On the Performance of Wind Farms in the United Kingdom

David MacKay FRS
Cavendish Laboratory, University of Cambridge
and
Department of Energy and Climate Change, London

May 28, 2013 – Draft 6.0

Abstract

This paper identifies a significant flaw in a recent study "The Performance of Wind Farms in the United Kingdom and Denmark," published by the Renewable Energy Foundation, which claimed that wind farms in the UK wear out sooner than expected, and that recently-commissioned farms are substantially less efficient than older farms. The statistical model that underlies the method used in the study to infer the age-performance function of wind farms is *non-identifiable*, which means that no matter how much data is available, the age-performance function cannot be deduced by that model; the underlying model can fit the data in an infinite number of ways, with age-performance functions that fall or rise arbitrarily steeply. The method used in the study is believed to have resolved this non-identifiability in arbitrary ways; as a consequence, most of the conclusions of the Renewable Energy Foundation study are believed to be spurious.

There is nothing wrong with the valuable data that was presented in the Renewable Energy Foundation's study, however, and it will be possible to use that data – either with different models, or alongside additional data (for example, weather data) that resolve the non-identifiability.

Simple graphing of the data from the older farms suggests that their load factor has declined by roughly 2% per year, rather than the 5%, 6.5% or 12.8% per year asserted by the Renewable Energy Foundation. (All these percentages are fractional reductions per year.) How much of this downward trend should be attributed to intrinsic deterioration of the wind farm, and how much to a down-turn in the wind conditions over the period 2002–2011 is not resolved by my paper, but could readily be resolved by an approach that includes wind data.

The REF study claims that "normalized load factors" of wind farms decline to 15%, 8%, 13%, or 7% in the 10th year operation (with the number depending on the details of the fitting method). All these numbers appear inconsistent with the raw data which show that the actual load factors of 10-year-old farms are about 24% $\pm$ 7%. The raw data show that even 15-year old farms have actual load factors of about 24% $\pm$ 7%.

1 The study's remarkable claims

The Renewable Energy Foundation and Gordon Hughes have performed a valuable service by collating, visualizing, and making accessible a large database containing the performance of renewable generators in the UK, especially wind farms (`www.ref.org.uk`).

A recent study published by the Renewable Energy Foundation [Hughes, 2012] claimed that wind farms wear out sooner than expected, and that recently-commissioned farms are less efficient than older farms. Some specific conclusions of Hughes's paper that were found surprising by wind-industry experts were:

1. that the average load factor of onshore wind farms, adjusted for wind availability, declines significantly as they get older – Hughes's paper includes graphs showing declines in performance of roughly 6.5% per year or 12.8% per year, depending on the details of the fitting method (these rates of decline are based on the two graphs in Hughes's figure 10, which respectively show the normalized load factor declining from 24% in year 1 to 13% in year 10, which is equivalent to a decline of 6.5% per year over 9 years, or declining from 24% in year 1 to 7% in year 10, which is equivalent to a 12.8% decline per year over 9 years);

2. that the normalised load factor for UK onshore wind farms declines from a peak of about 24% at age 1 to 15% at age 10 and 11% at age 15 (since $15\% \simeq 24\% \times (1 - 0.051)^9$, I describe this decline from 24% to 15% in 9 years as being equivalent to a 5.1% decline per year; similarly, a decline from 24% to 11% in 14 years is equivalent to a decline of 5.4% per year, and a decline from 15% to 11% in 5 years is equivalent to a decline of 6.0% per year);

3. that the average normalised load factor of new UK onshore wind farms at age 1 (the peak year of operation) declined significantly from 2000 to 2011; and

4. that larger wind farms have systematically worse performance than smaller wind farms.

Not only are these claims surprising; they also seem inconsistent with the data on the load factors of actual wind farms. Surely, if it is true that "the normalised load factor for UK onshore wind farms declines from a peak of about 24% at age 1 to 15% at age 10 and 11% at age 15" then should one not expect that a the actual measured load factors of 10-year old wind farms would be around 15% at age 10 or 11, and around 11% at age 15? But the actual measured annual load factors at ages 10, 11, and 15, plotted directly from the same raw data used by Hughes in figure 1, are much larger. The actual load factors of 10-year-old wind farms are not 15% but $24\% \pm 7\%$ (based on 51 farms, figure 1a). The actual load factors of 11-year-old wind farms are not 15% but $24\% \pm 7\%$ (based on 47 farms, figure 1b). The actual load factors of 15-year-old wind farms are not 11% but $24\% \pm 7\%$ (based on 27 farms, figure 1c). Moreover, when looking at the load factors as a function of capacity in figure 1d, my eye struggles to perceive any evidence for the claim that "larger wind farms have systematically worse performance than smaller wind farms" (although no wind farms with capacity larger than 25 MW have yet reached these ages in the UK).

How can this paradox be resolved?

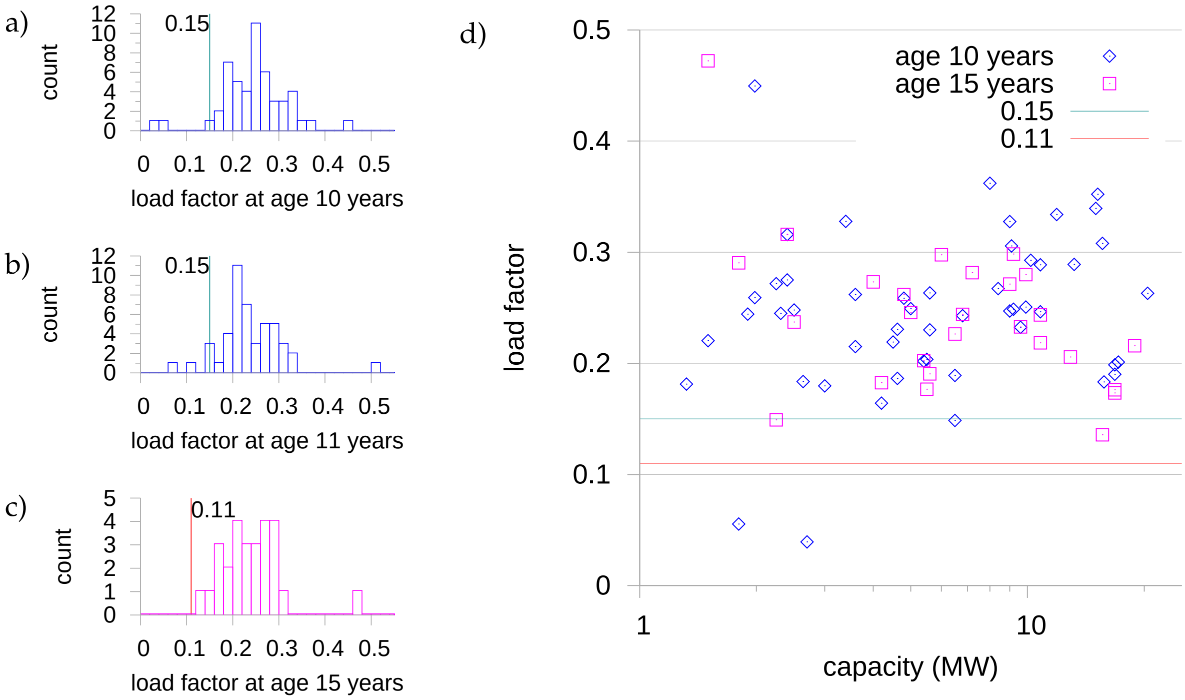

Figure 1: Actual load factors of UK wind farms at ages 10, 11, and 15.
a) Histogram of average annual load factors of wind farms at age 10 years. For comparison, the blue vertical line indicates the assertion from the Renewable Energy Foundation's study that "the normalised load factor is 15% at age 10."
b) Histogram of average annual load factors of wind farms at age 11 years.
c) Histogram of average annual load factors of wind farms at age 15 years. For comparison, the red vertical line indicates the assertion from the Renewable Energy Foundation's study that "the normalised load factor is 11% at age 15."
At all three ages shown above, the histogram of load factors has a mean and standard deviation of 24% $\pm$ 7%.
d) Average annual load factor versus capacity at age 10 and 15. For comparison, the horizontal lines indicate the assertions from the Renewable Energy Foundation's study that "the normalised load factor is 15% at age 10," and "the normalised load factor is 11% at age 15."

Multiplicative model
Hughes (2012),
Figure 10

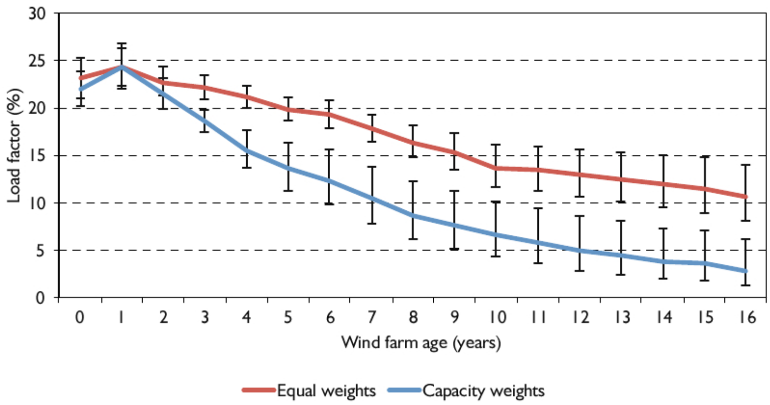

model type	data-weighting	year 1	year 10	rate of change (logarithmic)		rate of change (additive)
				$\frac{\log(l_{10}/l_1)}{9}$	$\left(\frac{l_{10}}{l_1}\right)^{\frac{1}{9}} - 1$	$\frac{l_{10}-l_1}{9}$
additive	equal⋆	24%	15%	−0.052	**−0.051**	−1.0
additive	capacity⋆	22.5%	8%	−0.115	**−0.109**	−1.6
multiplicative	equal	24%	13%	−0.068	**−0.066**	−1.2
multiplicative	capacity	24%	7%	−0.137	**−0.128**	−1.9

Figure 2: Four UK onshore age-performance curves from Hughes (2012) and Hughes (personal communication). The table summarises these results by extracting the rate of change of load factor between year 1 and year 10. This rate of change can be measured in three ways, and the highlighted column shows the rate of change measure I have used consistently in the present paper. Numbers in rows ⋆ are from Hughes Figures 1 and 2. For the five different fits of the model to the data, the discovered declines are equivalent to annual reductions of normalized load factor of 4.6%, 12.5%, 6.6%, 12.8%, and 5.1%.

To a statistician's eye, the most striking feature of the Renewable Energy Foundation's study is the *range* of different fits that Hughes [2012] reports, depending on minor changes of the model or the weighting of the data. As the table is figure 2 shows, the different fits that Hughes has reported correspond to annual reductions of normalized load factor of 4.6%, 12.5%, 6.6%, 12.8%, and 5.1%. This wild range of answers motivated the present investigation, which now offers an explanation for these apparently non-robust results.

2 The underlying model and its non-identifiability

The study by [Hughes, 2012] modelled a large number of energy-production measurements from 3000 onshore turbines, in terms of three parameterized functions: an age-performance function, which describes how the performance of a typical windfarm declines with its *age*; a wind-farm-dependent parameter describing how each *windfarm* compares to its peers; and a time-dependent parameter that captures national wind conditions as a function of *time*. The modelling method of Hughes is based on an underlying statistical model that is *non-identifiable*: the underlying model can fit the data in an infinite number of ways, by adjusting rising or falling trends in two of the three parametric functions to compensate for any choice of rising or falling trend in the third. Thus the underlying model could fit the data with a steeply dropping age-performance function, a steeply rising trend in national wind conditions, and a steep downward trend in the effectiveness of wind farms as a function of their commissioning date (three features seen in Hughes's fits). But all these trends are arbitrary, in the sense that the same underlying model could fit the data exactly as well, for example, by a less steep age-performance function, a flat trend (long-term) in national wind conditions, and a flat trend in the effectiveness of wind farms as a function of their commissioning date.

2.1 The underlying model

Hughes's work is based on "an error components model with fixed effects for each wind farm and each time period". I now define this "underlying" model, emphasizing that the model actually used by Hughes differs in a few ways, which will be detailed later.

The *additive* version of the underlying model expresses the load factor l_{it} of wind farm i during period t as:

$$l_{it} = f(A_{it}) + u_i + v_t + \epsilon_{it}, \tag{1}$$

where A_{it} is the age of farm i at period t, and $f(a)$ is the age-performance curve – the main topic of Hughes's paper. The "fixed effects" u_i and v_t describe respectively the excellence of windfarm i, and the excellence of the national wind-production conditions in period t.

The *multiplicative* version of the underlying model changes the left hand side thus:

$$\ln(l_{it}) = f(A_{it}) + u_i + v_t + \epsilon_{it}. \tag{2}$$

In either case, the fixed effects are 'normalized' by constraining

$$\sum_i u_i = 0; \text{ and } \sum_t v_t = 0. \tag{3}$$

Both versions of the underlying model (1, 2) are non-identifiable. *If the fitted model's age-performance curve $f(a)$ is modified by adding a linear function with arbitrary slope, and if the fitted 'windfarm effects' u_i and 'period effects' v_t are adjusted proportionally as detailed below, then the quality of fit of the model to the data is exactly unchanged;* therefore the procedure of fitting this underlying model to the data has an important degree of arbitrariness – the age-performance curve can actually go up or down, as steeply as you like, if the other parameters of the model are adjusted appropriately. The loss function used for model-fitting by Hughes [2012] was the weighted sum of squares of the residuals; the non-identifiability applies for any loss function.

I now prove this statement mathematically, for the multiplicative model. Assume that the function $f(a)$, the windfarm-parameters $\{u_i\}$ and the period-parameters $\{v_t\}$ have been adjusted to particular values that fit the data as described by the residuals $\{\epsilon_{it}\}$ in equation (2), repeated here:

$$\ln(l_{it}) = f(A_{it}) + u_i + v_t + \epsilon_{it}.$$

Let the birthdate of wind farm i be b_i, and let c_t be the time associated with the period whose index is t, so that a wind farm's age at time t is $A_{it} = c_t - b_i$, so we have:

$$\ln(l_{it}) = f(c_t - b_i) + u_i + v_t + \epsilon_{it}. \tag{4}$$

Now consider changing the age-performance function f to

$$\tilde{f}(a) = f(a) - D \times (a - \alpha), \tag{5}$$

where D is an arbitrary constant (the added 'slope'), and α is a constant to be defined in a moment; and changing the windfarm-parameters $\{u_i\}$ to

$$\tilde{u}_i = u_i - D \times (b_i - \beta); \tag{6}$$

and changing the period-parameters $\{v_t\}$ to

$$\tilde{v}_t = v_t + D \times (c_t - \gamma), \tag{7}$$

where β and γ are set to the means of $\{b_i\}$ and $\{c_t\}$

$$\beta = \sum_{i=1}^{I} b_i / I \text{ and } \gamma = \sum_{t=1}^{T} c_t / T. \tag{8}$$

These final two definitions ensure that the sums of u and v are unchanged:

$$\sum_i \tilde{u}_i = \sum_i u_i; \quad \sum_t \tilde{v}_t = \sum_t v_t, \tag{9}$$

so the normalization constraint (3) still holds. When these changes are made, what are the new residuals, $\{\tilde{\epsilon}_{it}\}$? They are defined by

$$\ln(l_{it}) = \tilde{f}(c_t - b_i) + \tilde{u}_i + \tilde{v}_t + \tilde{\epsilon}_{it}. \tag{10}$$

Substituting for (5), (6), and (7),

$$\ln(l_{it}) = f(c_t - b_i) - D \times (c_t - b_i - \alpha) + u_i - D \times (b_i - \beta) + v_t + D \times (c_t - \gamma) + \tilde{\epsilon}_{it}, \tag{11}$$

then cancelling terms

$$\ln(l_{it}) = f(c_t - b_i) - D \times (\cancel{c_t} - \cancel{b_i} - \alpha) + u_i - D \times (\cancel{b_i} - \beta) + v_t + D \times (\cancel{c_t} - \gamma) + \tilde{\epsilon}_{it}, \tag{12}$$

we find

$$\ln(l_{it}) = f(c_t - b_i) + u_i + v_t + \tilde{\epsilon}_{it} + D \times (\alpha + \beta - \gamma). \tag{13}$$

Thus, if we define

$$\alpha = \gamma - \beta \tag{14}$$

then comparing the new residuals' equation (13) with the old residuals' equation (4), we find that

$$\tilde{\epsilon}_{it} = \epsilon_{it}, \tag{15}$$

that is, the fit of the model to the data is unchanged.

Thus if the "fixed effects" may be freely varied, steeply increasing or steeply decreasing age-performance curves $f()$ can be arbitrarily created with this model – the fixed effects can be adjusted to exactly compensate for *any* linear addition to $f(a)$.

If the constraints (3) are relaxed then the number of unconstrained degrees of freedom increases from the one (D) identified here to two or three, with the additional free parameters being the β and/or γ that appear in equations (6) and (7) – they no longer need to satisfy the constraints in equation (8); given β and γ, however, the parameter α remains constrained by (14).

2.2 Relationship between the theoretical result above and Hughes's work

How was the non-identifiability resolved in Hughes's implementation? I do not have access to the source code, so I am not certain, but three possible resolutions can be imagined. Either:

1. the implementation was also of a non-identifiable model, and the outcome of the fitting therefore depends on spurious details such as the initialization of the model's parameters, or other choices made by the programmer; or

2. one or more additional ad-hoc constraints, not mentioned in the model definition above, were included and the free parameter D was thus pinned down to a particular value that depends on the details of the constraints chosen; or

3. minor differences between the underlying model described above and the model actually implemented lead to the model actually implemented being weakly-identified rather than non-identifiable.

Hughes (personal communication) has indicated to me that he believes the third The precise model used by Hughes [2012] was slightly different in a couple of ways from the one defined in section 2.1 in this paper. First, whereas my model definition assumed that $f()$ is a function of time represented with the same resolution as the times $\{c_t\}$ at which data are obtained, Hughes [2012] used a function $f()$ whose argument was rounded to an integer number of *years* and the data are given at *monthly* intervals. Since the degree of freedom I identified in the underlying model is a long-term linear trend $D \times (a - \alpha)$, I don't imagine this time-resolution issue is of crucial significance, but it is certainly a detail that would be sufficient to turn a non-identifiable model into a (weakly) identifiable model. Second, my specification included the normalization constraint (3), which I found in the model definition on the top line of page 27 of Hughes [2012], and my proof was careful to ensure that this constraint was satisfied for any setting of the free parameter D. Richard Green has pointed out that on page 31, Hughes states that another normalization constraint was used, namely a particular v_t, for the period of August 2007, was pinned to zero "in order to ensure that

$$\sum_i v_i \simeq 0.\text{''} \tag{16}$$

We infer that the constraints (3) were not necessarily enforced. The introduction of the constraint

$$v_{\text{August 2007}} = 0 \tag{17}$$

and the relaxation of the constraints (3) implies that Hughes [2012] actually used a slightly different model from the one I have analysed here; so the proof of non-identifiability presented here does not precisely apply to the model actually used by Hughes [2012]. Nevertheless, the fundamental issue identified here must be relevant to the results in Hughes [2012]. If the model actually used in Hughes [2012] *is* in fact identifiable, thanks to the years-versus-months time-resolution issue, or thanks to the inclusion of ad-hoc constraints such as (17), then an instructive way to view those details or ad-hoc constraints is that they are providing an arbitrary way of determining the degrees of freedom D, β, and γ identified in the present paper. Of these three, D is the significant one. I do not think that the use of ad-hoc constraints such as (17) is a sensible way of inferring the parameter D, which has, through equation (5), such a radical effect on the key quantity of interest, the age-performance function $f()$. So my overall conclusion remains: the model used by Hughes [2012] *either* is non-identifiable and therefore produces spurious results; *or* is identifiable thanks to minor details involving different granularities of representation of time, or other arbitrary constraints, and produces results that depend spuriously on those minor details or constraints.

2.3 Predictions

We can make testable predictions from the hypothesis that the underlying model's non-identifiability is leading to spurious results.

This hypothesis predicts that when a particular fit asserts that the age-performance function drops steeply, the fitted model will also assert a steeply rising trend in national wind conditions, and a steep downward trend in the effectiveness of wind farms as a function of their commissioning date; and the greater the steepness of the drop in the age-performance function, the greater will be the rising trend in national wind conditions and the downward trend in the effectiveness of wind farms as a function of their commissioning date.

3 What do the data show?

The data *can* shed light on the age-dependence of wind-farm performance in two ways. First, we can change the model by introducing additional information, constraints, or regularizers, or we can add complementary data that gives independent information about some of the parameters. For example, if we believe that the weather effect v_t does *not* have a large long-term linear trend (perhaps on the basis of additional wind-speed data from independent sources); or if we believe that the windfarm-excellence parameters u_i are unlikely to depend very strongly on the birth-date of the windfarm, then there is scope to pin down the otherwise-unconstrained degree of freedom "D". Iain Staffell and Richard Green (personal communication, and paper submitted for publication) have added complementary data giving independent information from NASA about local wind speeds throughout the duration of the data, and their paper will therefore give a definitive answer about the age-dependence of wind-farm performance.

Second, we can also investigate the question "how does a wind farm's load factor vary with its age?" empirically by simply *plotting the wind farm's load factor as a function of its age*. The graphs that follow focus on a subset of 94 farms in the dataset collated by Hughes [2012], namely the farms that were in existence in 2004. (This subset was chosen since it selects the oldest farms and the farms for which the longest runs of data are available.) I computed the load factor over three-month periods; Figure 3 shows these raw data as a function of wind-farm age on twelve separate graphs, one for each month (with "January" denoting the three-month period November–December–January, and so forth). Each graph also shows a simple least-squares fit of an exponential curve, and the slopes of those fits are displayed in table 1. (Data from the first 12 months of operation of a farm were omitted from the least-squares fit. Each data point was given a weight proportional to the capacity of the wind farm.)

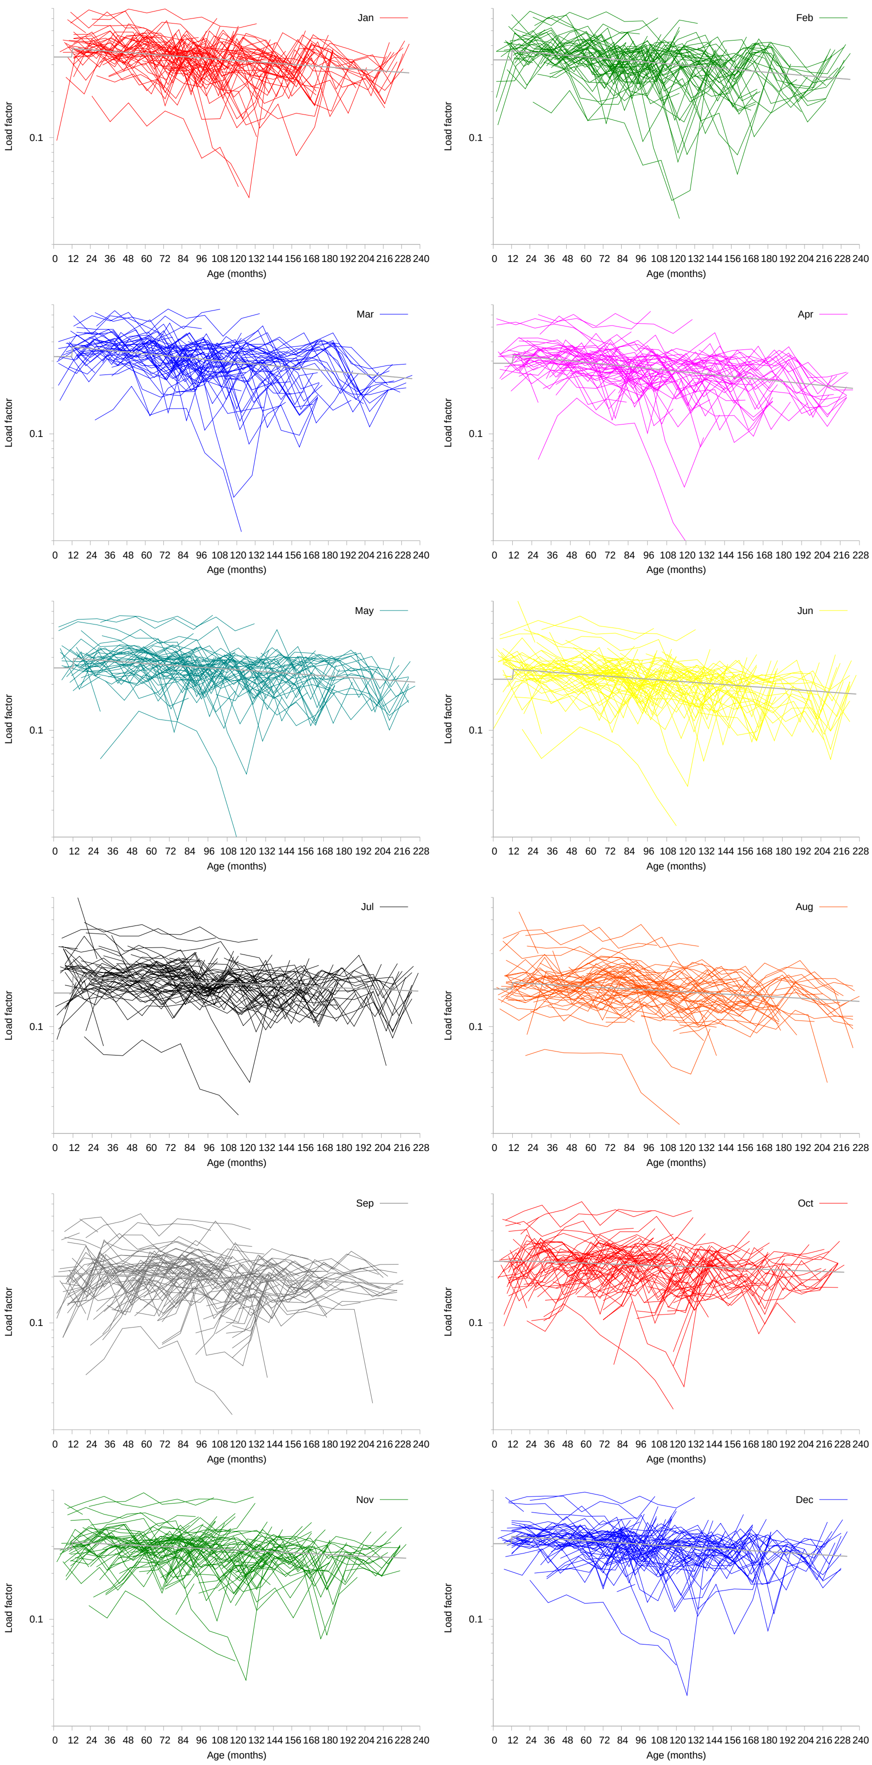

Figure 3: Load-factors of wind farms (averaged over three-month periods) versus their age. The straight grey lines show an exponential curve, fitted by least-squares to the data from age 12 months onwards, with each data point weighted by the capacity of the corresponding farm. The farms shown here are restricted to the ones that were in existence in 2004.

month	decline rate (%/year)	load factors at 1 year	load factors at 10 years
1	1.98	0.382	0.320
2	2.34	0.371	0.300
3	2.57	0.371	0.294
4	2.83	0.327	0.254
5	2.00	0.295	0.246
6	2.09	0.250	0.207
7	1.20	0.212	0.190
8	1.57	0.194	0.168
9	0.90	0.211	0.195
10	0.97	0.256	0.235
11	1.32	0.319	0.284
12	1.64	0.348	0.301
average		0.295	0.249

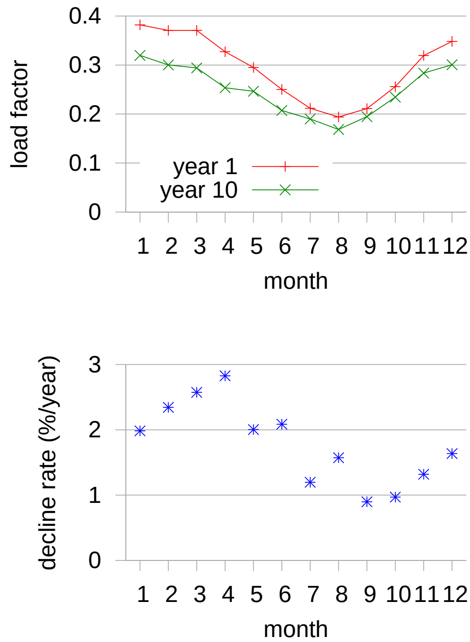

Table 1: Fitted parameters of exponential curves in figure 3. In the table and graphs, "month = 1" denotes the three-month period November-December-January. "month = 2" denotes December-January-February, and so forth. The upper graph shows the fitted load factor in years 1 and 10, and the lower graph shows the rate of decline (the *fractional* percentage decrease, per year). According to these fits, the average (annual) load factor decreases by 1.85% per year, from roughly 29.5% at year 1 to roughly 25% at year 10.

According to these fits, the late-winter load factors start higher and decrease slightly faster than the late-summer load factors. The average (annual) load factor decreases from roughly 29.5% at year 1 to roughly 25% at year 10 (approximately equivalent to a decline rate of 1.85% per year).

Hughes [2012] reported two fits of his model to the data, one with the datapoints weighted by capacity (as above), and another with the data equally weighted. Similarly I also made an equally-weighted fit to the data. The results are shown in figure 4 and table 2 in the appendix. The change in weighting of data-points changes the fit only slightly: the fitted year-1 load factors are slightly higher, and the year-10 load factors are essentially unchanged. This slight change in fit points in the opposite direction to the assertions made in Hughes [2012], who suggested that large windfarms age *more* rapidly than small windfarms – in fact the slight change in fit to the raw load factor data suggests that the larger windfarms age slightly *less* rapidly. These results motivate further work to explore the capacity-dependence of wind-farm performance.

The appendix of this paper shows, alongside the age-dependent data in figure 3, another display of the same data, as a function of calendar date instead of wind-farm age.

The simple plots and fits I have shown here are not the final answer to the question of wind-farm aging. I have only used data from the 94 oldest farms. Some of the 94 farms plotted here perform better than the identified trend, and some worse. I am sure that more-sophisticated models will be applied to this excellent data set and further interesting and useful insights will emerge.

Acknowledgements

I thank Gordon Hughes, Richard Green, David Newbery, David Milborrow, Ramesh Ghiassi, Alex Davies, and James Lloyd for helpful discussions.

References

G. Hughes. The performance of wind farms in the United Kingdom and Denmark. Technical report, Renewable Energy Foundation, 21 John Adam Street London WC2N 6JG, 2012.

A Appendix: Raw data

In all graphs:

- data points that were omitted from Hughes's [2012] analysis are also omitted here;
- and Hughes's [2012] corrections to capacity are applied.

The graphs show the load factor on a logarithmic scale. Lines with slopes 0%, −7%, and −14% per year are shown to aid comparison with the fits reported in figure 10 (page 33) of Hughes [2012], which show, from year 1 to year 10, declines averaging 6.8% per year and 13.7% per year, depending on the weights used when fitting the data.

The graphs labelled "January" show data where load factors for November, December, and January have been averaged together.

Since all the farms shown were in existence in 2004, if all farms tend to deteriorate at x% per year, then one would expect any farm's graph of load factor versus **age** to have a slope of x% per year (unless the weather has systematically drifted with time), and all the graphs of load factor versus **year** also to have a slope of x% per year (unless the weather has systematically drifted with time).

A.1 January data (smoothed 3 months)

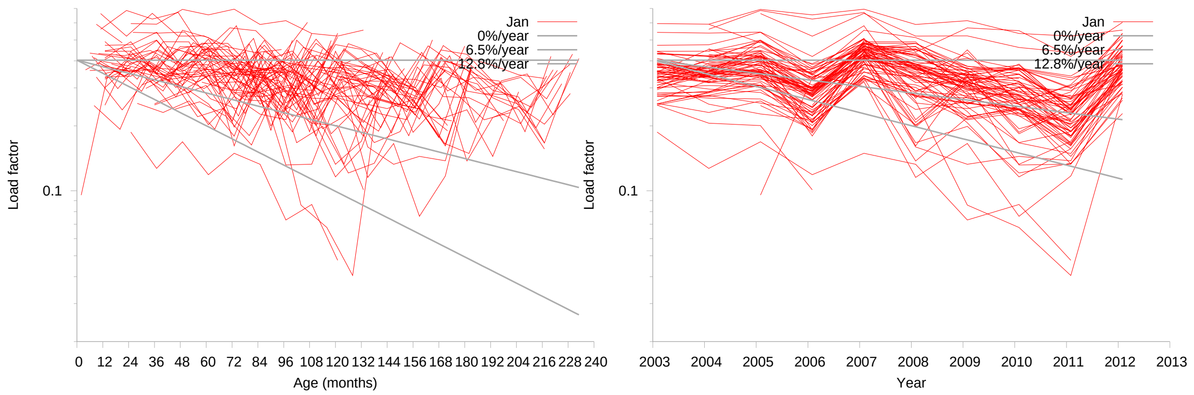

A.2 February data (smoothed 3 months)

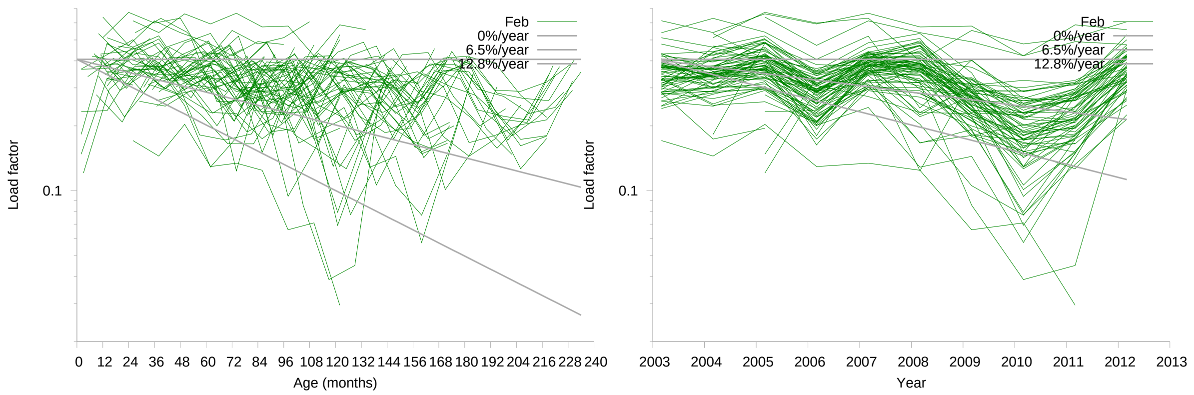

A.3 March data (smoothed 3 months)

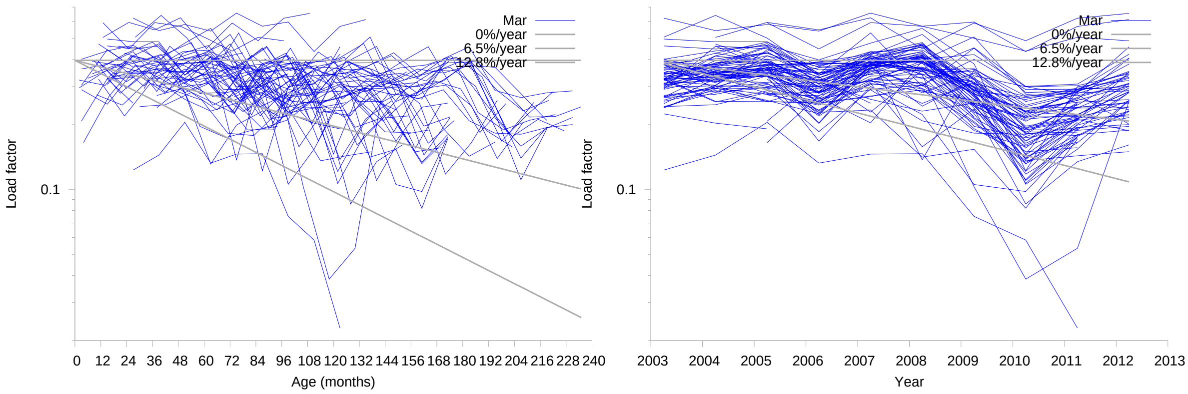

A.4 April data (smoothed 3 months)

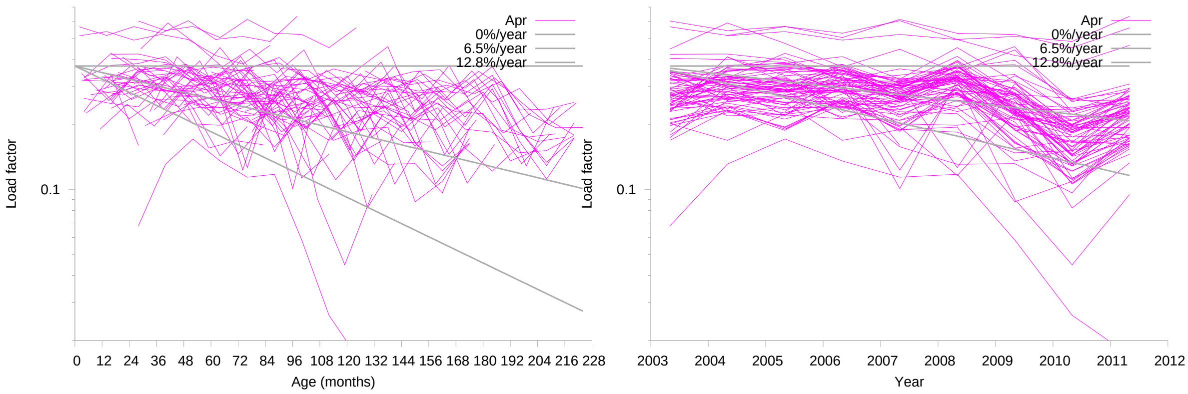

A.5 May data (smoothed 3 months)

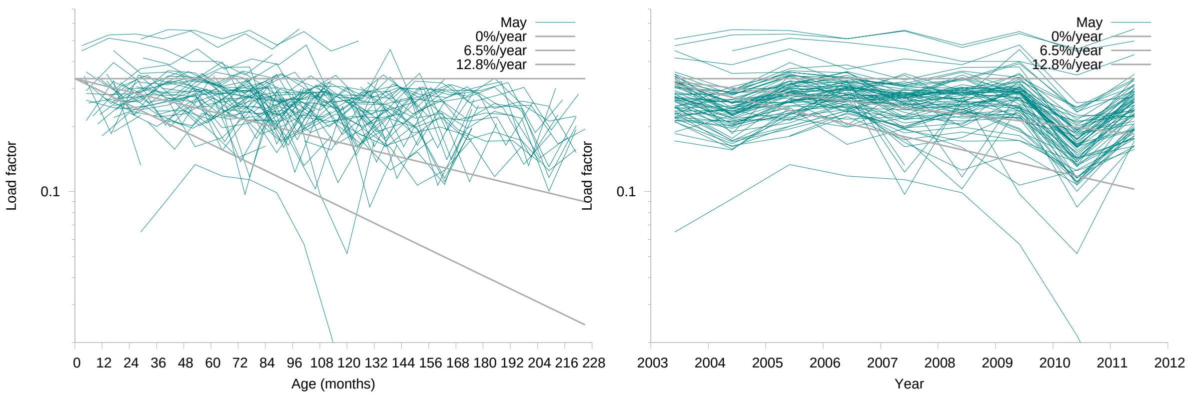

A.6 June data (smoothed 3 months)

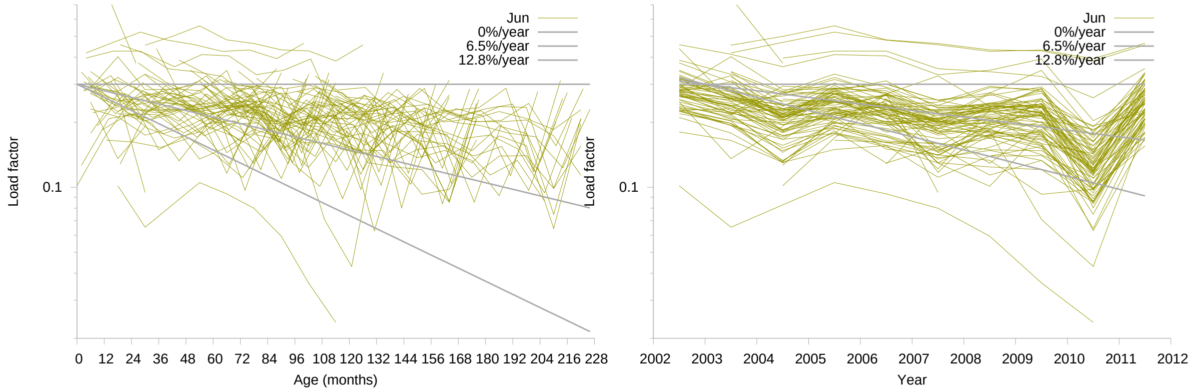

A.7 July data (smoothed 3 months)

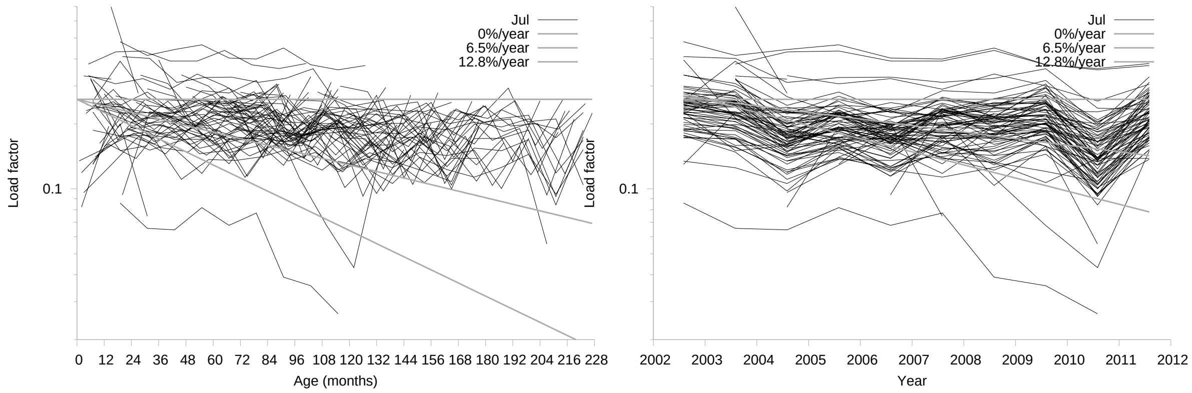

A.8 August data (smoothed 3 months)

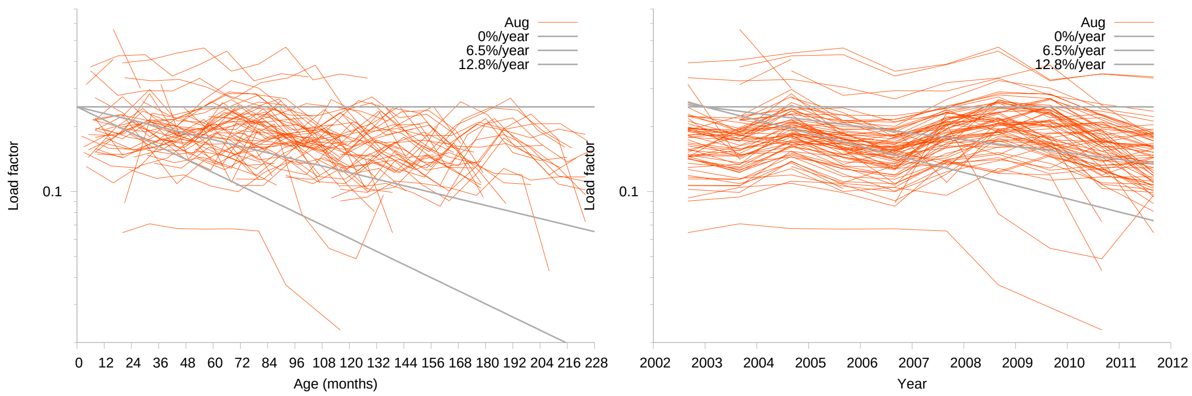

A.9 September data (smoothed 3 months)

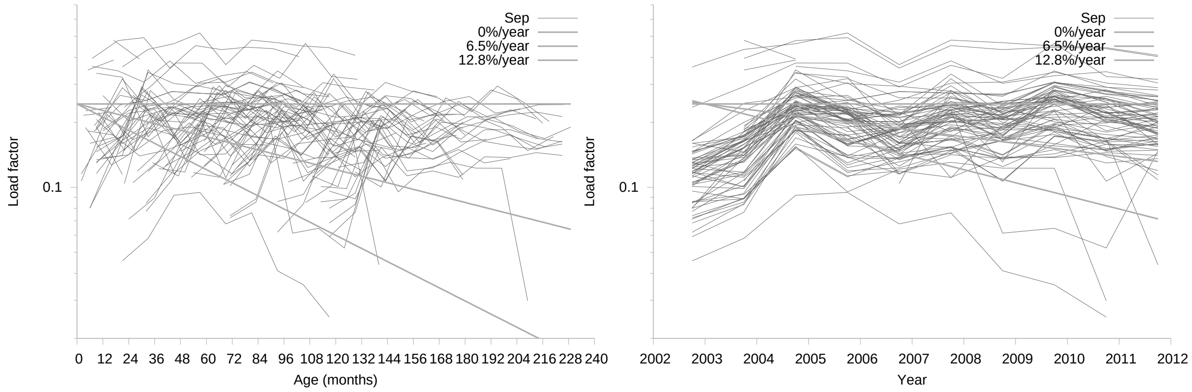

A.10 October data (smoothed 3 months)

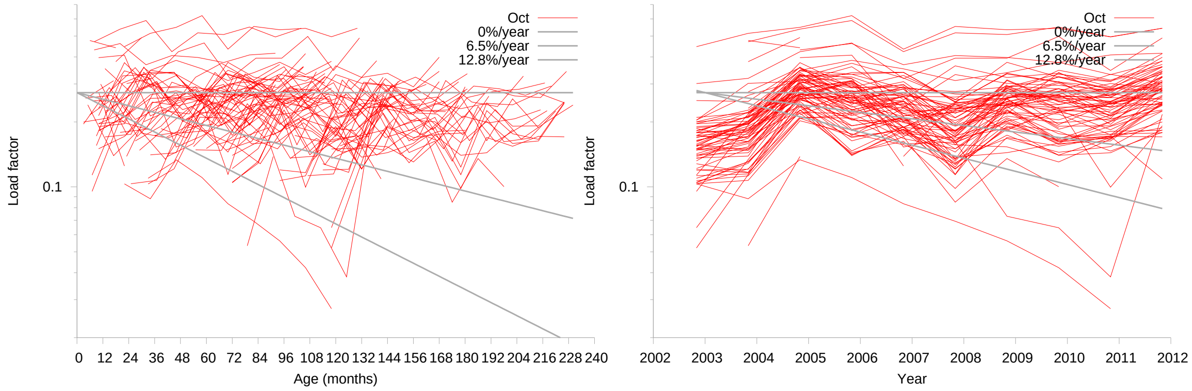

A.11 November data (smoothed 3 months)

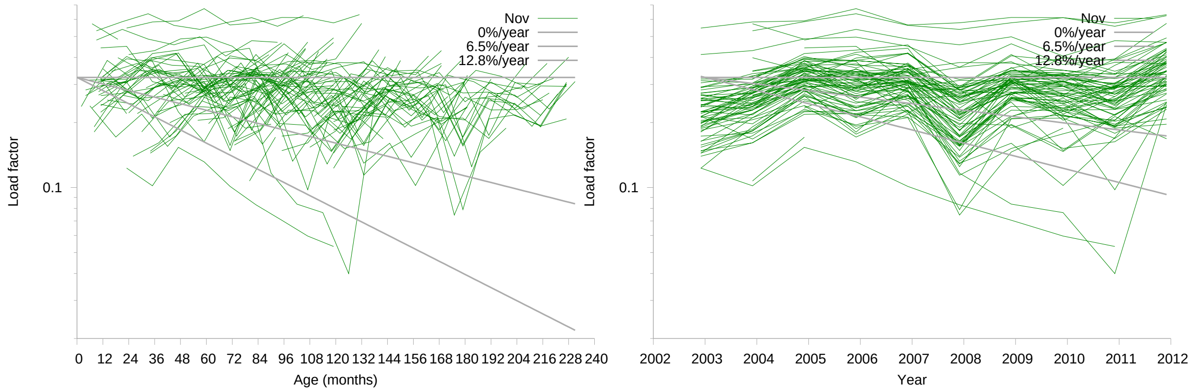

month	decline rate (%/year)	load factors at 1 year	10 years
1	2.32	0.397	0.322
2	2.79	0.392	0.304
3	3.04	0.391	0.297
4	3.18	0.341	0.256
5	2.41	0.308	0.248
6	2.80	0.271	0.210
7	2.21	0.237	0.194
8	2.38	0.215	0.173
9	1.36	0.223	0.197
10	1.38	0.266	0.235
11	1.61	0.324	0.281
12	1.85	0.355	0.301
average		0.310	0.252

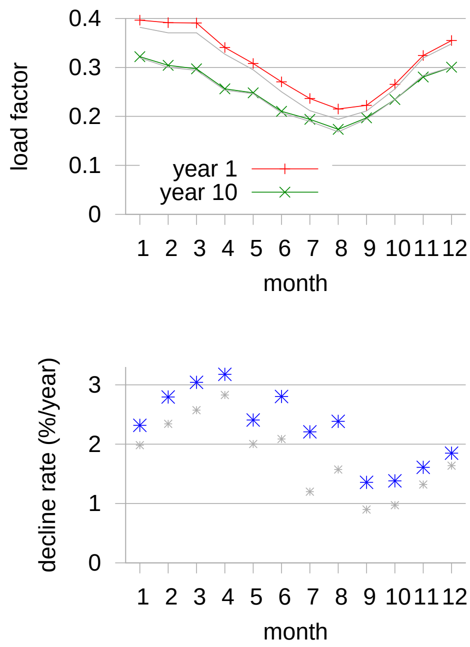

Table 2: Fitted parameters of the exponential curves in figure 4. In the table and graphs, "month = 1" denotes the three-month period November-December-January. "month = 2" denotes December-January-February, and so forth. The upper graph shows the fitted load factor in years 1 and 10, and the lower graph shows the rate of decline (the *fractional* percentage decrease, per year). In both graphs the grey curves or points show the capacity-weighted fits from table 1.

A.12 December data (smoothed 3 months)

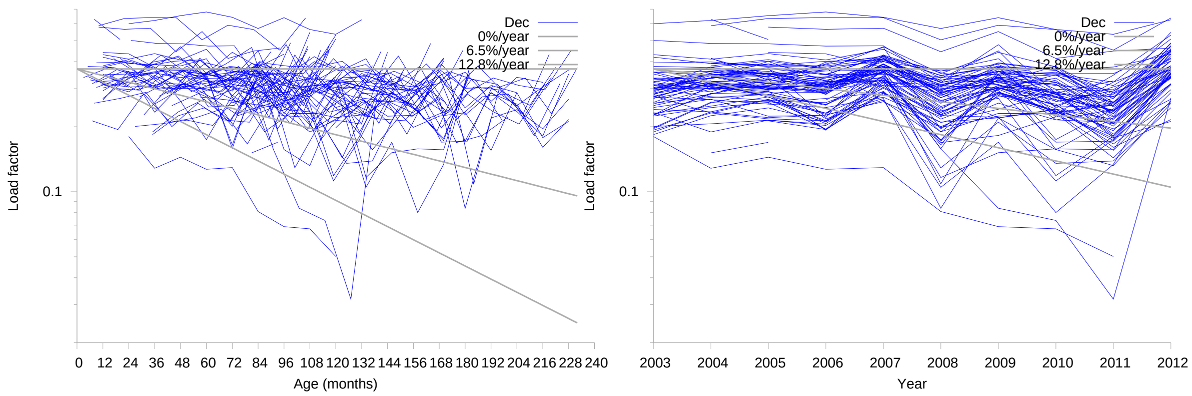

A.13 Plots of actual load factors on birthdays

See figure 5, figure 6 and figure 7.

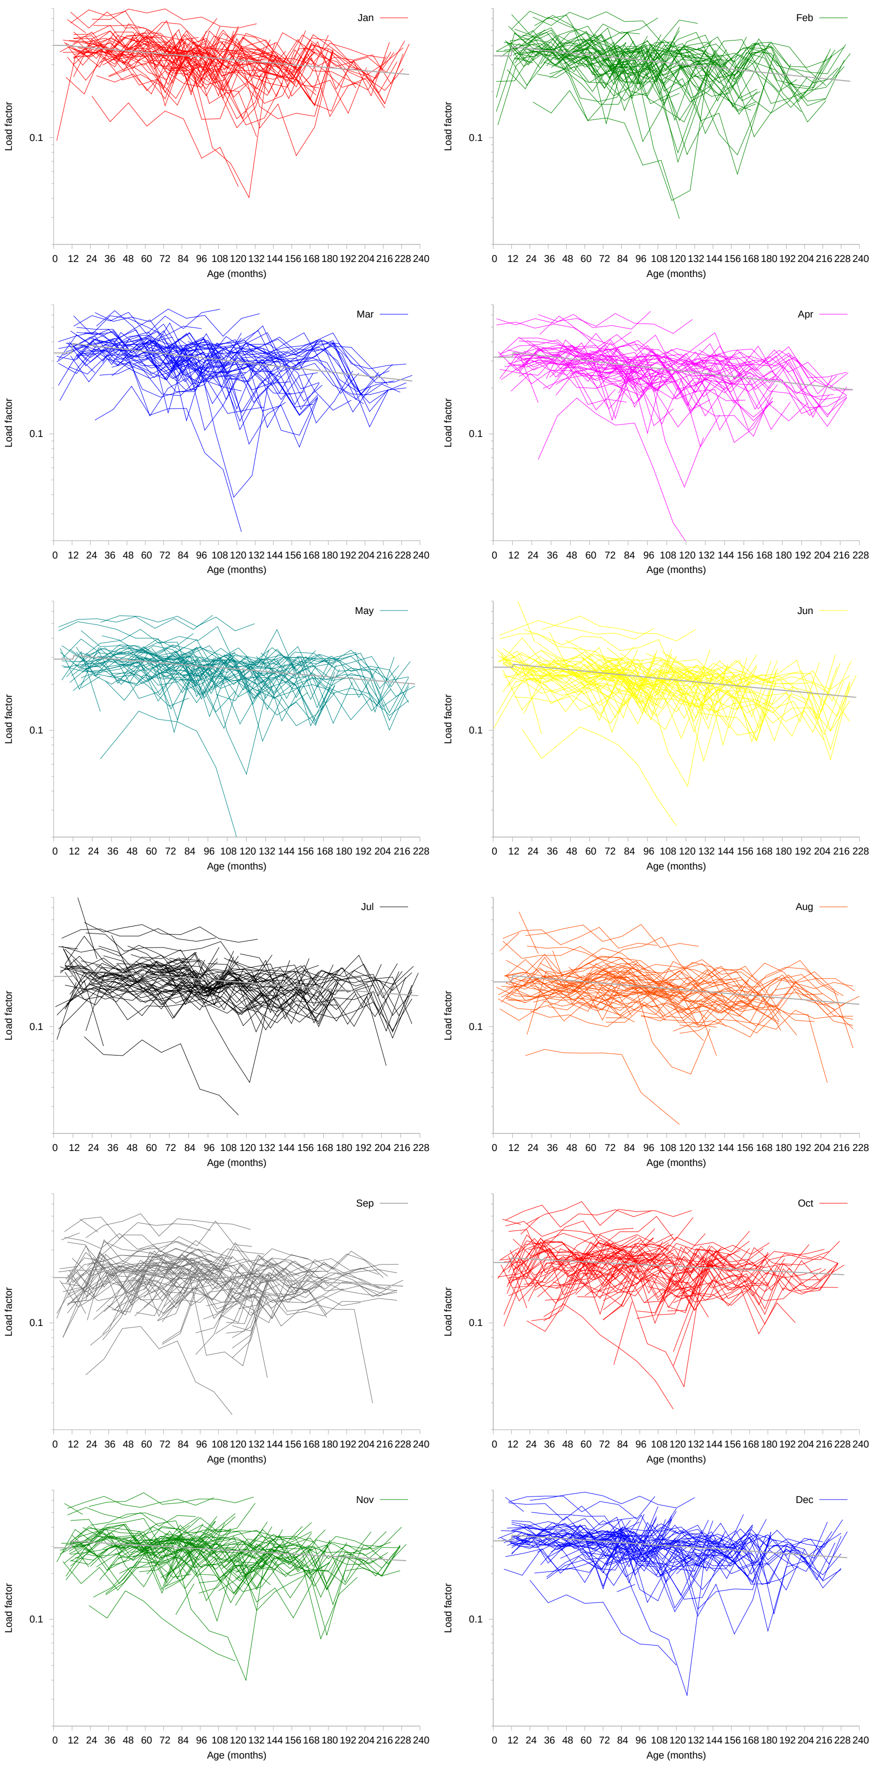

Figure 4: Load-factors of wind farms (averaged over three-month periods) versus their age. The grey lines show least-squares fits of exponential curves, using equal weighting of data points rather than capacity weighting.

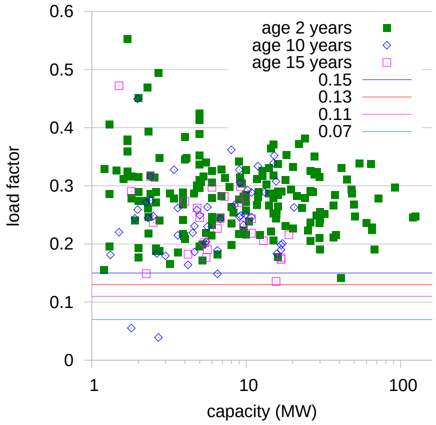

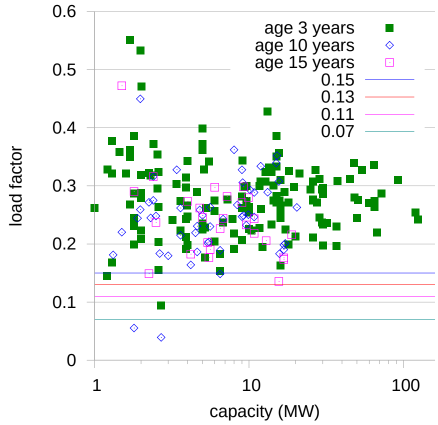

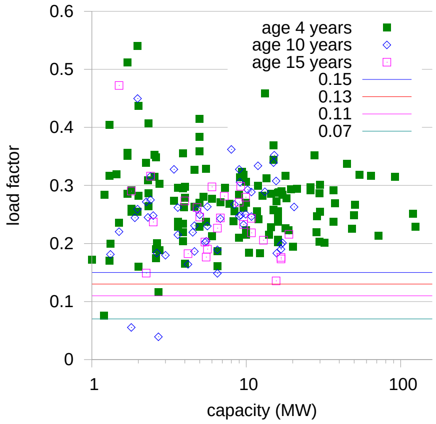

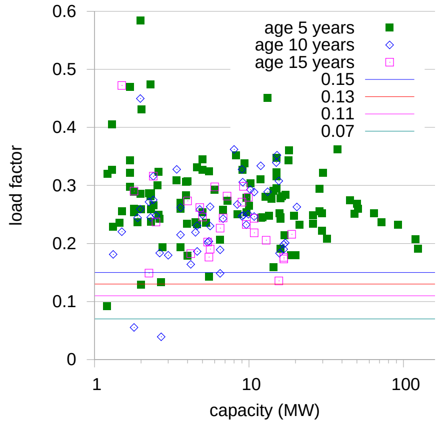

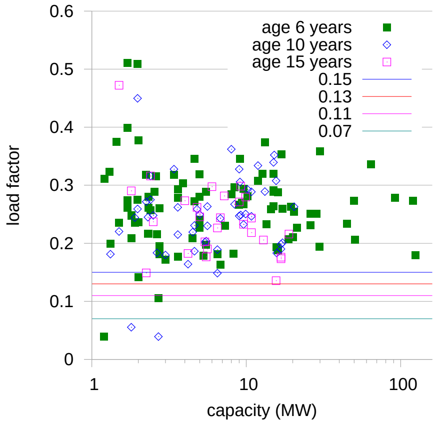

Birthday	Load factor – mean and standard deviation equal-weighted	capacity-weighted
2	28.2±6.4	26.9±4.9
3	27.8±6.7	26.6±5.0
4	27.4±6.9	26.7±5.0
5	27.4±7.0	25.5±5.3
6	26.3±7.1	25.7±5.4
10	24.5±7.1	25.4±6.1
11	23.7±6.9	24.6±5.5
15	24.1±6.7	22.3±5.1

Figure 5: Load-factors of wind farms on their 2nd, 3rd, 4th, 5th, and 6th birthdays, versus their capacities. For these graphs, all UK wind farms have been included. The data for 10th and 15th birthdays from figure 1 are also shown.

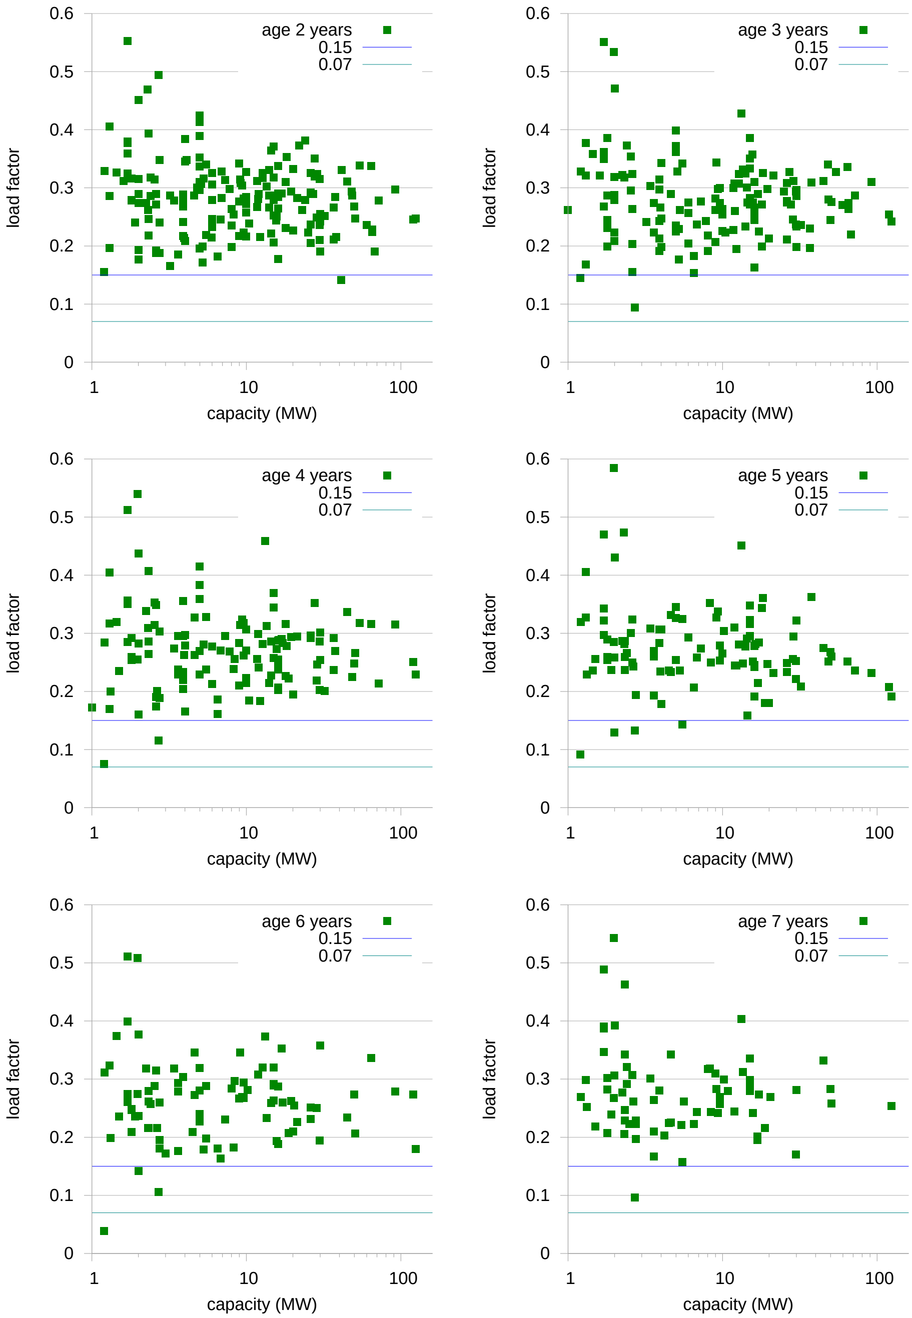

Figure 6: Load-factors of wind farms on their 2nd, 3rd, 4th, 5th, 6th, and 7th birthdays, versus their capacities. For these graphs, all UK wind farms have been included.

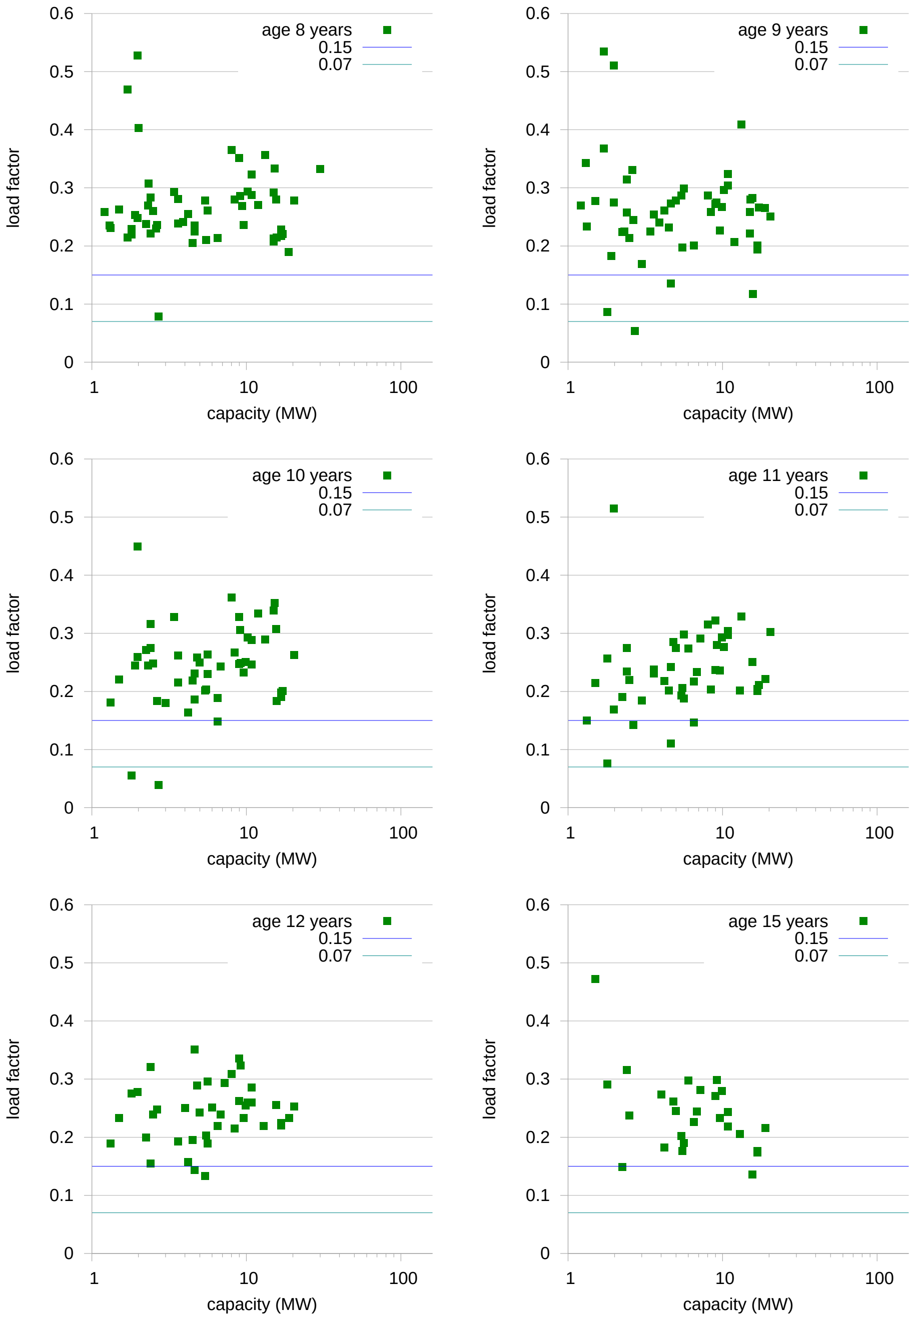

Figure 7: Load-factors of wind farms on their 8th, 9th, 10th, 11th, 12th, and 15th birthdays, versus their capacities. For these graphs, all UK wind farms have been included.