Black Start from Non-Traditional Generation Technologies

Network Innovation Allowance
June 2019

Case study:
Wind variability

In partnership with:

tnei

nationalgridESO

Contents

Executive summary

TNEI Services Ltd (TNEI) was commissioned by National Grid ESO to investigate the capability of non-traditional technologies in the restoration of the GB power system in the event of a partial or total system shutdown. The project is a Network Innovation Allowance (NIA) project initiated by National Grid ESO, with support from SP Energy Networks. The overall aim of the NIA project was to provide insight into the capability of several prevalent non-traditional technologies: wind, solar, storage, demand side response (DSR) and electric vehicles (EV), to provide ancillary services to National Grid ESO in the event that the GB network requires a Black Start.

Responding to the significant changes in the energy landscape in the past decade, National Grid ESO are seeking to understand how renewable generation and distributed energy resources (DER) could facilitate the restoration of the GB power system with the decline and decommissioning of traditional Black Start providers (larger, synchronous power stations). The creation of smaller, distributed power islands is of particular interest as a result, whereby these would be initiated on distribution networks and grow to energise the transmission network. This project has considered the technical capability of the technologies, the challenges of creating and maintaining small power islands with high penetrations of renewables and DER, and how to better predict the reliability and availability of renewable generation in such a scenario.

The project has three distinct deliverables:

- **Report 1:** Overview of the capability of non-traditional technologies to provide Black Start and restoration services;
- **Report 2:** Investigation of the challenges around power system strength and stability specifically in relation to power islands with high penetrations of renewables and converter-based technology; and
- **Report 3:** A sophisticated planning tool specifically designed to simulate distributions for the reliable output of wind over periods of hours to days, and how these distributions vary on timescales of months and years.

This report is Report 3, one of the three deliverables from the “Black Start from Non-Traditional Technologies” project. It is concerned with characterisation of the highly variable nature of the resource for non-dispatchable renewable generation, and wind in particular. It presents the results of extensive quantitative assessment and statistical modelling of wind variability, conducted to explore and describe the extent to which wind generators have the capacity to provide the reliable and sustained power outputs necessary for establishing and growing power islands.

The analysis in this report suggests that wind farms could be able to displace some conventional black start providers in investment planning timescales. For typical combinations of wind farms, large volumes of capacity would be required to achieve this, with some specific combinations requiring much less than others, where favourable geographic diversity ensures a higher level of minimum output.

Furthermore, the analysis suggests that wind could be very effective for displacing traditional providers in operational planning timescales, i.e. where the ability of the procured wind generation portfolio to deliver a stable minimum output in the short term can be accurately predicted. This could be very significant for the ESO given the high cost and carbon impact associated with guaranteeing the readiness of conventional out-of-merit thermal power stations for Black Start.

These results are consistent with the very well-established view that wind generation has a ‘capacity credit’ in the sense of being able to maintain system reliability (as measured by loss of load expectation or expected energy unserved) when replacing conventional generation capacity. More specifically, the results demonstrate that the concept of a positive capacity credit for wind also holds true in different contexts – albeit defined slightly differently, notably Black Start. This is also, almost certainly, applicable to other aspects of system security including procurement of system services such as frequency response and reserves.

Combining the output of wind with electricity storage and solar power could enable even greater displacement of traditional service provision, although this hasn’t been considered in detail in this analysis.

National Grid ESO | July 2019 Black Start from Non-Traditional Generation Technologies

The detailed findings of the report are:

1 Aggregations of wind generators can reliably deliver a sustained constant level of power, even for relatively long periods of up to 5 days, and even where the level of certainty required that the specified power level will in fact be consistently delivered is 99.9 per cent. However, that level of power decreases significantly as the duration is extended to days, and even more rapidly as the required reliability is increased from 50 per cent to 99.9 per cent, with values being very small in many of the combinations of duration and risk that were tested.

An example of this is shown in table 0.1 and figure 0.1 below which lists the average number of MW of wind capacity that would need to made available in advance (i.e. procured as part of the service) in order secure a sustained 500 MW of power. The averaging is across 100s of possible combinations of procured wind farms, and all possible start-times (and days of the year) are included in the quantile calculations. Results are shown for different:

- **Durations:** the length of time for which the set of procured wind generators are expected to deliver at least as much the specified power output, 500 MW in this example, without dipping below this level at any time. Naturally, if the specified power is only needed for one hour, then the capacity required to reliably deliver it is lower than if it was required for 24 hours.
- **Percentiles:** the 50th percentile value represents the amount of procured capacity that has a 50 per cent chance of success in generating at least 500 MW consistently throughout the required duration, given an unknown start time (and day of year). The 99.9th percentile value represents the amount of procured capacity necessary for a 99.9 per cent chance of success in delivering at least 500 MW consistently for the required duration, given an unknown start time.

Table 0.1

The average amount of procured wind capacity (in MW) required to secure a sustained 500 MW of power, for various durations and risk levels

Duration	50th percentile	90th percentile	99th percentile	99.9th percentile
1 hour	1,587	5,742	14,368	27,873
4 hours	1,771	6,679	16,750	32,599
8 hours	2,019	7,816	18,731	36,635
12 hours	2,267	8,781	20,410	40,343
24 hours	3,023	10,798	23,856	52,446

Figure 0.1

The average amount of procured wind capacity required to secure a sustained 500MW of power, for various durations and risk levels

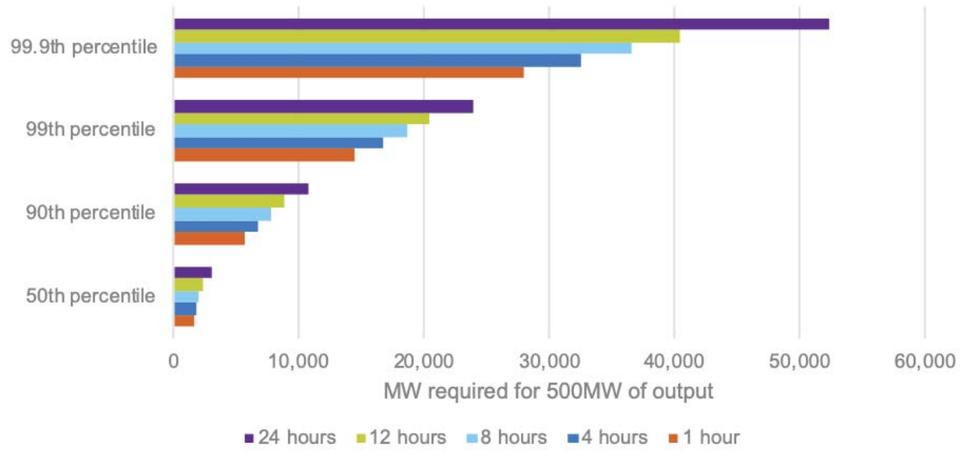

Executive summary

2 There can be very large differences in the ability of two sets of wind farms to provide reliable output, despite their total capacity being very similar, due to differences in their geographical arrangement. This would certainly need to be considered carefully by the Energy System Operator before procuring Black Start services from wind generators. Despite this variability, the overall trend in ability of wind farm sets to provide reliable constant outputs as a function of total procured capacity is almost exactly linear.

3 The level of power that can be sustained by a specific set of wind farms for a specified duration, and with a specified risk of failure, is dependent on a number of factors. In order of decreasing impact, these factors are: (i) the initial wind speeds (i.e. resource availability state); (ii) the month of the year; (iii) long-term random anomalies in resource availability patterns; (iv) the time-of-day.

If taking average values across all of these factors, then the ability of wind to contribute to restoration is not very significant. For example, as table 0.1 above shows, if attempting to provide 500 MW of wind with only a 1% risk of failure, then between 30 to 40 times more wind is required to be procured compared to the consistent output required. However, for specific combinations of these factors, the availability of wind to contribute can actually be very high. In particular, if the initial wind speeds are high, the analysis shows that wind can be relied upon with a reasonably low risk of failure to provide quite high levels of outputs for long durations. NGESO could exploit this when planning for system restoration in operational timescales – for example, if NGESO has a reliable forecast that wind speeds will be high for one or two days, then this may allow them to defer dispatching out-of-merit conventional Black Start providers without compromising their Black Start readiness.

4 A complex, multi-step statistical time series model can be built of the power output from a specified set of wind farms. The exact same model structure (with different parameter values) can be used to represent wind farm sets ranging from single wind farms to the full set of 22 farms studied. This model has been demonstrated to produce synthetic series that show a very good degree of agreement with historical series, when subjected to many diverse statistical comparisons.

5 The model can be used dynamically on many planning and semi-operational time scales to update predictions of the deliverable power level at specified durations and risk levels. However, if there is a need to update predictions with a nearly real-time frequency, the model should be expanded to include numerical, meteorological forecasts, ideally in ensemble form.

National Grid ESO | July 2019 **Black Start from Non-Traditional Generation Technologies**

1 Introduction

The capability of non-traditional technologies, particularly variable renewable generation, to contribute to system restoration is significantly limited by the uncontrollable variability of these generation sources. Within this project, we have explored a case study of wind contributing to system restoration.

When growing a power island during system restoration with traditional providers, the ESO would be able to add block loads to the island based on reliable ramping rates and levels of output from these providers. There are some uncertainties over longer term planning timescales, for example, associated with communications failures for these providers or unforeseen maintenance. However, in operational planning and scheduling timescales, once the ESO knows a generator has been made ready, it will be able to rely on that provider's output to contribute towards restoration with a high level of certainty.

This is not typically the case with prospective service provision from variable generation like wind and solar. This generation could, in principle, be used by the ESO in two different ways:

- It could be used to supplant the provision of the service from traditional providers. In this case, the ESO needs to know what level of output can be relied upon from the wind and solar (at a pre-determined level of certainty) and, therefore, the reduction in procurement from other traditional providers than can be achieved.
- It could be used to accelerate restoration, without reducing service procurement from traditional providers, by simply allowing this additional wind generation to spill onto the system and adding further block loads to match. However, in this case, ESO will need to be careful not to add block load too quickly, as if wind conditions subside before other generation becomes available this could lead to a shortage.

In both cases, the ESO will need to have an idea (i.e. forecast) of the extent to which it can rely upon wind to provide a minimum consistent level of output. The aims of this analysis, therefore, are:

- To determine the output of wind that can e relied upon during system restoration, i.e. the maximum sustainable output from a group of windfarms over some period of time. If time is modelled as progressing in discrete intervals/ steps (e.g. hours or 30 minutes), then the quantity of interest is the minimum, across all time steps in the restoration period, of the maximum power that can be delivered during each time-step.
- To explore how this quantity varies depending on how much risk the ESO is willing to accept (e.g. 50 per cent chance of success vs 99 per cent chance of success).
- To understand how this varies based on the required duration of the sustained output.
- To understand how this varies predictably between seasons and time of day, and also more randomly between generally good and poor wind resource seasons.
- To see how this varies depending on the combination of wind farms used.
- To demonstrate a model which could be used by the ESO when planning system restoration using variable resources.

1 Introduction

To address these questions, the first part of the analysis presented here involved detailed statistical analysis of wind farm output time series, generated from historical meteorological data. Following this, we present the development of a generative statistical model capable of producing unlimited synthetic wind resource data, and demonstrate some of this model's value. The model was trained on the quasi-historical series analysed in the first part on this case study. As it stands, a model like this could be used by the ESO on a planning time-scale to help understand the value of wind for restoration. Further, with the inclusion of probabilistic meteorological forecasts – such as those currently used by the ESO – the model could be adapted to become a state-of-the-art forecasting tool that predicts, on operational time-scales, the level of wind power output that can be maintained for a specified duration at a specified level of certainly.

We have used Scotland as a case study for this, as a single Black Start zone. This means the results are uniquely determined by both (i) the nature of the wind farms in Scotland (e.g. capacities and turbine models) and (ii) the wind resource in Scotland. We have considered both large onshore and offshore wind farms in Scotland (although in principle the same approach could be applied to smaller distribution connected sites).

The rest of this report is structured as follows:

1. a description of the data sources we have used
2. summarising the insight gained from exploratory data analysis
3. a description of the methodology used for the generative model
4. some example results using the generative model
5. a discussion of further applications of this model, and useful extensions that could be made in the future, including use as a forecasting tool.

2 Data sources

To complete this analysis, we needed granular time series (e.g. hourly) of the power output from a diverse mix of wind farms in the Black Start zone we are studying (Scotland), so that we can model evolution in time in such relatively short time-steps. In addition, these had to be long time series, to ensure that they fully account for both the short and long-term dynamics of the wind resource, given that the wind climate can change quite significantly year-on-year in addition to the short-term variability in weather. As a result, in many cases this might require "historical" time series for wind farms that actually pre-date the commissioning of the wind farms.

To get around this problem, we have generated "synthetic" time series, using the Renewables Ninja[1] website, developed by Iain Staffell of Imperial College London and Stefan Pfenninger of ETH Zurich. This site provides the user with hourly-resolution time series of the (maximum) power outputs of virtual wind farms and solar power plant at any location in the world, while also accounting for the way that turbines and solar panels convert the renewable resource into power (e.g. the effect of wind turbine power curves).

These time series are generated using the user's choice from two reanalysis meteorological datasets: NASA's MERRA dataset and the SARAH dataset from CM-SAF, with the former chosen for this study, due to its superior accuracy for European locations. Reanalysis datasets use a combination of highly complex meteorological models with many sources of observational data and satellite imagery to produce a self-consistent and plausible 'best guess' of the physical state of the atmosphere at all hours over the last few decades, and from this the true value of many meteorological variables. The model allows for the generation of time-stamped series of wind farm (and solar PV) output for up to 17 historical years, namely 2001–2017. That is, a wind power output time-stamped as e.g. 22:00 on 26 April 2010 is the maximum power a wind farm at that location could have produced at this time, given the meteorological conditions according to the chosen reanalysis model. This overall approach is similar in many ways to the way wind power is represented by the ESO in capacity adequacy studies.

The primary variable is obviously wind speed, but others are also used to dynamically calculate air density. The smoothing effect of high frequency fluctuations in wind speed about the hourly average values is accounted for, along with the fact that these fluctuations are only partially correlated across the spatial extent of the wind farm – although the model assumes fixed spatial dimensions. The tool also allows the user to specify the turbine hub height, and to choose the turbine model from an extensive, but not exhaustive, set of options.

As inputs to this model, we have compiled data for 22 large onshore and offshore wind farms in Scotland, including location, capacity, turbine type, and turbine hub height – drawing from a wide variety of online data sources, such as government databases. The set is comprised of the 20 largest operational onshore wind farms in Scotland, and the two large offshore projects that are currently operational, and their details are listed in table 2.1 below.

Table 2.1

Key attributes of the Scottish windfarms modelled in the case study

Wind farm	Latitude	Longitude	Capacity (MW)	Hub height	Turbine type
Arecleoch	55.0688	-4.7971	120	95	Gamesa G80 2000
Bhlaraidh	57.2191	-4.5819	110	76	Vestas V112 3300
Black_Law	55.7622	-3.7626	124	82	Siemens SWT 2.3 101
Clyde_1	55.4416	-3.5427	350	80	Siemens SWT 2.3 82
Clyde_2	55.5041	-3.5479	172.8	80	Siemens SWT 3.0 101
Corriegarth	57.1901	-4.3596	69	78	Enercon E82 3000
Crystal_Rig_2	55.8995	-2.5293	138	75	Siemens SWT 2.3 101
Dersalloch	55.3042	-4.4882	69	50	Siemens SWT 3.0 101
Dunmaglass	57.2515	-4.2558	94	70	GE 2.75 103
EOWDC	57.2167	1.9833	93.2	120	Vestas V112 3300

[1] https://www.renewables.ninja/

2 Data sources

Wind farm	Latitude	Longitude	Capacity (MW)	Hub height	Turbine type
Fallago_Rig	55.8313	-2.6880	144	80	Vestas V90 3000
Farr	57.3339	-4.1030	92	60	Bonus B82 2300
Gordonbush	58.0595	-3.9591	70	69	Vestas V80 2000
Griffin	56.5812	-3.7338	156	80	Siemens SWT 2.3 101
Hadyard_Hill	55.2568	-4.7319	119.6	70	Siemens SWT 2.3 101
Harestanes	55.2390	-3.5742	136	80	Gamesa G90 2000
Kilgallioch	54.9884	-4.8017	240	88	Gamesa G90 2000
Robin_Rigg	54.7642	-3.6955	174	80	Vestas V90 3000
Stronelairg	57.1269	-4.5264	228	125	Vestas V112 3300
Whitelee_1	55.6812	-4.2791	322	65	Siemens SWT 2.3 82
Whitelee_2	55.6771	-4.2868	108	65	Alstom Eco 110
Whitelee_3	55.6394	-4.3176	109	65	Alstom Eco 110

This data is drawn from quite a diverse range of publicly available sources. Wherever possible, we have used the Government's Renewable Energy Planning Database[2]. However, this has some gaps, which we have filled using information from sources like Wikipedia as well as the information websites of individual wind farm owners.

The wind farms are shown geographically in figure 2.1. This illustrates that there is a reasonably diverse geographic coverage of the entire Black Start zone.

[2] https://www.gov.uk/government/publications/renewable-energy-planning-database-monthly-extract

National Grid ESO | July 2019 Black Start from Non-Traditional Generation Technologies

Figure 2.1
Geographic location of wind farms

By running this data through the Renewables Ninja model, we have generated 17-year long time series for each of these 22 wind farms. We have also considered 500 different combinations of these wind farms, aggregating the time series together as required. The different combinations are plotted in figure 2.2, which shows the levels of maximum installed capacity that are possible for different combinations of wind farms. Each dot represents a different unique combination of wind farms.

2 Data sources

Figure 2.2
Total capacities of each modelled windfarm combination

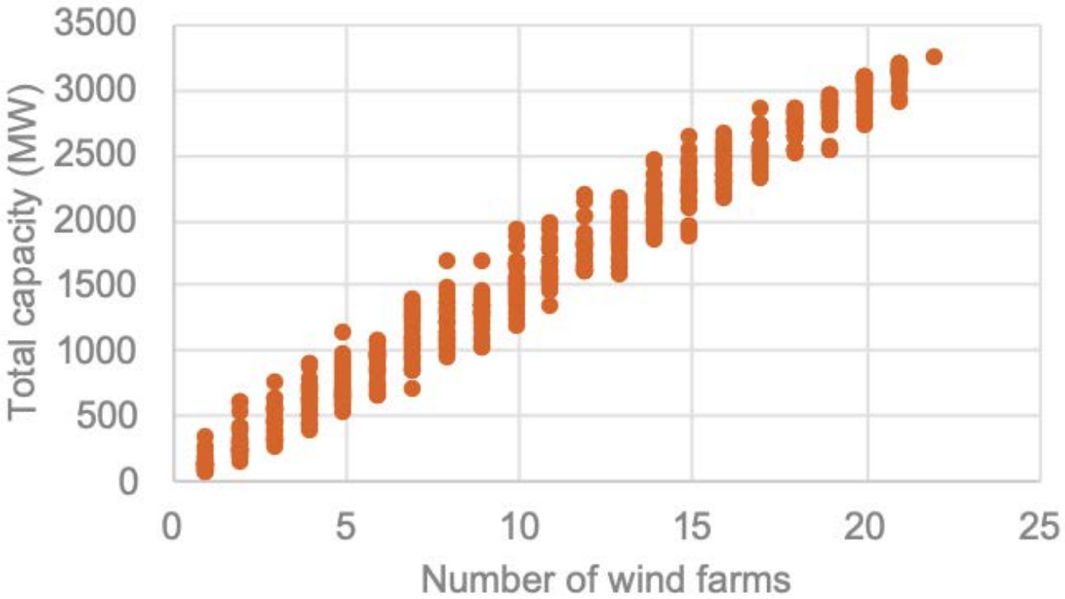

By studying multiple wind farm combinations, we can explore how their reliable level of combined output varies depending on which combination is selected. For example, there are 38 combinations of wind farms that would give a maximum installed capacity of between 900 MW and 1,100 MW.

3 Exploratory data analysis

3.1 Introduction

The first five aims introduced at the start of this section are initially explored through analysis of the time series data. This is achieved by:

1. selecting one of the 500 combinations of wind farms
2. selecting a "start date" at random from the 17-year ries for this wind farm combination, or more specifically a "start datetime"
3. identifying the maximum sustainable output produced by this wind farm combination in the d hours after this start date
4. repeating this process 1,000s of times, selecting a different start datetime for each repetition.

As stated above, the maximum sustainable output is the minimum, across d hours, of the maximum (average) power that the wind farm could have generated during each hour, as a result of the meteorological conditions that existed at that time. The implication is that the windfarms would produce as much as they can during the hour with the least favourable meteorological conditions, but would curtail their output to this consistent level during the hours with more favourable conditions.

The outlined method means that, for any wind farm combination, there is a large set of values for the maximum output that could be sustained consistently for d hours by that combination. This collection of values reflects the variation in wind patterns within the 17-year series, and may be viewed as an "unconditional" probability distribution for the maximum sustainable power immediately following a blackout event at an arbitrary time. In the remainder of this discussion, we focus on some pre-selected percentiles of this probability distribution:

- The 50th percentile. The precise meaning here is: "For a blackout at an arbitrary 'datetime', and in the absence of any information about the meteorological conditions immediately prior to the event, and without a meteorological-model forecast for the hours following the event, there is a 50 per cent probability that the maximum sustainable power for d hours is as least as big as this 50th percentile value." In other words, if the wind collection of farms promised to consistently deliver this level of power, there is a 50 per cent chance that it would fail to deliver on that promise during at least one hour. It should be noted that the data available means we can only consider the probability of the hourly averages being lower than promised, rather than failing to deliver for e.g. 1 minute only.
- The 90th percentile – analogous to the description above, but a lower value, so that there is only a 10 per cent chance that the collection of wind farms would fail to deliver this power during at least one hour of the period d.
- The 99th percentile – analogous to the description above, but an even lower power value, so that there is only a 10 per cent chance that the collection of wind farms would fail to deliver this power during at least one hour of the period d.
- The 99.9th percentile – analogous again to the first description, but so low such that there is only a 0.1 per cent chance that the collection of wind farms would fail to deliver this power during at least one hour of the period d.

In addition, we have explored several different required durations, d[3]:

- one hour from the start datetime
- four hours from the start datetime
- eight hours from the start datetime
- 12 hours from the start datetime
- 24 hours from the start datetime.

An indicative selection of results for the different durations and levels of risk are presented in the following subsections.

3.2 Distributions over all possible initial conditions

The procedure for calculating distributions for maximum sustainable power levels presented above was executed, initially covering all possible start datetimes within the 17-year series.

The results are presented as scatter plots in figures 3.1 to 3.5 below. In each plot, there is a coloured point for each unique combination of wind farms and the risk-level (i.e. quantile of the distribution across start datetimes) is constant. The x-axis value for each point is the total capacity of the windfarm combination, while the y-axis gives the quantile of the maximum sustainable power for that combination. There are five colours, representing the five investigated values for the required duration.

For ease of comparison attempts were made to display results on the same axis ranges, to the greatest extent possible. This occurs by default with regard to the x-axis, and for the y-axis we have used 3 ranges, that are repeated throughout the following subsections. These ranges are:

A 0–1,200 MW;
B 0–200 MW; and
C 0–2,500 MW.

[3] Within the generative model, we explore even longer durations such as 120 hours.

3 Exploratory data analysis

Perhaps the most surprising observation that can be made from the set of plots is that the maximum sustainable power shows a highly linear relationship to the procured capacity. For this reason, straight lines were fitted (using ordinary least squares minimisation) to describe the maximum sustainable power levels as linear functions of the procured capacity, but these are not shown in the figures to avoid an overload of information. Instead, the gradient and intercept values for the various durations and percentiles are presented in table 3.1. As expected, the table shows that gradients decrease as durations and percentiles increase, and given the scale of power values, the intercepts are negligible.

Figure 3.1 presents results for the 50th percentile, using the range A, and it can be seen that points are positioned fairly close together in a straight line, with the different durations showing a moderate degree of separation. The gradient value of 0.31 for a duration of one hour is certainly consistent with load factors for within GB, but this falls to 0.01 for the longest durations and highest percentiles. Similar observations can be made for 90th percentile values, presented in figure 3.2, although naturally the gradients are all reduced, by factors in the range of three–four. The roughly linear relationship is also true for 99th quantile, shown initially in figure 3.3 using y-axis range A, and again in figure 3.4 using the y-axis range B, for greater clarity. The gradients for the 99th percentile are roughly 10 per cent of their values for the 50th percentile, but the degree of separation among the colours is noticeably reduced. This greater variability between different combinations of windfarms with similar capacities, relative to the impact of duration, is most prominent for 99.9th percentile values, presented in figure 3.5.

Figure 3.1

50th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all possible conditions. Y-axis range A.

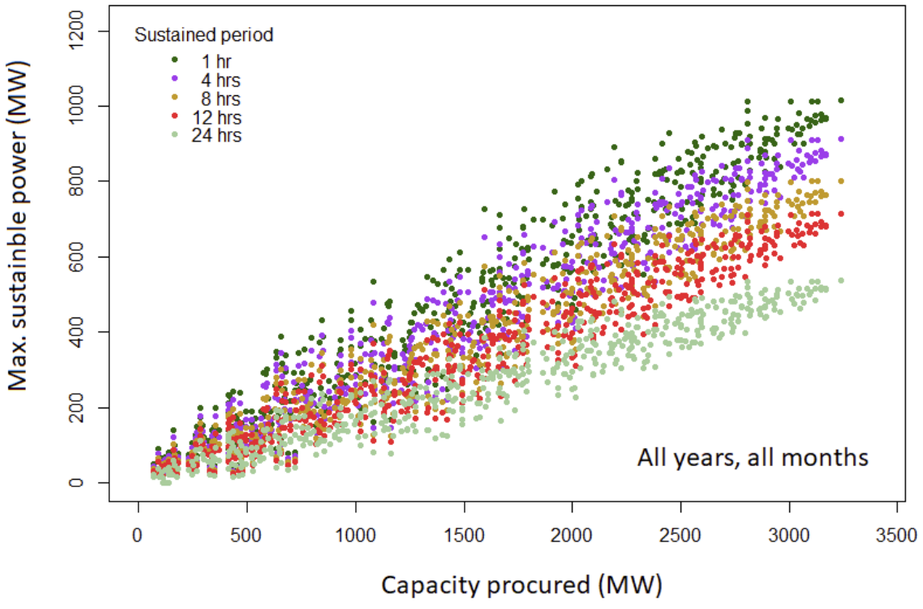

Figure 3.2

90th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all possible conditions. Y-axis range A.

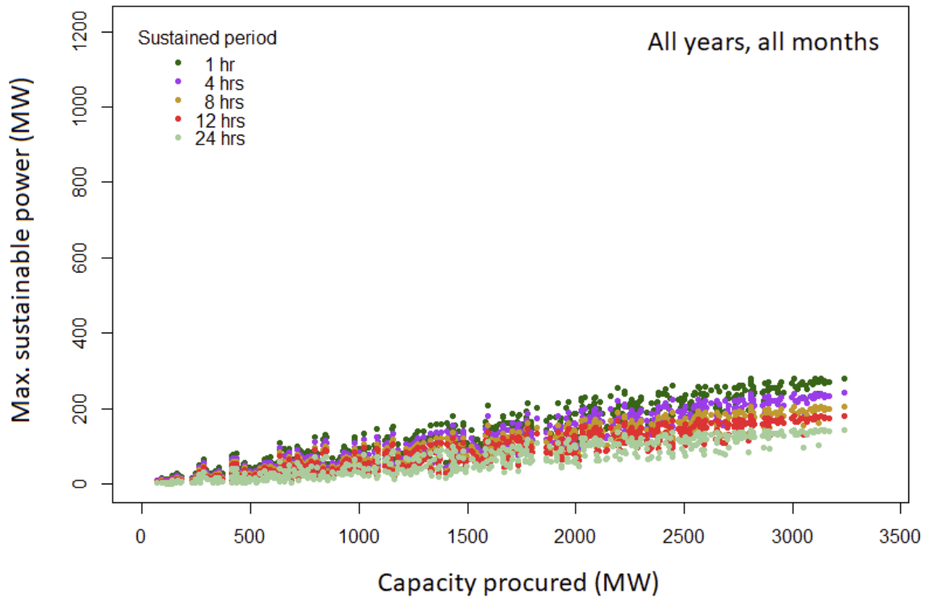

Figure 3.3

99th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all possible conditions. Y-axis range A.

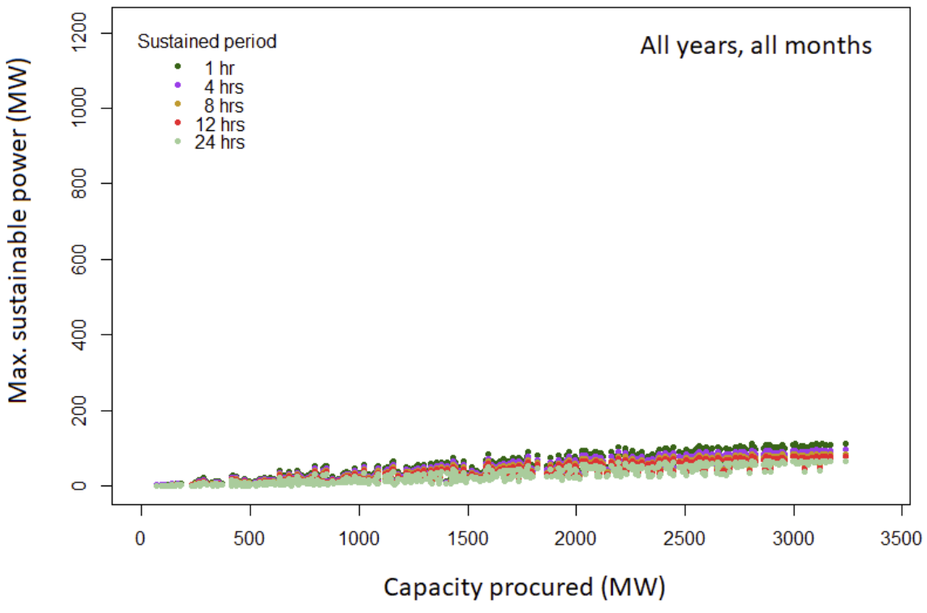

3 Exploratory data analysis

Figure 3.4
99th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all possible conditions. Y-axis range B.

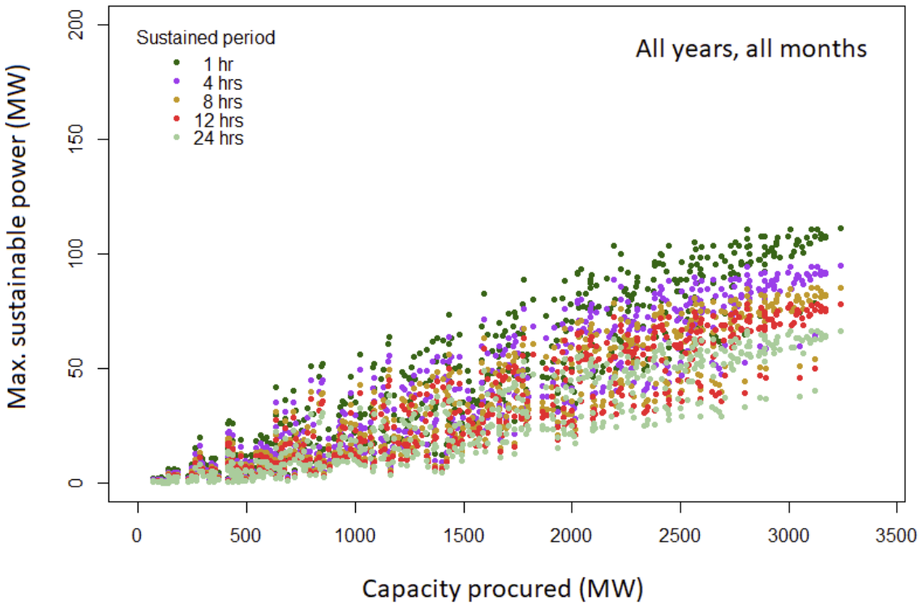

Figure 3.5
99.9th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all possible conditions. Y-axis range B.

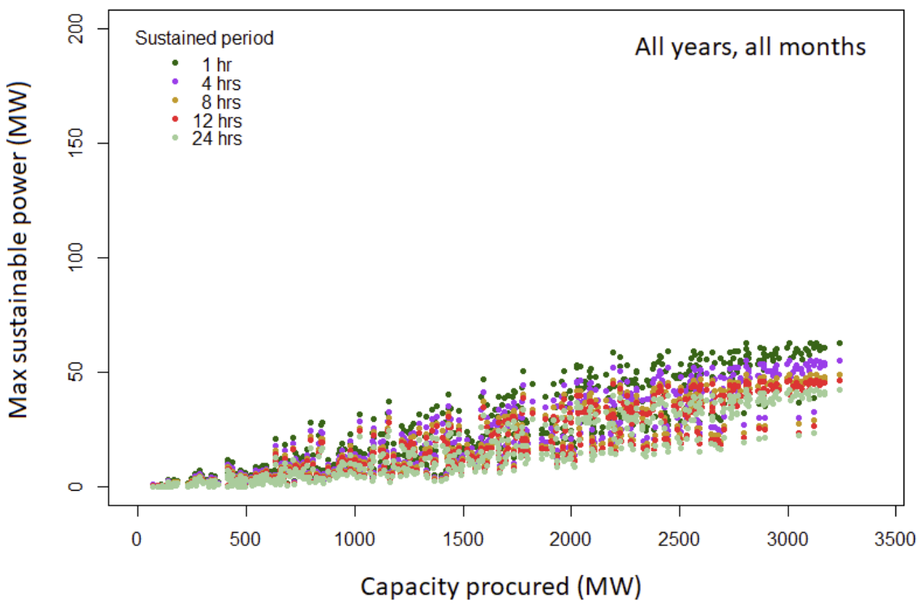

Table 3.1
Gradient and intercepts for best-fit straight lines for maximum sustainable power vs. procured capacity, all years, all months

Duration	50th percentile	90th percentile	99th percentile	99.9th percentile
1 hour	0.31, 9.35	0.09, -1.78	0.03, -5.68	0.02, -4.88
4 hours	0.28, 6.82	0.07, -3.01	0.03, -5.70	0.02, -4.62
8 hours	0.24, 4.87	0.06, -3.92	0.03, -6.02	0.01, -4.64
12 hours	0.22, 3.25	0.06, -4.58	0.02, -5.91	0.01, -4.57
24 hours	0.16, 0.21	0.05, -5.85	0.02, -5,54	0.01, -4.73

3.3 Seasonal results

Very similar scatter plots and table are presented in this subsection, but here we present results exploring the effect of time-of-year on the ability of wind power to support system restoration. We did this by producing scatter plots where start times were restricted to be either in July (figures 3.6 and 3.7) or in January (figures 3.8 and 3.9) – found to be on average the worst and best resource months, respectively. For brevity, we present results for only the 50th and 99th percentiles here.

It is clear from the figures and table 3.2 that:

i The same basic patterns observed in the year-round plots are present here, albeit scaled-up in the case of January and scaled down in the case of July;

ii The scaling factors involved are very significant. Indeed, for the 50th, 90th and 99th percentiles, the January gradients are roughly twice as large as the July gradients, but the differences are mostly insignificant (when rounded to two decima places) for the 99.9th percentile. Gradients for January are roughly 40 per cent greater than the all-month unconditional values.

Figure 3.6
50th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all starting datetimes in July. Y-axis range A.

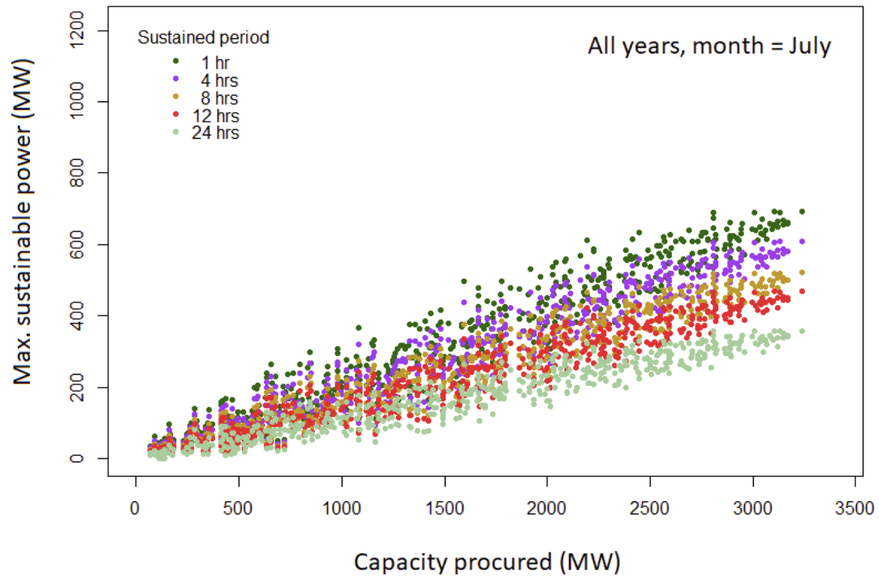

3 Exploratory data analysis

Figure 3.7
50th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all starting datetimes in January. Y-axis range C.

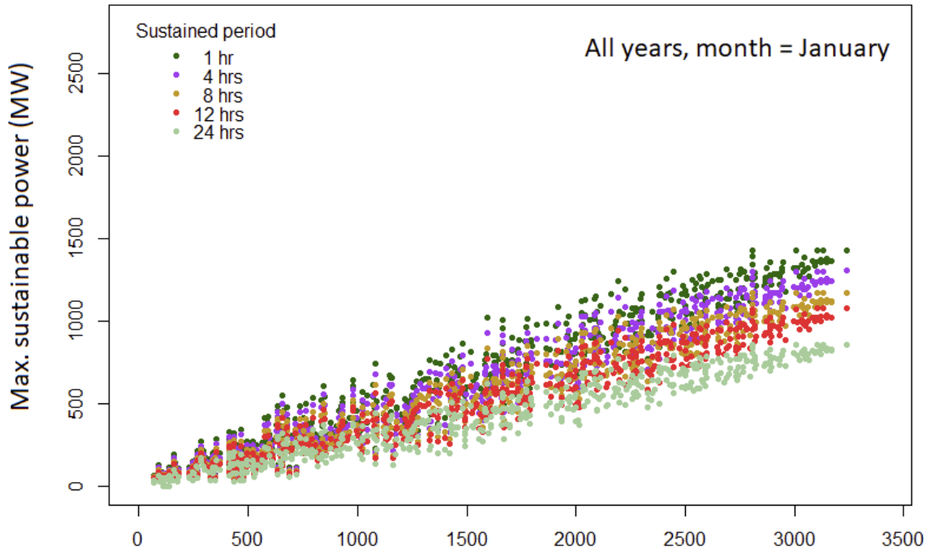

Figure 3.8
99th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all starting datetimes in July. Y-axis range B.

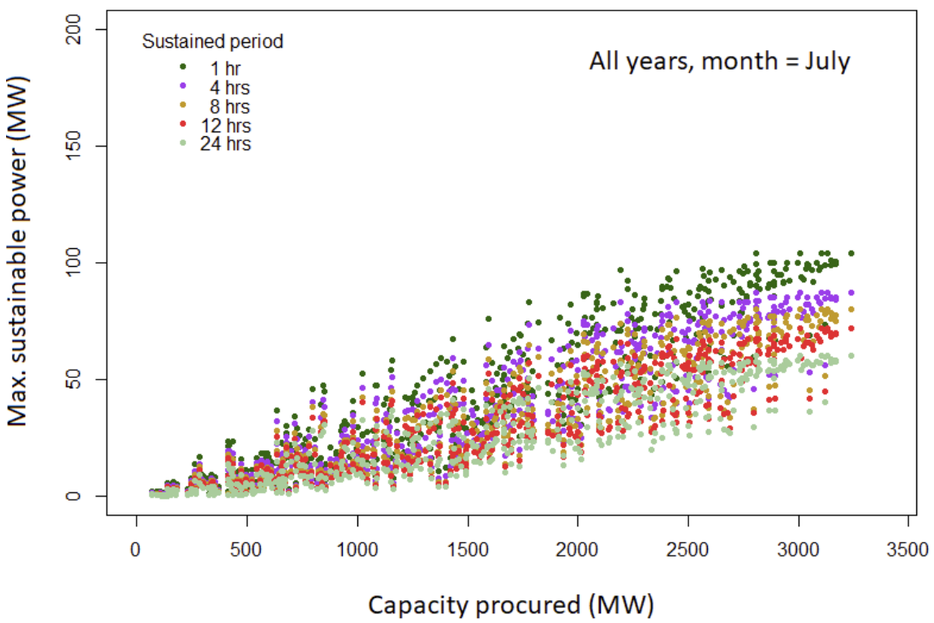

Figure 3.9
99th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all starting datetimes in January. Y-axis range B.

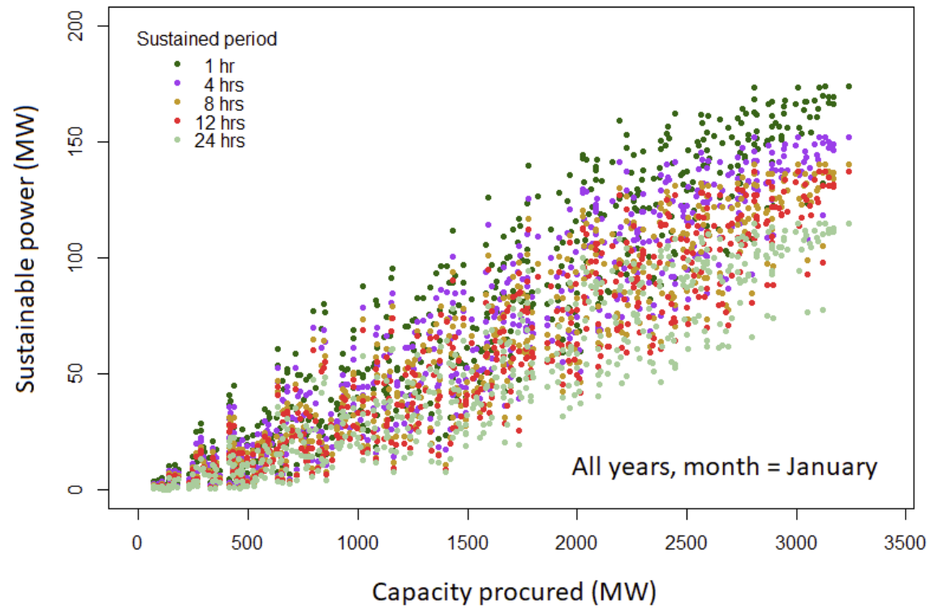

Table 3.2
Gradient and intercepts for best-fit straight lines for maximum sustainable power vs. procured capacity, all years, July and January

Month	Duration	50th percentile	90th percentile	99th percentile	99.9th percentile
July	1 hour	0.21, 6.17	0.07, -2.44	0.03, -6.02	0.02, -4.94
	4 hours	0.18, 4.36	0.06, -3.78	0.03, -5.52	0.02, -5.72
	8 hours	0.16, 2.68	0.05, -4.35	0.03, -6.05	0.02, -6.63
	12 hours	0.14, 0.56	0.05, -4.95	0.02, -5.26	0.02, -7.03
	24 hours	0.11, -2.10	0.04, -6.89	0.02, -4.59	0.02, -7.35
January	1 hour	0.43, 13.55	0.13, 0.17	0.05, -7.14	0.04, -9.48
	4 hours	0.39, 11.48	0.11, -1.50	0.05, -7.66	0.03, -8.00
	8 hours	0.36, 8.78	0.10, -1.56	0.04, -8.66	0.03, -7.01
	12 hours	0.33, 6.28	0.08, -2.35	0.04, -10.14	0.02, -6.58
	24 hours	0.26, 1.56	0.07, -5.62	0.04, -9.51	0.02, -6.07

3 Exploratory data analysis

3.4 Long-term random variability

This subsection similarly presents results of investigations where start datetimes were restricted. In this case, all months were allowed but start datetimes were restricted to specific years – 2015 and 2010, found to be the best and worst wind resource years, respectively. This was established by calculating annual mean wind power output levels for the aggregated output of every modelled windfarm. This aspect of variability was investigated since it is well-established that the wind resource tends to display long-term random variability, manifested as winters that are quite consistently windier, or less windy, than average – with the relatively low yields of 2010, for example, reportedly placing financial stress on wind developers. This can be expressed in formal statistical terms as stating that wind resource as a random variable possesses "long memory", which meteorologists are able to explain in terms of phenomena such as the North Atlantic Oscillation.

Figures 3.10–3.13 below demonstrate the impact of this variability, again showing only the 50th and 99th percentiles but presenting gradients for all four percentiles in table 3-3. It can be seen that the effect is certainly significant, but considerably less pronounced than the impact of time-of-year, and influencing 50th percentile values much more than the higher percentiles.

Figure 3.10

50th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all starting datetimes in 2015. Y-axis range A.

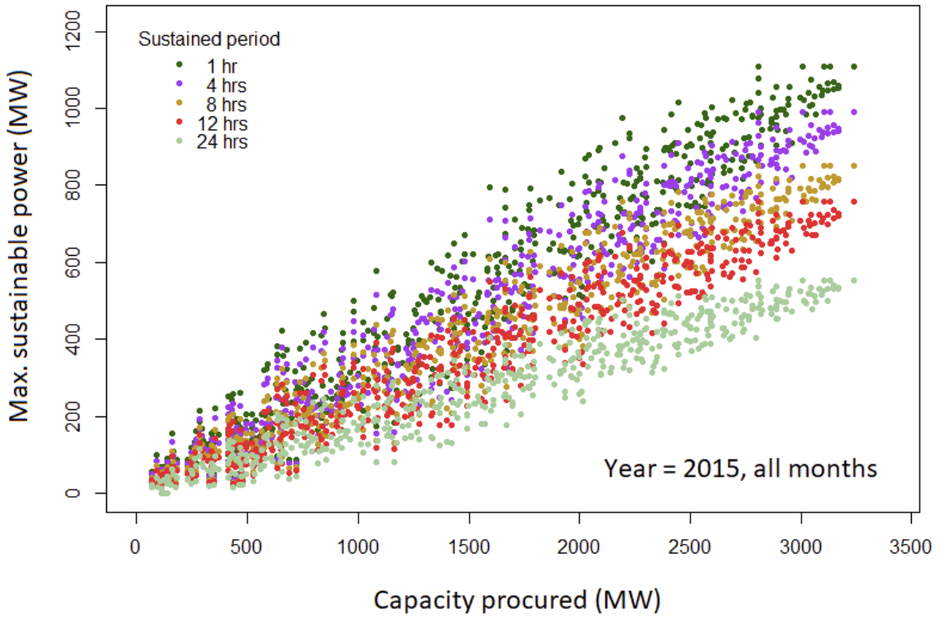

Figure 3.11

50th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all starting datetimes in 2010. Y-axis range A.

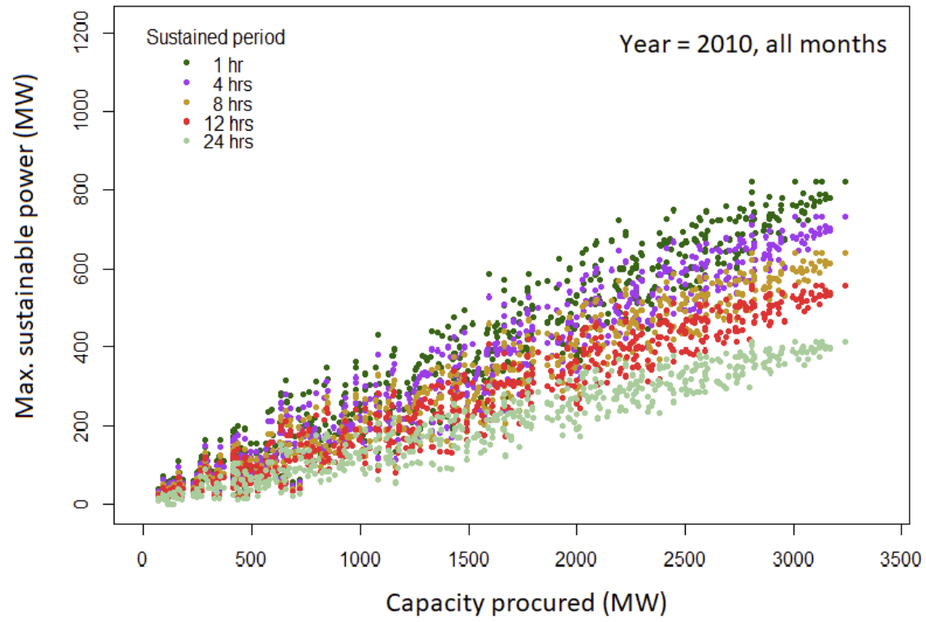

Figure 3.12

99th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all starting datetimes in 2015. Y-axis range B.

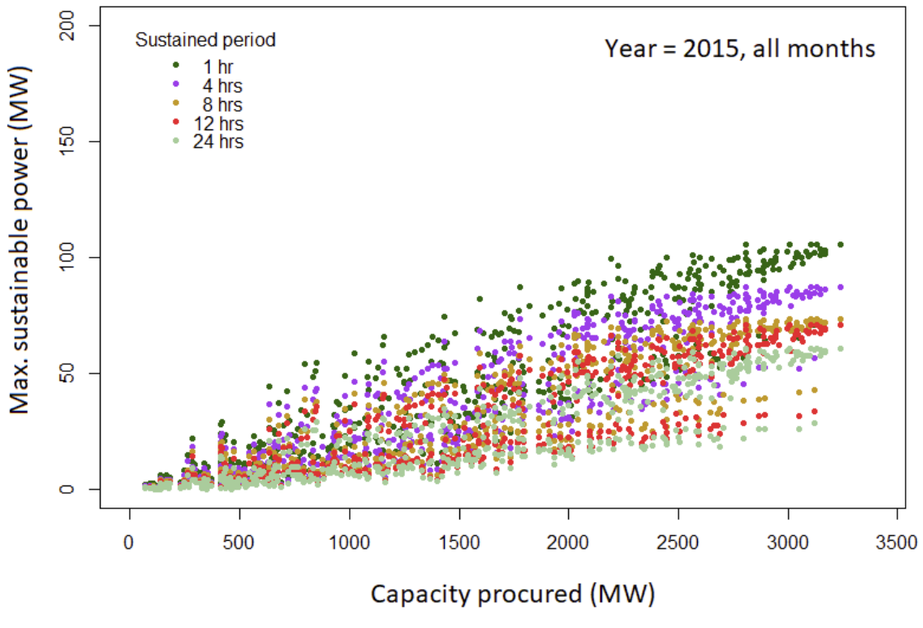

3 Exploratory data analysis

Figure 3.13
99th percentiles of the distributions of maximum sustainable power for each windfarm combination, over all starting datetimes in 2010. Y-axis range B.

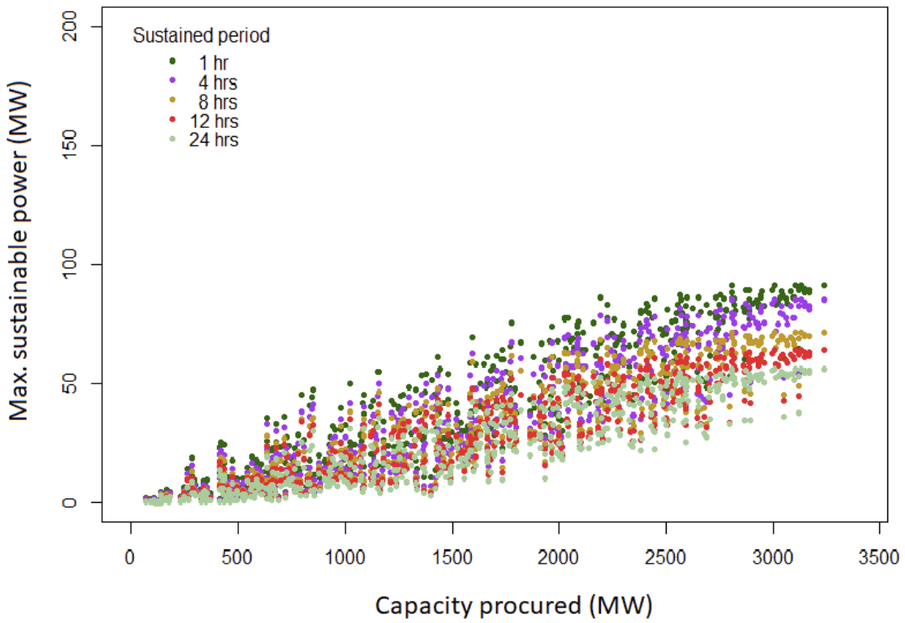

Table 3.3
Gradient and intercepts for best-fit straight lines for maximum sustainable power vs. procured capacity, all months, 2010 and 2015

Year	Duration	50th percentile	90th percentile	99th percentile	99.9th percentile
2010	1 hour	0.25, 8.49	0.07, -2.27	0.03 -3.94	0.02, -4.86
	4 hours	0.22, 5.89	0.06, -3.48	0.03, -5.32	0.02, -5.56
	8 hours	0.19, 3.51	0.05, -4.35	0.02, -4.54	0.02, -5.88
	12 hours	0.17, 2.99	0.05, -5.72	0.02, -4.08	0.01, -5.08
	24 hours	0.13, -0.99	0.04, -4.95	0.02, -4.28	0.01, -5.08
2015	1 hour	0.33, 10.35	0.09, -2.50	0.03, -4.48	0.02, -2.60
	4 hours	0.30, 8.21	0.07, -3.29	0.03, -5.70	0.02, -2.58
	8 hours	0.26, 6.92	0.06, -2.74	0.02, -6.19	0.02, -2.57
	12 hours	0.23, 4.28	0.05, -4.36	0.02, -5.83	0.01, -2.57
	24 hours	0.17, 1.88	0.04, -4.27	0.01, -2.98	0.01, -2.57

3.5 Initial wind resource state

This subsection presents results of the final aspect of variability investigated based on restricted start datetimes – the strength of the resource during the first hour of the sampled periods. The strength of the resource is defined here as the aggregated power output from every modelled windfarm. For contrast, we produced separate plots for start times where the strength of the resource was in the bottom 10 per cent of its values (across the 17-year series of such values), in figures 3.14 and 3.16, and for the top 10 per cent of these values, in figures 3.15 and 3.17. One difference here, compared to the last two subsections, is that the higher percentile value presented here is the 99.9th. This choice is due to the fact that, for starting times with high wind speed, the results are extremely unsensitive to the choice of distribution percentile. Gradients for the full set of duration-percentile pairs are again presented, in table 3.4.

The plots reveal, perhaps unsurprisingly, that the initial 'windiness' (i.e. comparing the top and bottom 10 per cent) has by far the strongest impact on results, at least an order of magnitude greater than the previous impacts. Indeed, the central gradient of the plots for the 50th percentiles differ by a factor of about 1,000 per cent–1,500 per cent and, rather than decrease for the higher percentiles, in this case the difference grows to about 7,000 per cent for the 99.9th percentile, and long durations – much bigger than the typical difference in initial conditions.

It is also worth noting that the colours are completely mixed in figures 3.15 and 3.17, indicating that the spatial distribution of the windfarms is more significant than the required duration at such high percentiles (i.e. extremes of low risk).

Figure 3.14
50th percentiles of the distributions of maximum sustainable power for each windfarm combination, over the bottom 10 per cent 'windiest' starting datetimes. Y-axis range A.

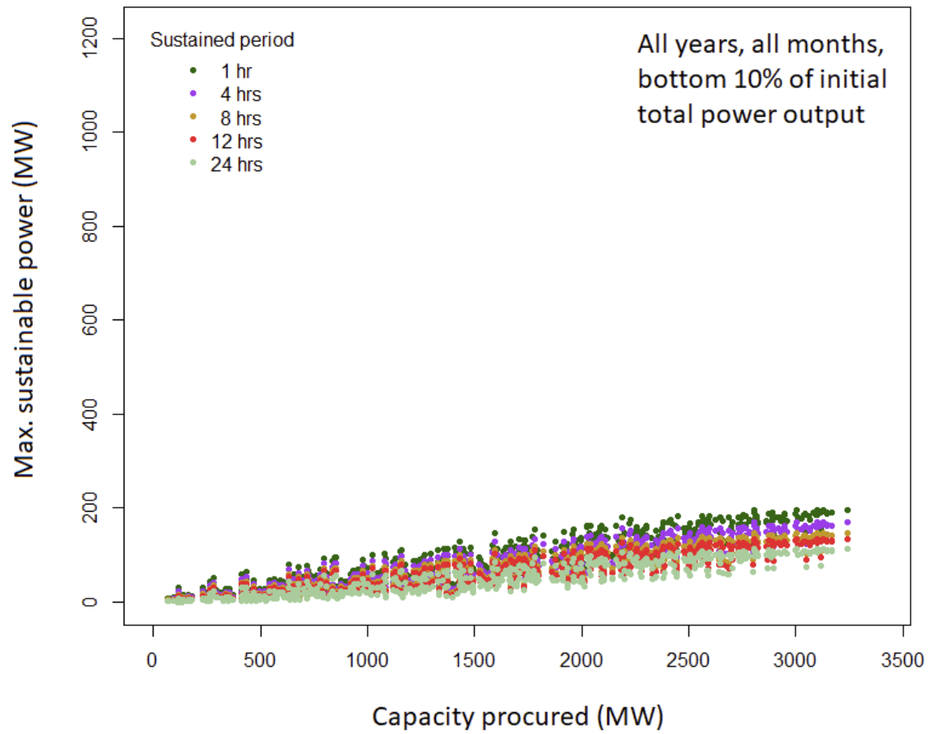

3 Exploratory data analysis

Figure 3.15
50th percentiles of the distributions of maximum sustainable power for each windfarm combination, over the top 10 per cent 'windiest' starting datetimes. Y-axis range C.

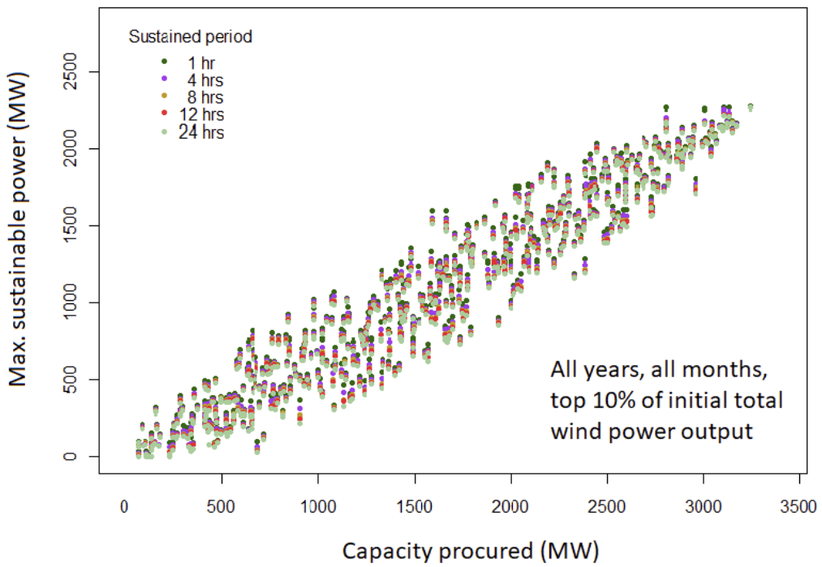

Figure 3.16
99.9th percentiles of the distributions of maximum sustainable power for each windfarm combination, over the bottom 10 per cent 'windiest' starting datetimes. Y-axis range B.

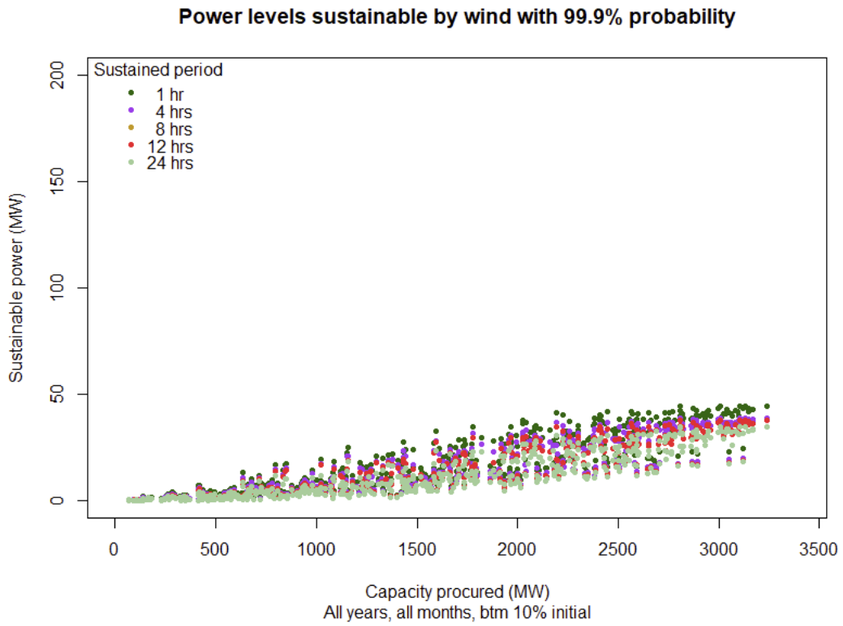

Figure 3.17
99.9th percentiles of the distributions of maximum sustainable power for each windfarm combination, over the top 10 per cent 'windiest' starting datetimes. Y-axis range C.

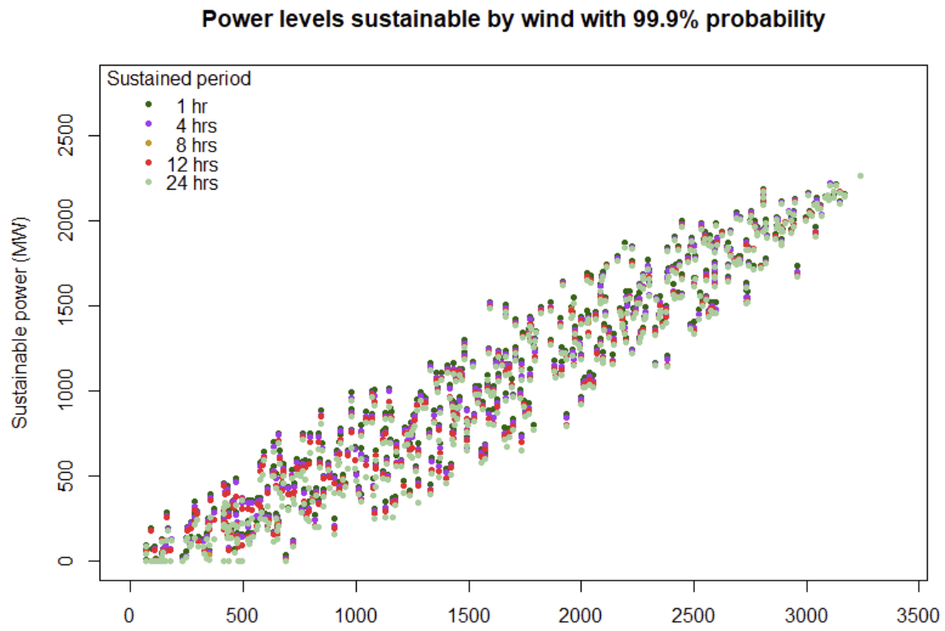

Table 3.4
Gradient and intercepts for best-fit straight lines for maximum sustainable power vs. procured capacity, top and bottom 10 per cent of windiest initial conditions

Initial Windiness Range	Duration	50th percentile	90th percentile	99th percentile	99.9th percentile
Bottom 10%	1 hour	0.06, -1.29	0.03, -5.52	0.02, -4.86	0.02, -2.51
	4 hours	0.05, -3.15	0.03, -5.59	0.02, -4.60	0.02, -2.50
	8 hours	0.05, -4.24	0.03, -5.94	0.01, -4.59	0.01, -2.50
	12 hours	0.04, -4.95	0.02, -5.84	0.01, -4.50	0.01, -2.50
	24 hours	0.4, -5.84	0.02, -5.47	0.01, -2.56	0.01, -2.35
Top 10%	1 hour	0.77, 23.06	0.70, 12.14	0.70, -21.88	0.70, -60.27
	4 hours	0.74, 21.70	0.69, 1.49	0.70, -40.65	0.70, -81.54
	8 hours	0.71, 21.00	0.69, -7.85	0.70, -51.71	0.70, -99.9
	12 hours	0.70, 17.21	0.69, -13.65	0.70, -58.31	0.70, -100.23
	24 hours	0.69, 7.64	0.69, -24.55	0.70, -69.31	0.72, -148.31

4 Statistical time series modelling

4.1 Motivation

This section describes the various components of a multifaceted statistical model fitted to the wind power time series. Such statistical models can be used for short term forecasting, based on recent meteorological conditions, but the objective in this case was to develop a model capable of generating unlimited new time series data, where the simulated series displays the same temporal patterns, correlations and seasonal dependencies on many time scales as the historic series.

One reason to do this is that the model fitting process provides considerable insight about the nature of the random process being modelled[4]. For example, in our case, it has provided an illustration of the different timescales over which wind output varies.

The other reason to develop a statistical model is to avoid being limited by the finite sample size of the historical data. Even though the series are very long, at about 150,000 hours, if one is interested in e.g. a specific month and time of day, the data becomes significantly smaller. Also, the historical series have a sample size of only 17 with regard to differences between e.g. annual means, or winter seasonal means etc.

4.2 Model fitting

4.2.1 ARMA-type models

The basic modelling approach adopted is to assume that the wind power that can be generated at any time t is a random variable (i.e. a partially predictable quantity with more than one possible outcome) with a Normal distribution. These distributions are assumed to be characterised by a dynamically varying but entirely predictable mean, plus a random 'noise' term, ε_t, with zero mean and a (initially assumed) constant variance. The dynamic mean values are assumed to be functions of recent wind power values, and of recent noise values, and the nature of this function defines the precise model type. More specifically, these functions are assumed to belong to the set of ARMA models (Auto Regressive Moving Average) and their many and varied extensions, referred to collectively as here as ARMA-type models. Finally, the noise terms $\varepsilon_{(t-1)}$, ε_t, $\varepsilon_{(t+1)}$,… are all assumed to be statistically independent of eachother, and so together form what is known as a Gaussian white noise process. These model structures will be presented in the following subsections, but we do not cover the technical details of this approach in great detail here, for the sake of brevity.

One powerful feature of such models is that they allow the process being modelled to be simulated with ease, i.e. one can easily produce very large numbers of traces through time of how the phenomenon being modelled might progress into the future, given current and recent past conditions. In order to conduct this simulation, one begins by generating a set of numbers that satisfy the Gaussian white noise criteria described above, then progressively apply a number of filters and transformations, that represent an alternative way of conceiving of the model for dynamic means.

When fitting the model, we go in the opposite direction, i.e. start with historical wind power series and keep applying filters and transformations until we are left with 'errors' that satisfy the Gaussian white noise conditions. We assume that each possible aggregated combination of wind farms requires a unique set of model parameters, but strive to find a single model structure that is close enough to optimum for each combination, i.e. so there is only one model structure, if possible.

4.2.2 The logit transformation

The first step in fitting the statistical model is to transform the 'random process', which is wind output from the group of wind farms being assessed, into something which is approximately 'Normally' distributed, through the use of a 'logit' transformation. This is almost like transforming from wind output (in MW) back to wind speed in ms-1. The main difference is that, with a 'smoothed' power curve assumed by the Renewables Ninja model, a single power level can correspond to two wind speeds – one on both sides of the maximum power range, while a range of wind speeds correspond to a single (maximum) power level. This contrasts to the logit transformation where there is a one-to-one relationship between the inputs and outputs of the transformation.

The logit transform involves first normalising the wind output, X, so that it takes values only within the range 0–1, and possibly even a smaller range within that, then performing a log-based transformation as follows:

$$X' = \frac{X - a}{b - a}$$

$$Y = \log\left(\frac{X'}{1 - X'}\right).$$

The coefficients a and b are values which are close to the values of the minimum output from the group of wind farms and the maximum output, except that a is a little lower than the minimum output and b is a little larger than the maximum output. For the current model fitting, values were chosen through trial and error across 250 combinations, such that the histogram of Y, taking all ~150,000 values of output across all 17 years, is approximately normally distributed[5].

[4] The term 'random' in this context means 'not predictable with perfect accuracy', rather than 'completely unpredictable'. The term 'random process' refers formally to an ordered collection of random variables, where the outputs of separate random experiments are assigned to each indexed time step t within the set. When simulating from a statistical model consisting of several components, some of them may be kept fixed while others vary randomly, as an extension of the studies presented in the preceding subsections.

[5] This is measured using a statistical test called the Kolmogorov–Smirnov test, which compares the data to a perfect normal distribution with the same mean and variance. We selected values for a and b which provide the lowest score from the Kolmogorov-Smirnov test.

National Grid ESO | July 2019 Black Start from Non-Traditional Generation Technologies

4.2.3 Accounting for long-term deterministic trends

Although highly variable, wind speeds do follow both daily and annual cycles in their mean values when observed in the very long term. It is therefore necessary to account for these patterns when fitting the model. While it was obvious at the outset that the model must capture variability with time of year, it was not obvious that time-of-day effects are large enough to require representation. However, we found that while time-of-day effects are modest, they are indeed significant enough to justify the additional modelling complexity of representing these patterns.

The first step was to compute a matrix of mean values for the logit transformed output for every time of day and every month, across all 17 years of data. This results in a 12 by 24 matrix, with 12 rows for each month and 24 columns for each hour of the day.

These mean values were then "stretched" out to fill an entire 8,760-hour long year, with the pattern of averages repeating every day within each month.

If used directly in the model, this would result in discontinuities in output when moving from one month to the next. Therefore, this series of means is smoothed, by taking a weighted moving average of the values from each day as well as the previous and following 15 days, for the same time of day, as shown in the following equation.

$$Y_t^{**} = \sum_{i=-361}^{361} \lambda_i \cdot Y_t^{*}$$

Where Y_t^* is the stretched-out series of unsmoothed monthly/daily means, and λ_i is the weighting factor, which takes values of 0 for i = -361 and i = 361, a value 1 for i = 0, and varies linearly between these like an isosceles triangle, but taking a value of 0 for values of i which are not a multiple of 24 (i.e. for periods t which aren't the same time of day).

The smoothed series Y_t^{**} is then subtracted from the original logit transformed series Y_t, to give a series from which the long-term deterministic trends have been removed, Y_t'.

4.2.4 Modelling low-frequency anomalies

The next step is to account for the longer-term anomalies in wind output, accounting for the fact entire months typically have mean output values that are significantly lower, or higher, than their long-run seasonal mean. Accounting for these anomalies is achieved by calculating the mean for each month in each year (i.e. 12 values for each of the 17 years to give a total of 204 values), and removing them from the series. It was found that the 204 values of monthly fluctuation are approximately normally distributed, and attempts were made to fit an ARMA model to them. However, the result of this attempted model fitting was that they are best treated as a Gaussian white noise process, which is ideal for model simplicity.

Before the monthly anomalies could be subtracted from the series, they again required stretching out to 8,760 hours for each year ($Y_t^{\dagger}$). Smoothing was again applied before subtraction, as above, except in this case non-zero weights are admitted even for values of i which are not multiples of 24). We define this mathematically as:

$$Y_t^{\dagger\dagger} = \sum_{i=-361}^{361} \lambda_i \cdot Y_t^{\dagger}.$$

This series is then subtracted from the series produced in the previous step:

$$Y_t'' = Y_t' - Y_t^{\dagger\dagger} = Y_t^{\square} - Y_t^{**} - Y_t^{\dagger\dagger}.$$

4.2.5 Modelling short-term variability

After accounting for the longer-term behaviour, the final step of fitting the statistical model is to fit a model to describe the short-term variability. This is achieved using ARMA type time-series statistical modelling, the components of which are described below.

- **Autoregressive (AR) model:** The AR model is used to describe a time-varying random (or partially predictable) process where the value of the process during any given hour t depends linearly on its own values in previous hours t-i, plus a stochastic "error" term. Mathematically, the model is written as

  $$X_t = c + \sum_{i=1}^{p} \varphi_i \cdot X_{t-i} + \varepsilon_t;$$

  where φ_i are the model coefficients, c is a constant, and ε_t is either an error term (when model fitting) or collectively form a white noise process (when using the model to generate synthetic data). The values of i here are called the "lags" of the model, and can take any positive integer value. The model is written as AR(p), so that an AR(2) model would include the first and second lag. The model is essentially a filter on the white noise process, as described in previous sections. AR models are special, simpler, cases of the more general class of Auto Regressive Moving Average (ARMA) models, but it turns out that the modelled series in this case are best described by AR models.
- **Seasonal autoregressive integrated moving average (SARIMA) model:** This is an extension of the ARMA model type, which includes (random) seasonal variation, with some defined period T, as well as the ability to deal with non-cyclical trends, which are not relevant to the series to be modelled here. Such models are specified with a collection of integer-valued parameters, specifically as SARIMA(p,d,q)(P,D,Q), where p and P are the lags and the seasonal lags respectively. The term d relates to the removal of non-cyclical trends, and q refers to the moving average extension that were also found, after investigation, to be unnecessary for our data. Finally, D and Q are the seasonal counterparts to d and q. We have found that the best fitting model has d=q=D=Q=0, so we have therefore not provided any more description of their corresponding model structures in this report.

4 Statistical time series modelling

- **Generalised Autoregressive Conditional Heteroskedasticity (GARCH) model:** The GARCH model is similar conceptually to the AR model, except instead of describing the temporal evolution of the mean of a partially-predictable process, it describes the dynamic temporal evolution of the variance of an uncorrelated, Normally distributed random process. Such models are also filters applied to a white noise input with constant mean, but in this case output an uncorrelated noise process with dynamically evolving variance. Such models allow for "volatility clustering", such as is seen with wind power, i.e. periods of little change in output followed by periods of many large changes. The models, and their associated parameters, are defined by the following equations:

$$z_t \sim N(0, \sigma_0^2)$$

$$\varepsilon_t = \sigma_t \cdot z_t$$

$$\sigma_t^2 = \omega + \sum_{i=1}^{p} \beta_i \cdot \sigma_{t-1}^2 + \sum_{i=1}^{q} \alpha_i \cdot \varepsilon_{t-i}^2$$

  The ω term is a constant. The first summation (over lags up to order p) is essentially an autoregressive model for the dynamic variance, and the second term is a moving-average model, which together make for an Autoregressive moving-average (ARMA) model of the variance of the error term.
- **SARIMA-GARCH model:** This is a combination of the two models, which represents two filters being applied to the white noise input. One filter, the SARIMA model, determines the dynamic mean of the partially predictable output series, while the second filter, the GARCH model, determines the dynamic variance of that output series.

Multiple different combinations of maximum lags (i.e. model orders) for the SARIMA and GARCH model components were explored to determine the best fitting model. In the context of fitting time-series statistical models, best-fitting means there is negligible "autocorrelation" and "partial autocorrelation" in the residual (i.e. unexplained) error of the model, as well as in the squared residuals when fitting the GARCH model. If this is the case, it means that when Gaussian white noise is used to generate data through the model, the results should be statistically very similar to the historical series to which the model was fitted.

The final model for the short-term variability in wind output was ultimately fitted as a SARIMA(6,0,0)(5,0,0)-GARCH(1,0) model[6] with a period of T=24. This means that the dynamic mean of the series is determined by the actual values of the series in the previous six hours, as well as the value at the same time of day during the last five days. The dynamic variance of the series is determined by its value in the previous hours, and it was found thatfound that only the first previous hour needs to be considered for an optimal model. However, in both cases the dependence (as reflected by linear correlation) extends over far longer periods.

4.2.6 Additional transformation for the simulated data

Following application of each filter/model component described above to the historic data series, the resulting series of 'residuals' satisfied, to a very good approximation, the statistically desirable condition of possessing no linear correlations with each other. Additionally, the squares of the residuals were also free of linear correlation, indicating that the model adequately accounts for the volatility clustering in the original series.

As stated above, satisfying these conditions indicate that applying those filters (in the reverse direction compared to model-fitting) to a Gaussian white noise series should produce a synthetic series with statistical properties very similar to the historical series. Upon testing this, it was discovered that the distributions of synthetic values (for a few example windfarm combinations) weren't acceptably close to the corresponding distributions of historical values. The main reason for this is that the 'b' coefficient values in the initial/ final logit transformation need to be larger than the installed capacity of the wind farms in order to result in approximately normally distributed data. This gives rise to a significantly more gradual tapering-off in the simulated distributions in the "right-hand-side" tail of those distributions. It was found, through trial and error, that this discrepancy can be rectified almost completely with a transformation of the form: $y = c \cdot x^{\lambda}$. For each combination of wind farms, the optimal pair of transformation coefficients, c,λ, were found by trialling values of c in the range 0.5–1.5, and λ in the range 0.4–1.6, and again using the Kolmogorov–Smirnov test to determine which pair is optimal.

[6] SARIMA(6,0,0)(5,0,0) is equivalent to AR(1,2,3,4,5,6,24,48,72,96,120), although the latter is not commonly accepted notation. GARCH(1,0) can also be written as ARCH(1).

National Grid ESO | July 2019 Black Start from Non-Traditional Generation Technologies

4.3 Statistical model validation

This section presents a number of plots demonstrating the rigorous validation conducted on the adopted model structure. The validation was conducted on a number of windfarm combinations, with the number of aggregated windfarms ranging from 1–22, and several values in between[7]. For each combination, the validation exercise consisted of comparing many different aspects of the 17-year historical series with five synthetic 17-year series generated by the model, e.g. the distribution of annual means, and the distribution of changes in maximum wind power output over eight hours.

The results showed that the historical and synthetic series were very similar in all aspects tested, and for all windfarm combinations. Further, the differences between the real and historic series were consistent across all windfarm combinations, albeit with those differences generally being somewhat greater for smaller numbers of aggregated windfarms – i.e. spatial smoothing effects favour the model. This section presents a selected subset of the validation plots for only one example of aggregated windfarms – Blacklaw (124 MW) and Clyde 1 (350 MW), which is a poorer than average fit, due to comprising only two windfarms.

For brevity, the precise nature of the validation tests presented here are not described. Instead, the reader can simply observe that plots for historical and synthetic data are similar, or follow closely lines of gradient 1 and intercept 0 in the case of Quantile-Quantile plots[8].

The first plots presented, in figures 4.1 to 4.4 are Quantile-Quantile scatter plots (QQ plots) for the synthetic and historical series, for: (i) the actual series, (ii) differences over eight hours and (iii) differences over one hour.

It can be seen that there is very good agreement for the actual series, and almost perfect agreement for differences over eight hours, as is the case for differences over four hours, 24 hours, 48 hours and 120 hours (five days), although for brevity these are not included. It can also be seen that while agreement is generally good for the distributions of one-hour changes, some considerable outliers exist (that are different for different model runs, which is generally not the case for the other differences). It is postulated that these outliers reflect the fact that very sudden changes can occur in the historical data due to changes in wind speed to the opposite side of the turbine cut-off value, and that the statistical model does not fully replicate such changes – something extremely difficult to achieve. These outliers are not a serious problem in the current context, since it is highly likely that the ESO would require a consistent output from wind for periods considerably longer than two hours.

Figure 4.1

QQ plot of historical and simulated 17-year series (run 1)

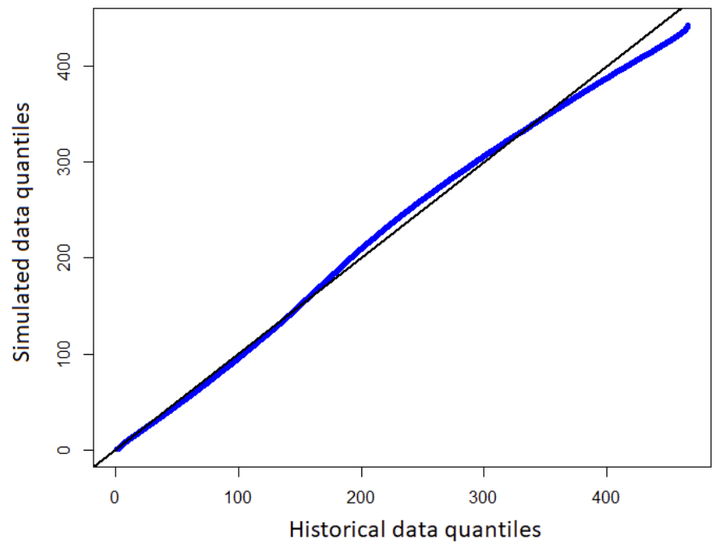

[7] Although we didn't validate all 500 combinations explored in the previous section.

[8] The interested reader may also research the validation methods for themselves, as they are very common in the field of time series modelling, e.g. https://en.wikipedia.org/wiki/Q%E2%80%93Q_plot.

4 Statistical time series modelling

Figure 4.2
QQ plot of changes over eight hours, historical and simulated series (run 1)

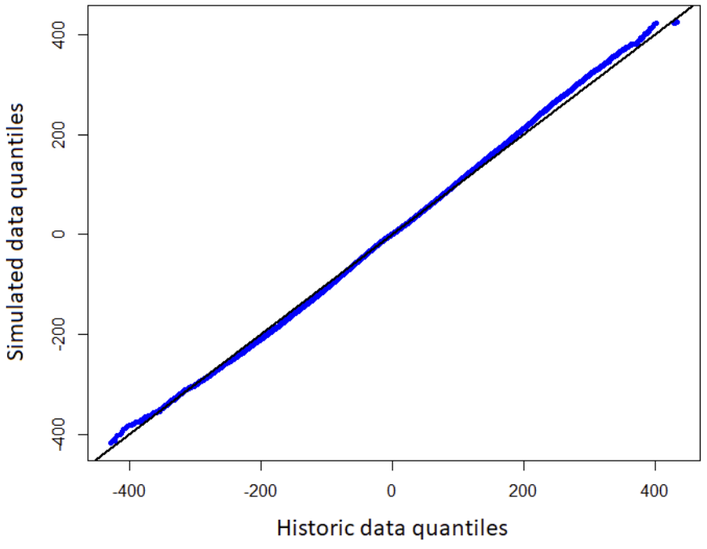

Figure 4.3
QQ plot of changes over one hour, historical and simulated series (run 1)

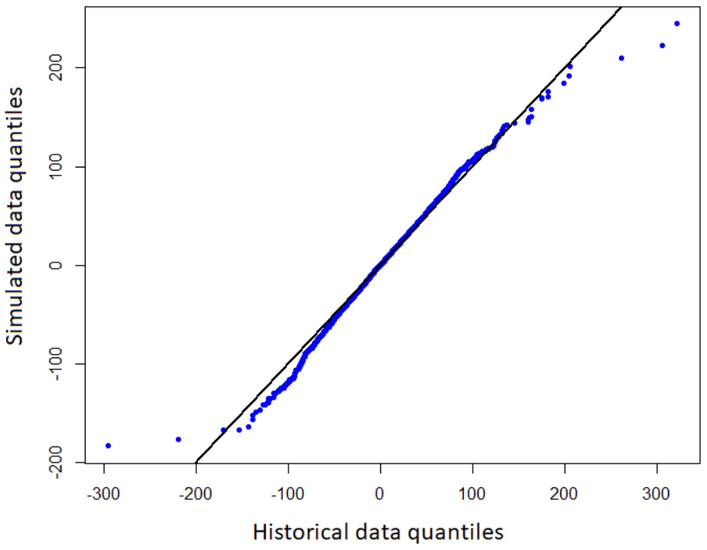

Figure 4.4
QQ plot of changes over one hour, historical and simulated series (run 2)

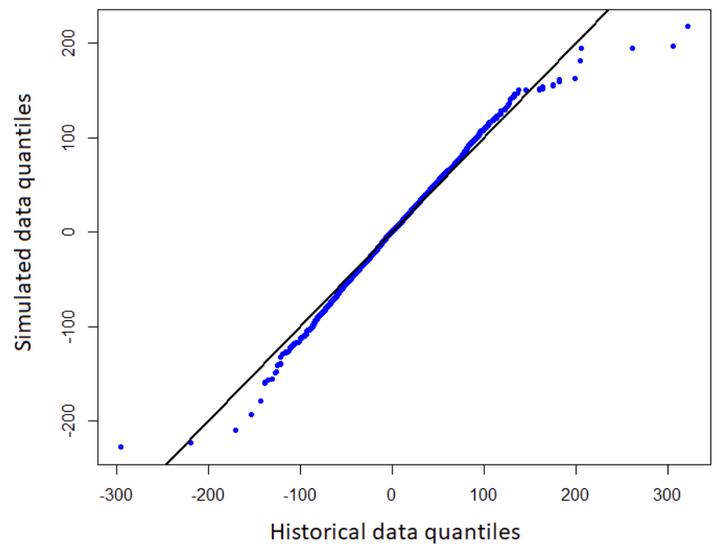

The next set of figures (4.5 and 4.6) compare time series segments for the series after aggregating in two different ways: the first is the series of annual means, while the second is the progression of January means. It can be seen that the agreement for annual means is excellent, while the agreement with January means is quite good, but with the synthetic series displaying smaller variability compared to the historical series. For series involving a greater number of windfarms, the degree of agreement tends to be considerably better, and it must be remembered that a simpler model involving e.g. only an ARIMA or even SARIMA model would perform very poorly if subjected to these tests.

Figure 4.5
Time series of annual means, historical and simulated series (runs 1–5)

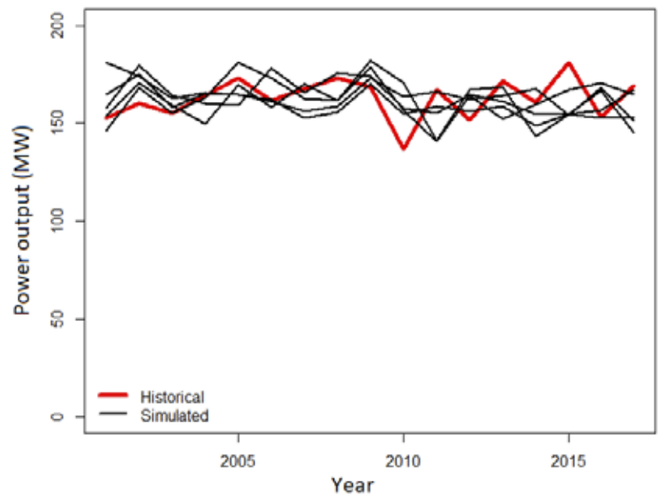

4 Statistical time series modelling

Figure 4.6
Time series of January means, historical and simulated series (runs 1–5)

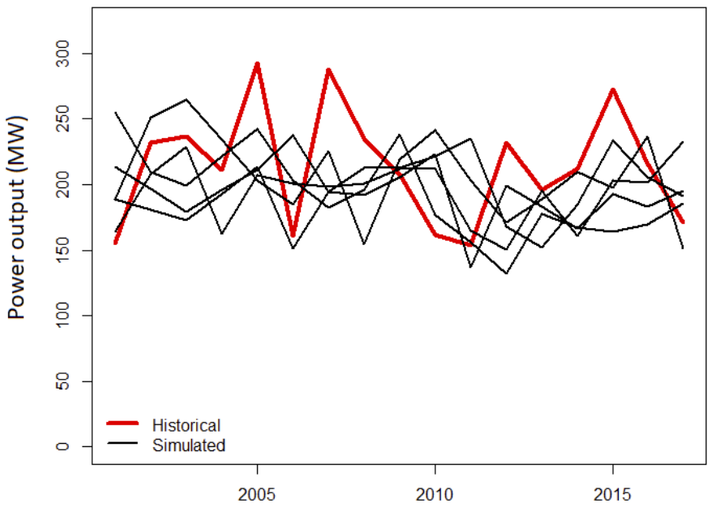

The final set of plots (figures 4.7 to 4.10) present Autocorrelation functions (ACFs) for the historical and synthetic data, with short-range and very long-range plots included. The value of the ACF function for argument k is the linear correlation of the time series with a lagged version of itself, calculated across all time steps in the series, between values at times t and at times t-k. So, where the horizontal axis in the plots has a value of e.g. one hour, the y-axis value is the linear correlation between each instance of X_t and X_{t-1}, and so on.

In fact, the plotting and analysis of ACFs are the primary tool in time series analysis and modelling. It can be seen from figures 4.7 and 4.8 that the simulated series ACF is in almost perfect agreement with the historical series ACF up to a lag of about 35 hours, but declines somewhat more rapidly after that. Figures 4.9 and 4.10 demonstrate that the shape of the synthetic series ACF is in close agreement with the shape of the historical series ACF, but the amplitude of the function is somewhat reduced for the synthetic series.

Figure 4.7
ACF for the historical 17-year series, first 50 lags

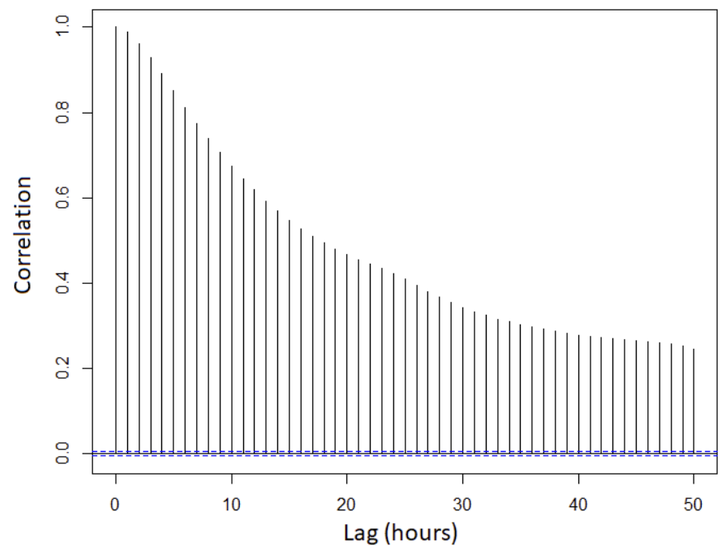

Figure 4.8
ACF for the simulated 17-year series (run 1), first 50 lags

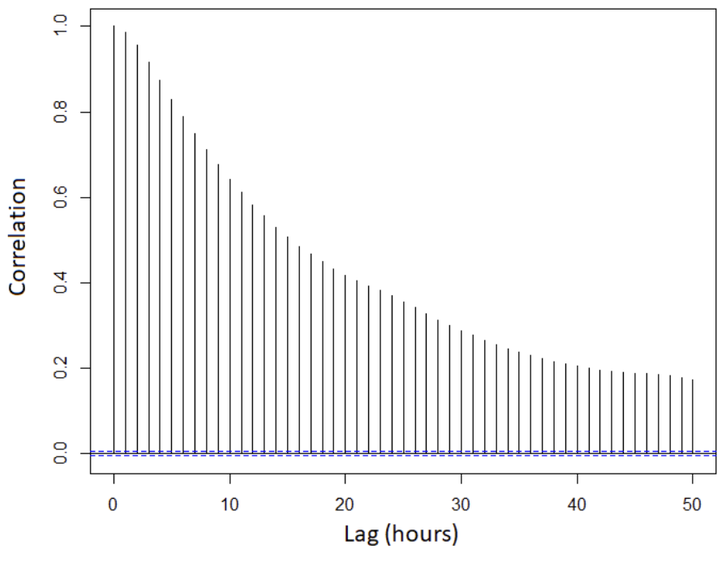

4 Statistical time series modelling

Figure 4.9
ACF for the historical 17-year series, lags 0–30,000

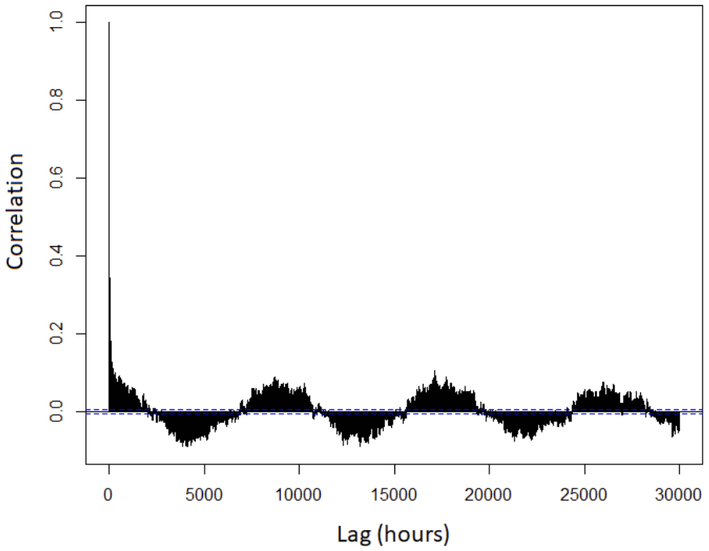

Figure 4.10
ACF for the simulated 17-year series (run 1), lags 0–30,000

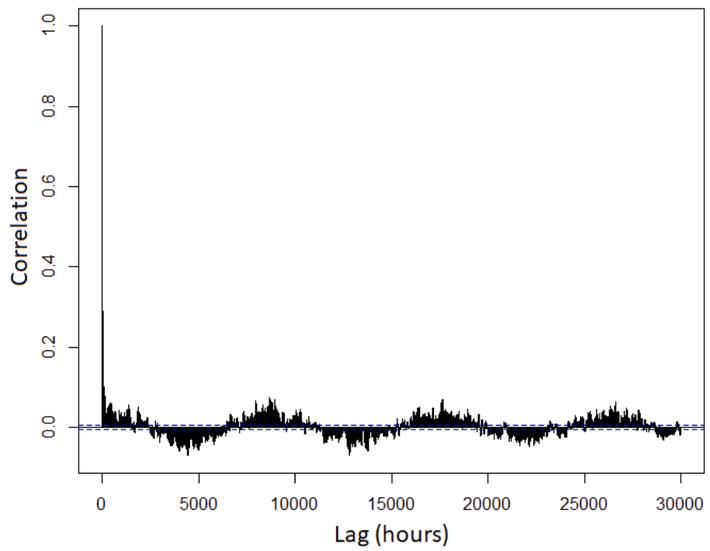

5 Example applications of the statistical analysis and modelling

In this section, we present some examples of the statistical model being used in a variety of useful ways to generate synthetic data, and provide a critical discussion of the model's full range of potential applications. Unless otherwise stated, the results presented in all figures in this subsection are for a single combination of 18 windfarms, with a total installed capacity of 2,534.6 MW.

5.1 Year-round long-term restoration planning

We consider, initially, the case where the ESO is conducting long term planning in relation to system restoration far ahead of real time. In this case, probably the most useful quantities to calculate are quantiles for the unconditional distribution, i.e. over all possible conditions, including all months of the year, times of day and good/poor resource years. To calculate these values, one can choose either of the methods presented in this report (i.e. sample from the historical data or generate synthetic samples using the statistical model). The former approach has the advantages of being less computationally expensive and being more directly representative of reality. As a result, it might be most appropriate where the ESO wishes to investigate many potential combinations of windfarms, and are interested in risk levels such as a five per cent or one per cent chance of failing to maintain the contracted wind power, i.e. the 95th and 99th percentiles. However, if interested in even lower risk levels, such as a 0.1 per cent chance of failure (the 99.9th percentile), the ability of the statistical model to simulate a large number of years give it a significant advantage.

Based on the data analysis presented in section 3, values like the gradients and intercepts presented in the tables might be used in long term planning, where the target amount of sustainable capacity, along with the acceptable level of risk and duration are estimated. This is demonstrated in table 5.1 below, where the values presented in table 3.1 were used to predict the amount of capacity that would have to be procured in order to consistently sustain an output of 500 MW for the required duration and at the required level of risk. This is presented graphically in figure 5.1.

Table 5.1
The average amount of procured wind capacity (in MW) required to secure a sustained 500 MW of power, for various durations and risk levels

Duration	50th percentile	90th percentile	99th percentile	99.9th percentile
1 hour	1,587	5,742	14,368	27,873
4 hours	1,771	6,679	16,750	32,599
8 hours	2,019	7,816	18,731	36,635
12 hours	2,267	8,781	20,410	40,343
24 hours	3,023	10,798	23,856	52,446

5 Example applications of the statistical analysis and modelling

Figure 5.1
The average amount of procured wind capacity (in MW) required to secure a sustained 500 MW of power, for various durations and risk levels

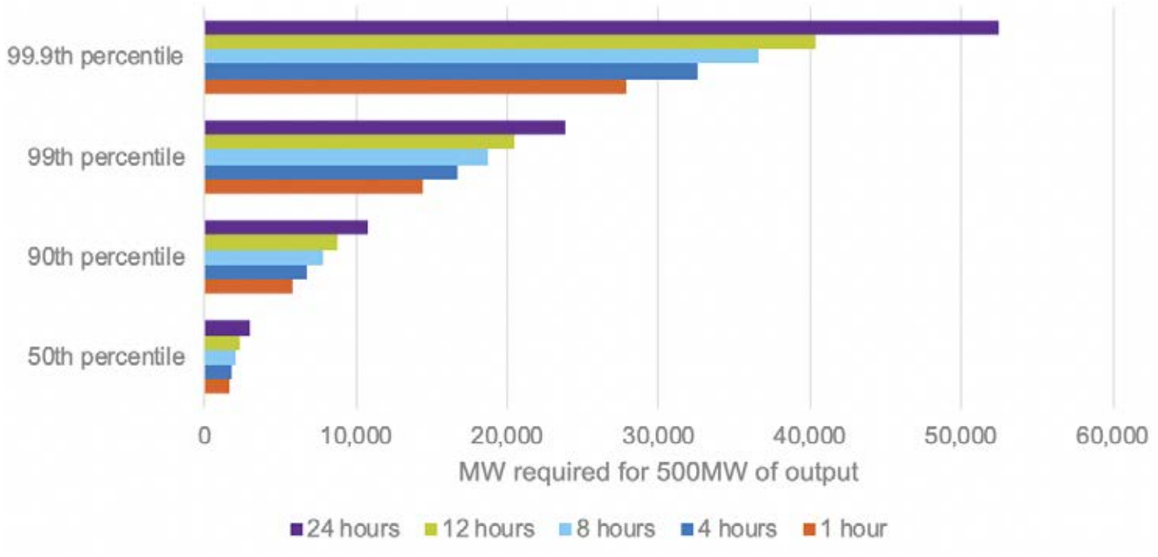

5.2 Restoration planning specific to season and time-of-day

Another application of interest for long-term planning is to produce quantile values that are specific to the time-of-day and the month or season. This would be particularly useful if it was known that the risk-based availability of other resources procured for system restoration also possess such seasonal dependencies, or if there are other seasonal/time-of-day factors that are relevant to variation.

If considering only variations in wind availability by month, it is not clear whether the historical data sampling or the statistical model approach would provide the best results. However, if conditioning the results on both the month and time-of-day, the statistical model approach seems to have a clear advantage, again due to the finite historical sample size. Examples of this are presented in figures 5.2–5.5, which are histograms of maximum sustainable power for all combinations of the following: starting datetimes of 12am in January and 12pm in July; required durations of eight hours and 120 hours (five days). These plots provide some interesting insights on the way different quantiles are affected to different extents by changes in starting datetime and required duration.

To produce these plots, we produced 10,000 simulated time series segments, referred to as traces, from the SARIMA-GARCH model for short-term variability, each with different starting conditions. To each trace we added a different simulated slowly-varying trace, representing the smoothed monthly anomalies. After that, the same deterministic profile for seasonal means was added to each trace (as we are conditioning on a fixed month and time-of-day for the distributions), and finally the traces were transformed via the reverse logit and power transform.

Figure 5.2
Histogram of max. sustainable power for eight hours, starting at 12am in January

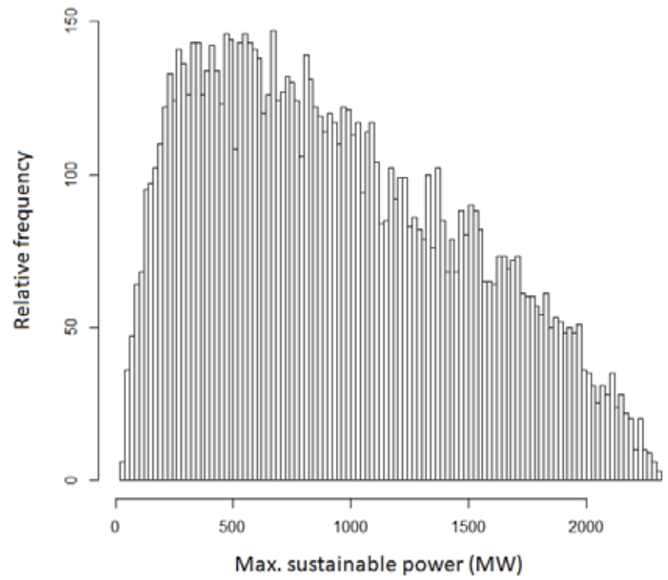

Figure 5.3
Histogram of max. sustainable power for eight hours, starting at 12pm in July

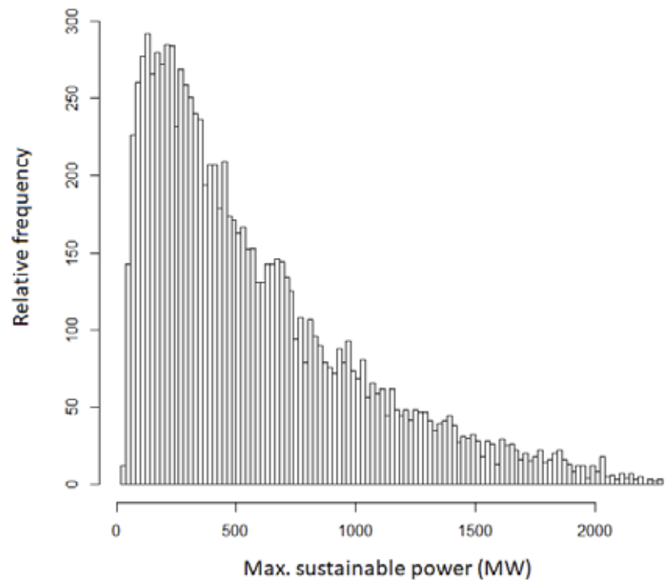

5 Example applications of the statistical analysis and modelling

Figure 5.4
Histogram of max. sustainable power for five days, starting at 12am in January

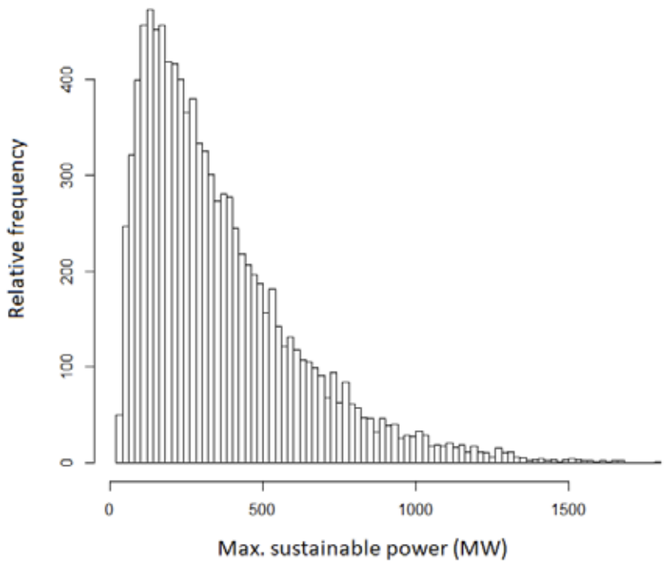

Figure 5.5
Histogram of max. sustainable power for five days, starting at 12pm in July

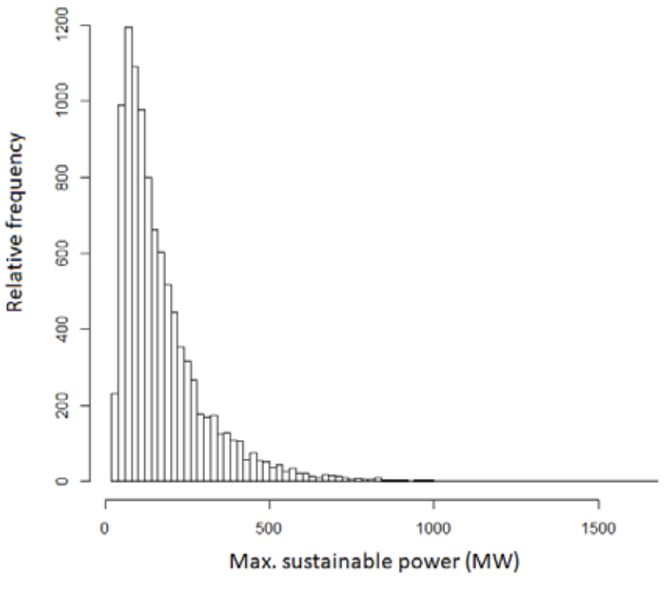

5.3 Restoration planning on a semi-operational time scale

In the next demonstration, presented in figures 5.6–5.11, we return to considering specific quantiles, in this case the 95th and 99.5th, and examine in isolation the effect of the low frequency anomalies (or in statistical terms, the long memory in the process) as described in section 3.4. This allows us to show, in isolation, the effect that these low frequency anomalies have.

To achieve this, we followed the procedure below:

- **i** simulate 1,000 traces from the short-term SARIMA-GARCH model, again with different starting conditions for each one;
- **ii** produce a single sample trace for the low-frequency process and add it to the 1,000 short-term traces;
- **iii** add a fixed trace of deterministic seasonal means to each trace, as in the previous example;
- **iv** transform the values to wind power outputs, through the reverse logit and power transformations;
- **v** calculate the minimum power for each trace;
- **vi** calculate the required quantile for the distribution of trace-minima;
- **vii** repeat stages i–vi, 100 times, with the same set of traces produced in step i, the same deterministic means added in step iii, but different samples drawn for the low-frequency variability in step ii
- **viii** store the 100 quantile values from step vi as a distribution, and plot them as a histogram.

This application of the simulation from the statistical model allows the ESO to mimic, on a planning time scale, the situation on a semi-operational scale where the assumed availability of wind for system restoration, is re-assessed dynamically e.g. each month or each fortnight, depending on the extent to which the season appears to be good or poor in terms of the wind resource. The results clearly demonstrate that the low-frequency variability is very significant, and as a result there could be considerable value in this type of dynamic re-assessment of wind's ability to support restoration.

Figure 5.6
Histogram of 95th percentiles of maximum sustainable power for eight hours, starting at 12am in January, over possible conditions of low-frequency variability

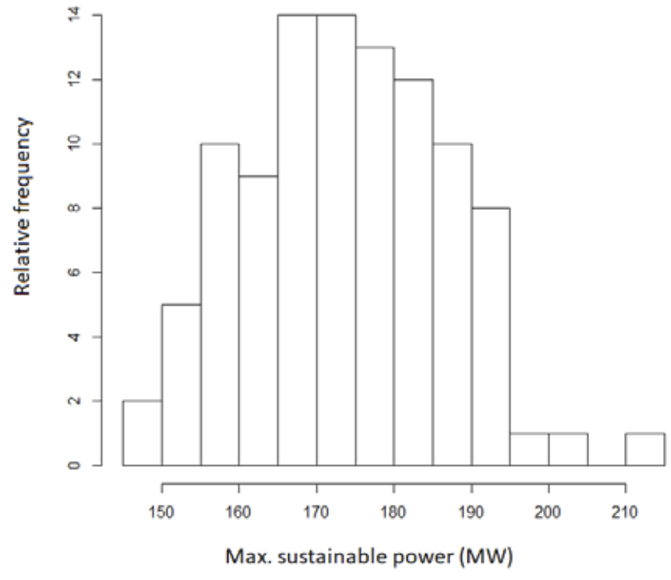

5 Example applications of the statistical analysis and modelling

Figure 5.7
Histogram of 99.5th percentiles of maximum sustainable power for eight hours, starting at 12am in January, over possible conditions of low-frequency variability

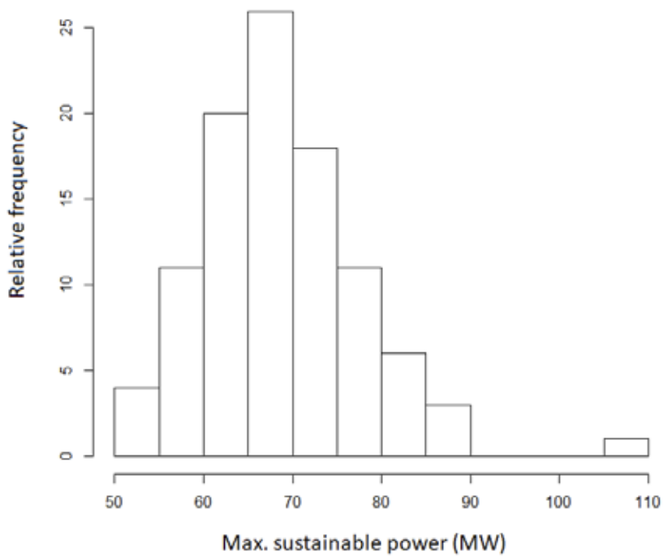

Figure 5.8
Histogram of 95th percentiles of maximum sustainable power for eight hours, starting at 12pm in July, over possible conditions of low-frequency variability

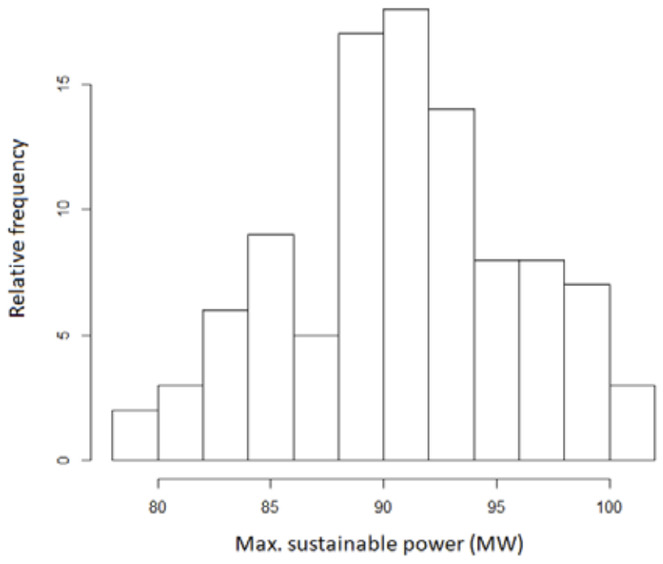

Figure 5.9
Histogram of 99.5th percentiles of maximum sustainable power for eight hours, starting at 12pm in July, over possible conditions of low-frequency variability

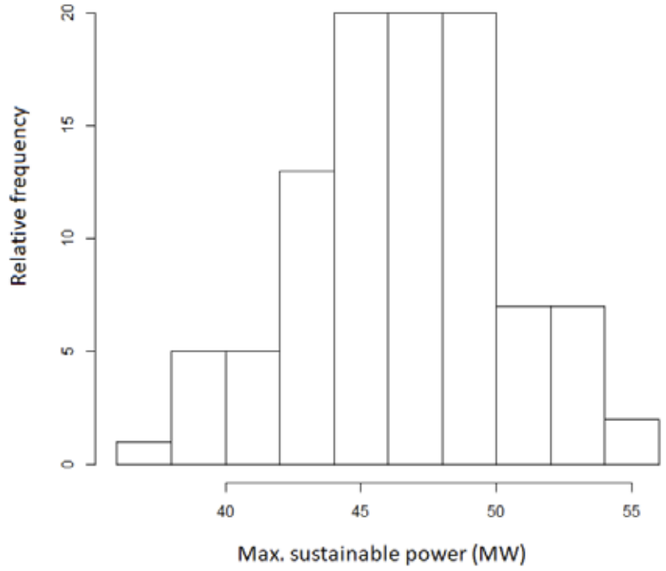

Figure 5.10
Histogram of 95th percentiles of maximum sustainable power for five days, starting at 12am in January, over possible conditions of low-frequency variability

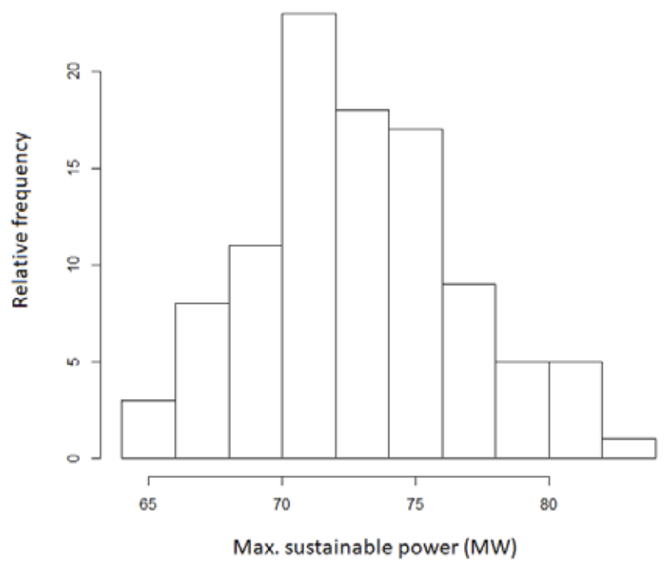

5 Example applications of the statistical analysis and modelling

Figure 5.11
Histogram of 95th percentiles of maximum sustainable power for five days, starting at 12pm in July, over possible conditions of low-frequency variability

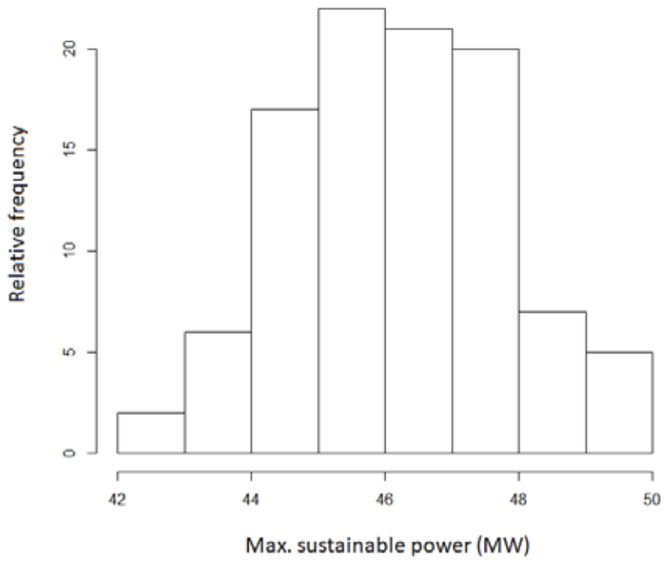

5.4 Nearly real-time restoration planning

It might also be the case that the ESO could re-assess the likely contribution of wind generation to system restoration on a more frequent basis, e.g. weekly or daily. This operational-scale process could be mimicked on a planning scale by using the statistical model in a similar way to the procedure outlined above. The differences for such applications are that:

i the low-frequency variation part of the model would be kept fixed during the 100 repetitions;

ii within each of the 100 repetitions the 1,000 traces from the SARIMA-GARCH model would start from a single sample for the initial conditions; and

iii each of the 100 repetitions would involve different sampled initial conditions for the SARIMA-GARCH traces.

The ability to conduct such sophisticated planning-scale studies is what makes the model so powerful.

It would be useful to have a model that can actually adjust the likely contribution of wind capacity to system restoration on an operational time-scale, rather than simply mimic this on a planning time-scale. The model presented here would indeed be able to do this well in the first case of monthly and possibly fortnightly reviews. However, the model could not be considered cutting-edge in its ability to do this on a daily basis or similar. In order to do that, the model would have to be expanded to include meteorological forecasts, ideally of a probabilistic nature, as is currently used for forecasting by the ESO. We believe that this would be a fairly substantial modelling challenge, but that the model's outputs would be of great value to the ESO.

5.5 Wider applications of the model

Although the data analysis and statistical model have been discussed in the context of Black Start only, both could have much wider applications to many aspects of transmission and distribution system planning and operation. The level of reliable output that can be provided by variable generation is an important consideration for network planning, ancillary services procurement (e.g. frequency response from wind), and dispatch/balancing. Some similar analysis was carried out as part of the consultation on the latest draft version of Engineering Report (EREP) 130[9], which is a supporting document to the P2 planning standard.

One clear application would be updating the "scaling factors" used within the National Electricity Transmission System Security and Quality of Supply Standard (NETS SQSS) to set the level of output from renewable generation for transmission system planning. Currently, two planning backgrounds are considered, one where wind (and other variable generation) output is set to 70 per cent of registered capacity, and the other where it is set to 0 per cent.

One important insight from this analysis and modelling, which would also be relevant for any other applications, is that relying on variable generation for any type of system planning or security inherently carries risk. It should therefore be thought about explicitly in terms of probabilities, recognising that the risk varies based on several factors including the specific combination of variable generators and the required duration of their output.

[9] The consultation documentation is available on the Distribution Code website at, under DCRP/19/02/PC http://www.dcode.org.uk/consultations/closed-consultations/

6 Possible extensions of the modelling and analysis

In this section, we describe some possible extensions of this modelling and analysis which could yield additional insight for ESO or support other applications.

6.1 Forecasting tool

As described in the previous section, the model developed is primarily a simulation tool for use on planning timescales, but could be combined in a mathematically rigorous way with meteorological forecasts using cutting-edge data science techniques, probably involving ensemble methods.

6.2 Other types of variable generation

The principles of the model presented here could relatively easily be applied to other variable generation, particularly photovoltaic (PV) solar panels. Of course, the details of the model would change, but the basic steps and philosophy would be very similar. In fact, the Renewables Ninja reanalysis data can also be used to generate solar PV output data.

One potential challenge for PV is that although deterministic trends across the year are stronger, the short-term variability around this can be quite extreme, due to cloud cover. It is possible that this effect may not even be observable in data with hourly granularity, and possibly other approaches will be needed to account for this.

6.3 Energy storage

Combining energy storage with variable generation could be a potentially attractive way to boost the contribution that this generation could make[10]. The analysis we have done has studied the minimum level of maximum output from variable generation over different durations. It is possible that, particularly during periods of high wind volatility, there may be times when this is limited by sharp drops in output which only last for an hour or two. The higher output before and after this drop is therefore not helpful, as it does not allow the ESO to displace conventional plant or switch on additional block loads.

In such cases, there would be a clear role for storage to increase the minimum level of output by shifting variable output in time. For example, in figure 6.1, the drop in output during hours t+5 and t+6 limits the output of this variable generator to only around 20 MW. If the 'excess' output during hours t+3 and t+4 were used to fill up some energy storage, which could export during hours t+5 and t+6, this would allow for faster restoration or less reliance on traditional Black Start providers.

[10] Elsewhere in the project, we have highlighted that storage is potentially very well suited to supporting restoration due to its technical capabilities (e.g. for managing frequency or voltage).

National Grid ESO | July 2019 Black Start from Non-Traditional Generation Technologies

Figure 6.1
Illustration of role of storage in supporting output from variable generation

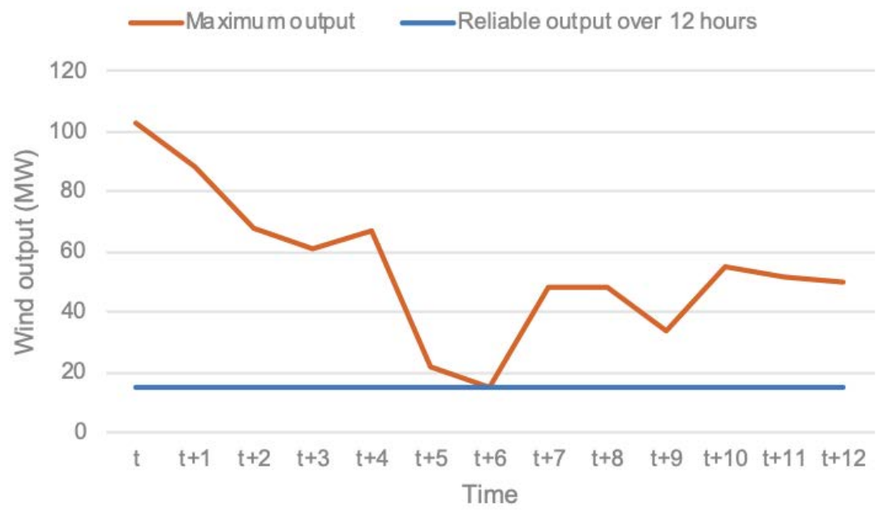

The extent to which this is possible would depend on the condition of the energy storage (e.g. a battery's state of charge) at the start of and during the restoration. One option might be to just mandate certain states of charge. However, this may be very complex to achieve, as the optimal state of charge for a battery will vary depending on the exact nature of its role in the restoration. For example, for a grid-forming battery, the optimal state of charge would be for it to be full, whereas for one that was supporting a variable generator, it might be for the battery to be around half-full so that it can both absorb and inject power as necessary.

Rather than try to directly control the condition of the storage, an alternative option would be for ESO to predict this based on historical data. This would probably involve a model similar to the one we have used for wind in this report. It is likely that there would be strong seasonal and time-of-day patterns in the state of charge of a battery, for example, based on how it responds to external market conditions, which are themselves closely correlated with weather, season, and time-of-day. This would require data about states of charge of batteries or the condition of other storage. If not directly available, it might be possible to determine this based on historical metering, or possibly from data submitted into the Balancing Mechanism.

6.4 Spatial modelling

In this report, we have only considered the impact of variability of a very specific combination of wind farms in a single Black Start zone. An obvious next tep, in principle, would be to continue carrying out similar analysis for greater combinations of wind farms both in Scotland and other zones, including current and future wind farms.

However, this is likely to require considerable human and computational effort. Instead, it would be worth exploring means for doing this more efficiently. For example, one option might be to train a machine learning algorithm to assess the output from a group of wind farms, based only on some information about their capacities and locations. This could be done using direct observations of the data (as in section 3), but could also be used to predict the parameters of a statistical model (as in section 4) which describes their output. This would require re-running the model and analysis in this report on many different combinations to produce a training data set, which could then be used with a relatively simple machine learning regression model. This would make the use of this model far more practical as a planning tool, rather than requiring lots of time-consuming model fitting to be completed every time a new set of wind farms are considered.

National Grid Electricity System Operator
Faraday House
Warwick Technology Park
Gallows Hill
Warwick
CV34 6DA
United Kingdom

Registered in England and Wales
No. 11014226

nationalgrideso.com

nationalgridESO