hypothetical turbine sited on perfectly flat and smooth terrain, several kilometres from other turbines, foliage or buildings. The real-world factors 1–4 listed above suggest that a performance ratio of 0.725 should be expected, which leaves a reduction of 4 percentage points (6%) attributable to the specific site conditions for UK turbines.

5.3. Weather-corrected load factors

Combining the observed and ideal load factor data allows us to calculate a weather-corrected load factor ($WCLF_{f,t}$) for every observation, giving a time series for each farm that should not be affected by wind conditions changing from month to month. The actual load factor ($LF_{f,t}$) for a given farm (f) and month (t) is divided by the ratio of its ideal load factor ($ILF_{f,t}$) for that month to the farm's mean ideal load factor over the entire 2002–12 period (whether or not the farm existed throughout the whole period), as in Eq. (3). This can be simplified to the performance ratio (PR) for each month multiplied by the farm's average ILF.

$$WCLF_{f,t} = LF_{f,t} \div \left(\frac{ILF_{f,t}}{\overline{ILF}_f} \right) = PR_{f,t} \times \overline{ILF}_f \quad (3)$$

The WCLF represents what a particular farm would have produced each month if wind conditions followed their long-term mean distribution. As with the uncorrected load factors, the absolute value encompasses the available wind resource, the quality of the local site conditions, and the turbine model installed, but the variation over time is no longer dominated by seasonal weather patterns.

This is demonstrated in Fig. 9a: removing the weather noise reveals a gradual decline in this farm's conversion efficacy, and allows periods of low availability to be easily identified. Fig. 9b shows the WCLF averaged across all farms of a given age (in months), revealing the aggregate level of degradation without the need for smoothing. The reduction in scatter reduces the uncertainty on this degradation rate from 0.04 when using nominal load factors (as in Fig. 2), to 0.01 points per year.

Weather correction is notoriously difficult and has the potential to skew results as it makes large changes in the month-to-month

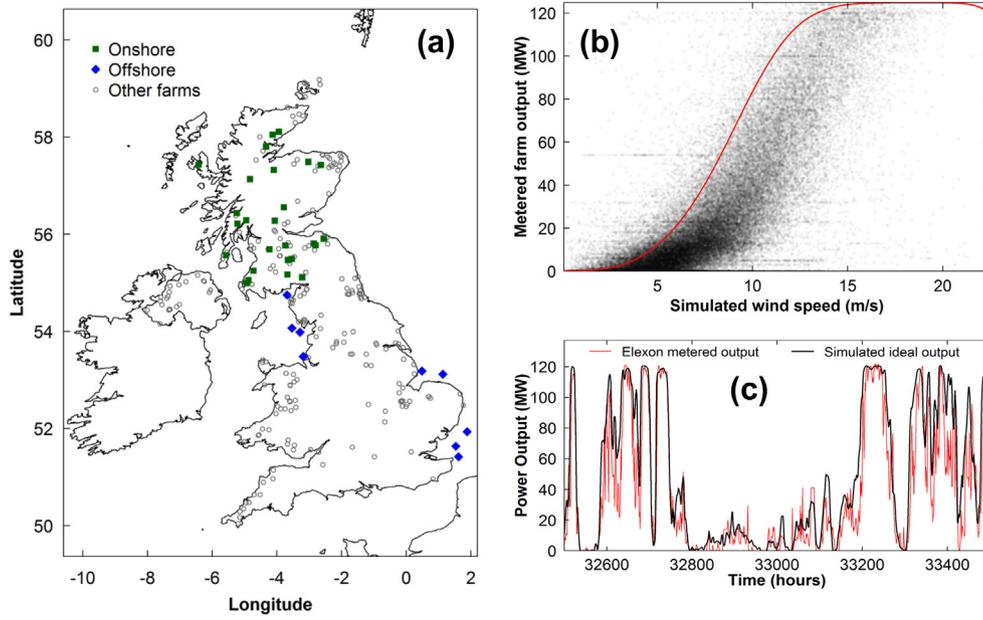


Fig. 6. (a) A map showing the location of the UK's wind farms, highlighting those which were used for validation, with charts showing (b) the hourly metered output from one farm against simulated wind speeds, and (c) a short time series comparing the observed and simulated ideal outputs.

load factor values. We find this correction procedure to be unbiased with wind speed (as in Fig. 8b); and comparing Figs. 2 and 9b shows that it has almost no effect on the fleet-average degradation rate (a change of 0.01 points per year). The distribution of WCLF decline rates at individual farms exhibits little change from that of the unmodified load factors presented in Fig. 3b, averaging -0.45 instead of -0.48 points per year for farms with more than ten years of data. As we find in Supplementary Section 5.3, the weather correction process can clean up the short-term fluctuations without having a systematic effect on the long-term trends.

5.4. Technology improvement over generations

Having rejected the idea that an underlying change in national wind speeds has distorted the results, we look at the evolution of the population of turbines. The oldest farms in our sample are the earliest to have been built, using (presumably) the worst technology, and hence are likely to have the lowest load factors.

Fig. 10 shows the individual degradation rate for each farm against the year it began operating, for all farms with more than 5 years of data. Bubble size is proportional to capacity, bubble colour represents the number of observations, and horizontal bars depict the standard error on each decline rate. Black lines show

the best fit to the data using a capacity-weighted loess regression, showing the central estimate and the range that covers 95% of observations.

The central fit to the raw load factors in Fig. 10a has a mean of -0.50 ± 0.07 points per year of operation, and although it shows noticeable variation between start years there is no long term trend. A few modern farms have seen increasing load factors over the first few years of their lives, but this is offset by the majority having declining outputs.

Correcting for the weather in Fig. 10b greatly reduces the uncertainty on individual farms and the scatter between them. It also reduces the mean decline rate to -0.36 ± 0.05 points overall farms. This decline rate appears stable until 2002, after which it reduces for more recently commissioned turbines. Farms built before 2003 have an average decline rate of -0.49 ± 0.05 points per year, whereas those built afterwards average -0.16 ± 0.08 .

The impact of weather correction is strongest on the farms commissioned most recently as they have the fewest years of data, and so the noise introduced by year-on-year variations in wind speed is strongest. This is most evident in the group of farms commissioned after 2003 which were pulled downwards by the low wind speeds experienced in 2010, and so see a notable improvement in Fig. 10b compared to a.

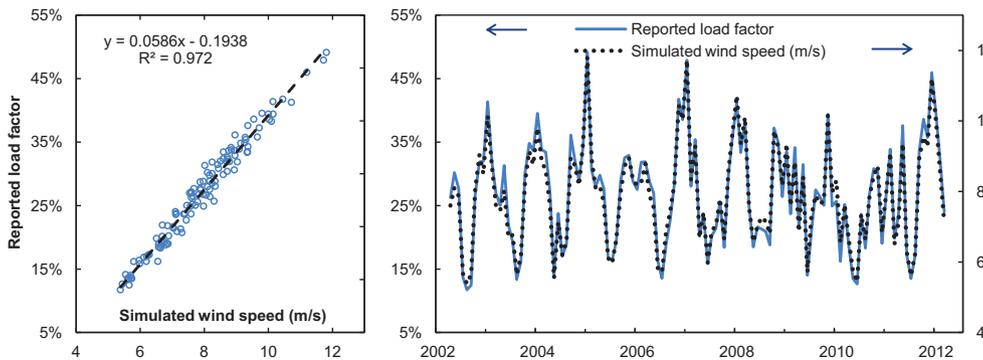


Fig. 7. Correlation between the national monthly average NASA wind speeds and load factors reported to Ofgem.

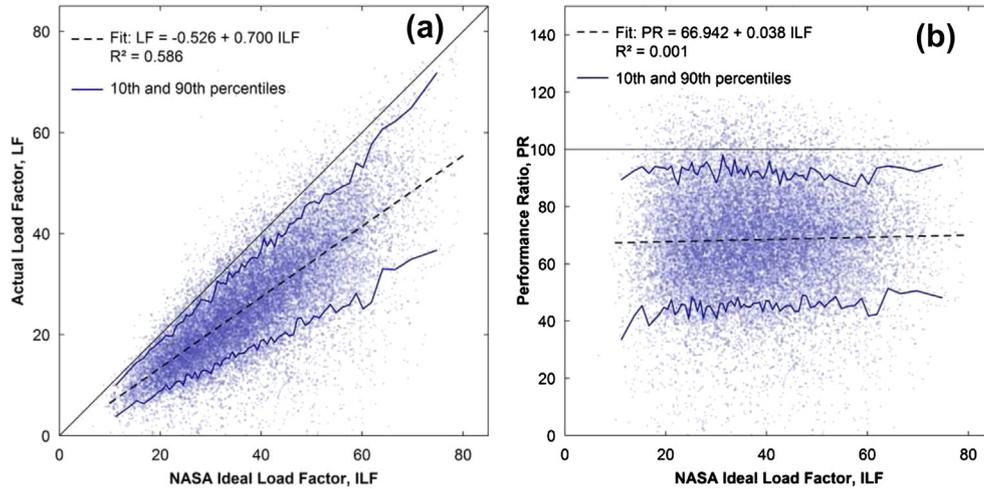


Fig. 8. Comparison of observed and ideal load factors for all farms and periods (a), and the dependence of performance ratio on the ideal load factor (b).

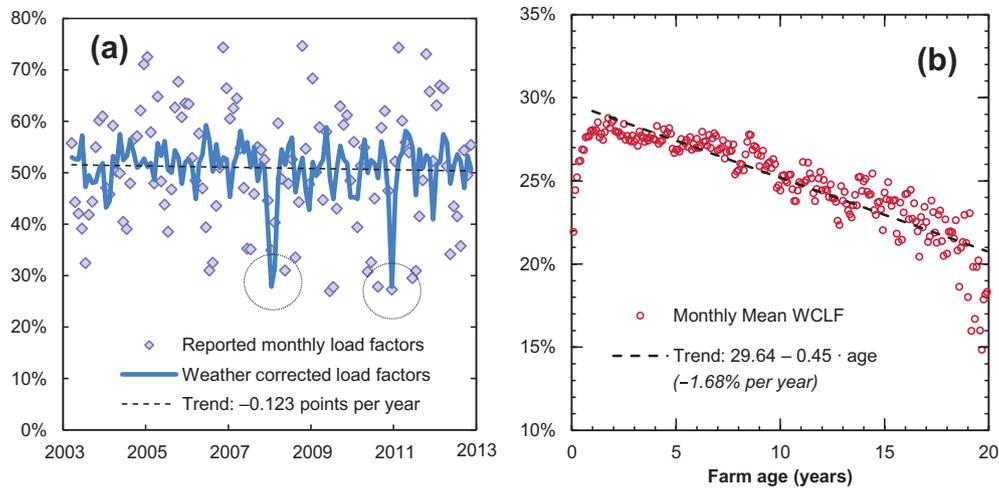


Fig. 9. (a) An example of reported and weather corrected load factors for a single wind farm (Burradale 2), highlighting periods of partial downtime; and (b) the decline in weather-corrected load factor across all farms aggregated by age.

5.5. Full regression of national fleet performance

Four systematic factors determine the actual output of a wind farm: wind speeds at the site, the quality of individual turbine locations (with regards to turbulence and masking), the model of turbine installed and age-related deterioration of performance. Other factors, such as the number of turbines suffering faults in a given month or undergoing planned maintenance, are less systematic.

As demonstrated by Hughes in Ref. [1], it is possible to separate the impact of turbine ageing from these other factors by using an error components model with fixed effects. The observed load factor, $LF_{f,t}$, of wind farm f in month t is estimated by least squares regression against the ideal load factor, $ILF_{f,t}$, with fixed effects for each site, s_f , and age of the turbine, $A_{f,t}$, minimising the sum of the squares of the error component, $\varepsilon_{f,t}$, as in Eq. (4):

$$LF_{f,t} = \alpha + \beta ILF_{f,t} + s_f + A_{f,t} + \varepsilon_{f,t} \quad (4)$$

The fixed effects (s_f and $A_{f,t}$) are the freeform equivalent of a linear trend. A numeric constant is determined for each site by the regression to control for any farm-specific factors which affect the actual wind speeds experienced at the site relative to the NASA

estimates. Similarly, constant modifiers are determined for each age of turbine (in years) to assess the impact of ageing. A linear or quadratic regression against age would not capture complex non-linear behaviour, and so ageing can be better understood through fixed effects. For it to be possible to solve this model, the fixed effect for one site (chosen at random) is held constant at zero, as is the effect for turbines aged 1, which together act as the reference point.⁴

In the model chosen by Hughes [1], the site fixed effects had to account for the model of turbine as well as local site conditions, and period fixed effects were used for each observation month to account for available wind resource (but with the same impact on every farm in that month). In our model, the systematic effect of location-specific wind speeds and turbine model are both incorporated in the ideal load factor data, while the site dummy variables measure the extent to which each farm's surroundings and layout systematically give it more or less wind than the simulation

⁴ In the results presented, the performance of Shooters Bottom (in Somerset) aged 1 is used as a reference.

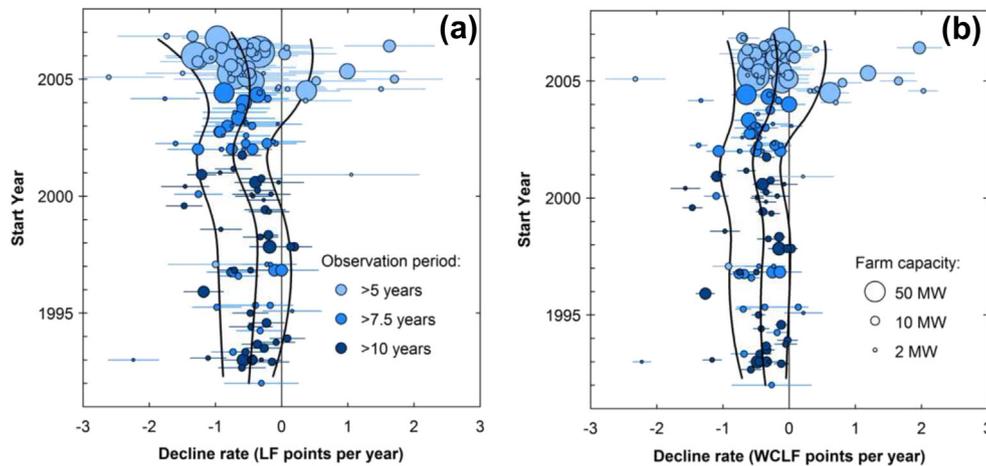


Fig. 10. Bubble plots showing the decline rate of individual farms against the year they began operating, considering (a) unmodified load factors (LF) and (b) weather-corrected load factors (WCLF). Thick vertical lines denote the central fit to each set of points and 95% interval.

predicts. The error term picks up the unsystematic factors affecting each observation.

The regression produced a constant offset of $\alpha = -1.634 \pm 1.529$ percentage points of load factor⁵ and an estimated coefficient on the ideal load factor of $\beta = 0.755 \pm 0.003$. This coefficient is above our average performance ratio of 68% as it is for one-year-old wind farms, and the age effects show that performance declines over time. This model, which uses a highly specific wind resource assessment, provides a better fit to the observed data than using period fixed effects, giving an R^2 of 0.802 compared to 0.657 attained in Ref. [1]. The full regression results are given in [Supplementary Section 5.6](#), along with alternative model formulations which we tested.

[Fig. 11](#) plots the fixed effects produced by this regression, re-centred to give actual load factors as opposed to deviations from the reference point. [Fig. 11a](#) shows the impact of turbine age on the load factor, taking account of spatial and temporal differences in the available wind resource, technology installed and local site conditions.

From [Fig. 11a](#) it can be seen that the uncertainty on the age fixed effects is small enough to be confident that there is indeed degradation, but is too large to be able to discern whether the trend beyond age 1 is linear, exponential, or some more complex function. For simplicity's sake we fit a linear trend to these effects, which falls from 28.5% at age 1 to the national average load factor of 26.3% for farms at the national average age (5.9 years), reaching 21.0% for the oldest farms aged 19.

[Fig. 11b](#) plots the site fixed effects added to the regression offset and the individual contributions from the ideal load factor at each farm (i.e. $\alpha + \beta \text{ILF}_f + s_f$). This combination yields what we call the individual farm effects, which are the model's estimate of each farm's load factor at age 1, accounting for location (estimated wind resource), technology (turbine model and its hub height), and surroundings (local site quality). Regressing these farm effects against the year each farm was built yields no significant trend, implying that any technical improvement has been masked by diminishing site quality.

Many variations on this model can be considered; for example, using an exponential rather than linear fit, substituting the expected load factor with wind speeds ($w_{f,t}$), or estimating the farms'

performance ratio rather than load factor. These models were found to give very similar results to the linear model in Eq. (2), with annual degradation rates that range from 1.5 to 1.9% per year. [Supplementary Section 5.6](#) details these results and shows that our extension of the original REF dataset did not skew the rate of decline.

6. Discussion

The finding that wind turbines lose around 1.6% of their output each year poses three questions:

1. What are the reasons for this deterioration?
2. Can we expect it to continue in future?
3. What are its wider impacts?

6.1. Reasons for declining output

The degradation rate we observe is perhaps to be expected, as it lies in the middle of the range experienced by gas turbine technologies: 0.75–2.25% per year [2]. As with gas turbines and other aerodynamic rotating machinery, a portion of the unrecoverable loss could be attributed to gradual deterioration, such as fouling of the blades (which will impede the aerodynamic performance) and a gradual reduction in component efficiencies (gearbox, bearings, generator). These may not be recoverable by maintenance procedures, but only by component replacement.

A (potentially larger) contribution could come from availability declining with age, either because older turbines fail more frequently or because they take longer to bring back online. Possible reasons for the latter are the likelihood that older machines suffer more serious failures, difficulty in obtaining components for obsolete models, and operators being less likely to hold comprehensive maintenance contracts. Availability will depend on the amount of effort the owner is willing to invest in maintenance, which may naturally fall over time as the asset is paid down, and will depend on electricity prices and O&M costs. The manufacturer's availability warranty provided with a new turbine (which may for example guarantee 97% uptime) is also likely to exceed the standard provided by third party O&M providers in later life.

Early turbine death is a third contributing reason. If one turbine in a farm of four fails completely at age 17, the farm will continue operating at a maximum of 75% of its original load factor,

⁵ This is made up from the offset relative to the reference farm (2.926) plus the average of all farm site effects (−4.560).

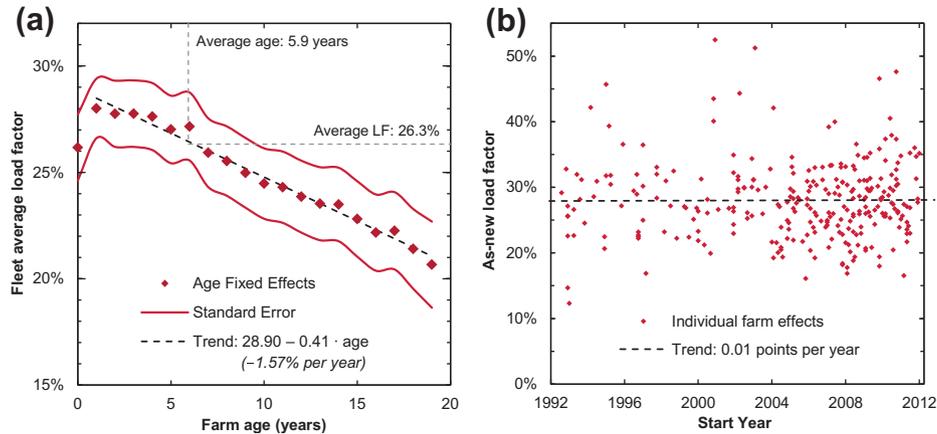


Fig. 11. Fixed effects from the regression of observed load factor against ideal load factors.

which would translate to an annualised degradation rate of around 1.6%.

Fig. 9a highlights these first two reasons as a gradual downwards slope in the bulk of weather-corrected load factor observations, and isolated periods of very low output due to downtime which are concentrated towards higher ages. Fig. 3a shows an example of the third: Blyth Harbour stands out at the lower-right of the chart as its load factor declined from 12% to just 2% between the ages of 12 and 17. By the end of its life only one of the nine turbines was generating, giving it the worst degradation rate of the farms we observed.

6.2. Future trends

What of the future? Although our main calculations are based on linear rates of decline, we cannot yet know if this assumption is correct. Figs. 9b and 11a suggest a slightly lower rate of ageing in the first six or seven years of a farm's life. This early period is typically covered by comprehensive warranties which guarantee near-maximal output, whereas older farms may be less intensively maintained and thus deteriorate more rapidly. The pattern seen in these Figures could be produced if all turbines (of whatever cohort) had an initial period of slow decline, followed by accelerating degradation as they become older.

Alternatively, it could be that advances in technology mean that recent turbines experience relatively little decline and will continue to do so in the future, whereas earlier models have declined relatively faster throughout their lives. As those early turbines were new in the 1990s their output is not available in our dataset, and so the experience of those years does not show up in Figs. 9b and 11a. If those early turbines had declined rapidly when new, the kink visible in these Figures would have been lessened or even removed.

Although the improvement in WCLF decline rates seen in farms commissioned after 2003 is statistically significant, that is no guarantee that this result will still hold if we were to repeat this study in a few years' time. Based on this, and further analysis in Supplementary Section 5.5, we must conclude that even though the initial signs are good, more data are required before we could say for sure that modern turbines are declining less rapidly than earlier cohorts.

6.3. Wider impacts

The cumulative lifetime output of a 100 MW wind farm with a 28.5% load factor would be 4.99 TWh over 20 years. If this farm

suffers a linear annual deterioration of -0.41 points after the first year, its lifetime output reduces to 4.37 TWh, a fall of 12.5%. This will increase the cost of electricity from wind generators, as less electricity is produced to recover the costs of construction. The economic value of the lost output is relatively low as it mostly occurs in the far future. With a discount rate of 10%, degradation increases the levelised cost of electricity by 9%, from approximately £90 [4,5] to £98 per MWh. This impact becomes greater if the economic lifetime increases or the discount rate decreases.

A second impact is that more capacity will need to be installed to produce a given level of output. The UK has a target for energy production from renewable sources (15% of all final energy by 2020), as opposed to a target for peak capacity. If turbine build rates peak in the coming years, the average age of the UK's wind farms will creep upwards, and so the output from a fixed capacity can be expected to decline. For every year the fleet ages, an additional 435 MW (4 large farms) would need to be brought online to maintain the original capacity of the UK's anticipated 30 GW fleet.

7. Conclusions

This paper demonstrates a generic and broadly applicable method for predicting a wind farm's monthly load factor, accounting for its location, hub height and the particular model of turbine installed. We use this to estimate the ideal monthly load factors for 282 of the UK's wind farms over the last decade, and compare these to the actual outputs over this period. This allows us to correct for the rapid improvement in wind turbine technology over the last two decades and the huge seasonal variability in wind speeds, thus revealing the subtle rate of degradation.

We find evidence of important, but not disastrous, performance degradation over time in a large sample of UK wind farms. When variations in the weather and improvement in turbine design are accounted for, we find that the load factors of UK wind farms fall by 1.57% (0.41 percentage points) per year. This degradation rate appears consistent for different vintages of turbines and for individual wind farms, ranging from those built in the early 1990s to early 2010s.

We use six methods of increasing complexity to find the following rates of degradation in absolute percentage points per year; and relative to the UK mean wind farm:

- Simple regression of all load factors against age (-0.44 ± 0.04 absolute) ($-1.69 \pm 0.17\%$ relative);
- Average trend in load factors for individual farms (-0.48 ± 0.36 absolute) ($-1.81 \pm 1.32\%$ relative);

- Correct for wind resource and regress weather-corrected load factor against age (-0.45 ± 0.01 absolute) ($-1.68 \pm 0.05\%$ relative);
- Trend in weather-corrected load factors for individual farms against their age (-0.45 ± 0.22 absolute) ($-1.70 \pm 0.82\%$ relative);
- Capacity-weighted fit to individual farms against their year of commissioning (-0.50 nominal, -0.36 weather-corrected) (-1.90% and -1.24% relative);
- Full fixed effects regression, accounting for site-specific wind speeds, turbine model and site quality (-0.41 ± 0.01 absolute) ($-1.57 \pm 0.06\%$ relative).

The combined average of these measures is -0.43 ± 0.05 percentage points per year, giving $-1.6 \pm 0.2\%$ annual degradation. The similarity of results from different methods gives us confidence that the underlying trend is robust: the decline in load factor with age is neither an artefact of systematic variation in wind speeds nor of the continual improvement in technology. Questions do however remain as to the exact form of this degradation, for example whether it is linear, quadratic or logarithmic with age; or how degradation rates are changing over time and whether they will be lower in the future. Access to data from more farms, and a more detailed wind resource assessment for each site will be fundamental to furthering our understanding of these issues.

The level of degradation we find is not insignificant, yet it is not unusual compared to conventional generation technologies. The fact that it has been omitted from calculations of the levelised cost of electricity from wind means that these estimates are around 9% below the true value (depending on assumed discount rate and economic lifetime). This is unlikely to be large enough to change the business case for wind power, but nonetheless it needs to be accounted for to give an accurate picture of its cost.

Acknowledgements

The authors would like to thank:

- Gordon Hughes, John Constable and the Renewable Energy Foundation, who enabled this research by generously making their data publicly available;
- Christopher Crabtree, Andrew Garrad and Paul Gardner (GL Garrad Hassan), Matthew Hanson (Pöyry), David MacKay (DECC), David Newbery, Jim Oswald (Oswald Consulting), Jim Platts and Peter Tavner for insightful discussions and feedback on our earlier drafts;
- The British Atmospheric Data Centre and NASA Goddard Flight Centre for the excellent services they provide;
- Anabelle Guillory from the Centre for Environmental Data Archival (CEDA) for help in acquiring the MIDAS marine data;
- Alasdair Skea and Simon Vosper from the UK Met Office for improving our understanding of the MIDAS equipment and data;
- The Alan Howard Charitable Trust and the EPSRC Grand Challenge Project 'Transforming the Top and Tail of the Energy Networks' (EP/I031707/1) for funding this work.

Appendix A. Supplementary data

Supplementary data related to this article can be found at <http://dx.doi.org/10.1016/j.renene.2013.10.041>.

References

- [1] Hughes G. The performance of wind farms in the United Kingdom and Denmark. London: Renewable Energy Foundation; 2012. URL: <http://tinyurl.com/cn5qnqg>.
- [2] Brooks FJ. GE gas turbine performance characteristics. GE Power Systems; 2000.
- [3] Jordan DC, Kurtz SR. Photovoltaic degradation rates—an analytical review. *Prog Photovoltaics: Res Appl* 2013;21(1):12–29.
- [4] MacDonald Mott. UK electricity generation costs update; 2010.
- [5] Arup. Review of the generation costs and deployment potential of renewable electricity technologies in the UK. London: Department of Energy and Climate Change; 2011.
- [6] Brinckerhoff Parsons. Electricity generation cost model—2011 update revision 1. London: Department of Energy and Climate Change; 2011.
- [7] Ribrant J, Bertling LM. Survey of failures in wind power systems with focus on Swedish wind power plants during 1997–2005. *IEEE Trans Energy Convers* 2007;22(1):167–73.
- [8] Tavner PJ, Xiang J, Spinato F. Reliability analysis for wind turbines. *Wind Energy* 2007;10(1):1–18.
- [9] Arwade SR, Lackner MA, Grigoriu MD. Probabilistic models for wind turbine and wind farm performance. *J Solar Energy Eng* 2011;133(4):041006–9.
- [10] Pöyry Energy Consulting. Implications of intermittency. Oxford. URL: <http://www.uwig.org/ImpactofIntermittency.pdf>; 2009.
- [11] Sinclair Knight Merz. Growth scenarios for UK renewables generation and implications for future developments and operation of electricity networks; 2008. URN 08/1021.
- [12] Green R, Vasilakos N. Market behaviour with large amounts of intermittent generation. *Energy Policy* 2010;38(7):3211–20.
- [13] Green R, Staffell I, Vasilakos N. Divide and conquer? k-means clustering of demand data allows rapid and accurate simulations of the British electricity system. *IEEE Trans Eng Manage* 2013 [in press].
- [14] Sturt A, Strbac G. A times series model for the aggregate GB wind output circa 2030. In: Presented at renewable power generation (RPG 2011), IET conference 2011.
- [15] García-Bustamante E, González-Rouco JF, Jiménez PA, Navarro J, Montávez JP. A comparison of methodologies for monthly wind energy estimation. *Wind Energy* 2009;12(7):640–59.
- [16] Kubik ML, Coker PJ, Barlow JF, Hunt C. A study into the accuracy of using meteorological wind data to estimate turbine generation output. *Renew Energy* 2013;51(0):153–8.
- [17] Kiss P, Varga L, Janosi IM. Comparison of wind power estimates from the ECMWF reanalyses with direct turbine measurements. *J Renew Sustain Energy* 2009;1(3):033105–11.
- [18] Hawkins S, Eager D, Harrison GP. Characterising the reliability of production from future British offshore wind fleets. In: Presented at IET conference on renewable power generation (RPG 2011) 2011. Edinburgh.
- [19] Kubik M, Brayshaw D, Coker P. Reanalysis: an improved data set for simulating wind generation?. In: WREF 2012. Denver, CO. URL: <http://tinyurl.com/c4ge72x>; 2012.
- [20] Ofgem. Electricity capacity assessment. Ref 126/12: London. URL: <http://tinyurl.com/97cuuxt>; 2012.
- [21] Renewable Energy Foundation. Analysis of wind farm performance in UK and Denmark [cited Dec 2012]; URL: <http://tinyurl.com/azu3m97>; 2012.
- [22] Lucchesi R. File specification for MERRA products. GMAO office note No. 1 (Version 2.3). NASA Goddard Space Flight Centre; 2012. URL: http://gmao.gsfc.nasa.gov/pubs/office_notes/.
- [23] 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(14):3624–48.
- [24] R Development Core Team. R: a language and environment for statistical computing. Vienna, Austria. URL: <http://www.R-project.org>; 2013.
- [25] Schallenberg-Rodriguez J. A methodological review to estimate technological wind energy production. *Renew Sustain Energy Rev* 2013;21(0):272–87.
- [26] Bañuelos-Ruedas F, Angeles-Camacho C, Rios-Marcuello S. Methodologies used in the extrapolation of wind speed data at different heights and its impact in the wind energy resource assessment in a region. In: Suvire GO, editor. Wind farm—technical regulations, potential estimation and siting assessment. Rijeka, Croatia: InTech; 2011. <http://dx.doi.org/10.5772/20669>.
- [27] Fox NI. A tall tower study of Missouri winds. *Renew Energy* 2011;36(1):330–7.
- [28] Nørgaard P, Holttinen H. A multi-turbine power curve approach. In: Presented at Nordic wind power conference 2004. Gothenburg, Sweden.
- [29] MacKay DJC. Could energy-intensive industries be powered by carbon-free electricity? Philosophical transactions of the Royal Society a: mathematical. *Phys Eng Sci* 2013;371(1986).
- [30] Harman K. Availability and operating efficiency – how well do wind farms actually perform?. In: Renewable UK conference. Glasgow: GL Garrad Hassan; 2010. URL: <http://tinyurl.com/adz94sp>.
- [31] Conroy N, Deane JP, Ó Gallachóir BP. Wind turbine availability: should it be time or energy based? – a case study in Ireland. *Renew Energy* 2011;36(11):2967–71.
- [32] Johnson C, Graves A, Tindal A, Cox S, Schlez W, Neubert A. New developments in wake models for large wind farms. In: AWEA windpower conference. Chicago, IL: Garrad Hassan America Inc; 2009. URL: <http://tinyurl.com/dxllkatz>.
- [33] Phillips JL, Cox SD, Henderson AR, Gill JP. Wake effects within and between large wind projects: the challenge of scale, density and neighbours—onshore

- and offshore. In: EWEC conference. Warsaw, Poland: GL Garrad Hassan; 2010. URL: <http://tinyurl.com/bm9nhk9>.
- [34] Barthelmie R, Larsen G, Pryor S, Jørgensen H, Bergström H, Schlez W, et al. ENDOW (efficient development of offshore wind farms): modelling wake and boundary layer interactions. *Wind Energy* 2004;7(3):225–45.
- [35] Ali M, Matevosyan J, Milanović JV. Probabilistic assessment of wind farm annual energy production. *Electric Power Syst Res* 2012;89(0):70–9.
- [36] Johnson C, Tindal A, Harman K, Graves A. Validation of energy predictions by comparison to actual performance. In: WINDPOWER conference. Garrad Hassan. URL: <http://tinyurl.com/aqkagoy>; 2008.