WIND POWER ECONOMICS
RHETORIC & REALITY

VOLUME II

The Performance of Wind Power in Denmark

Gordon Hughes

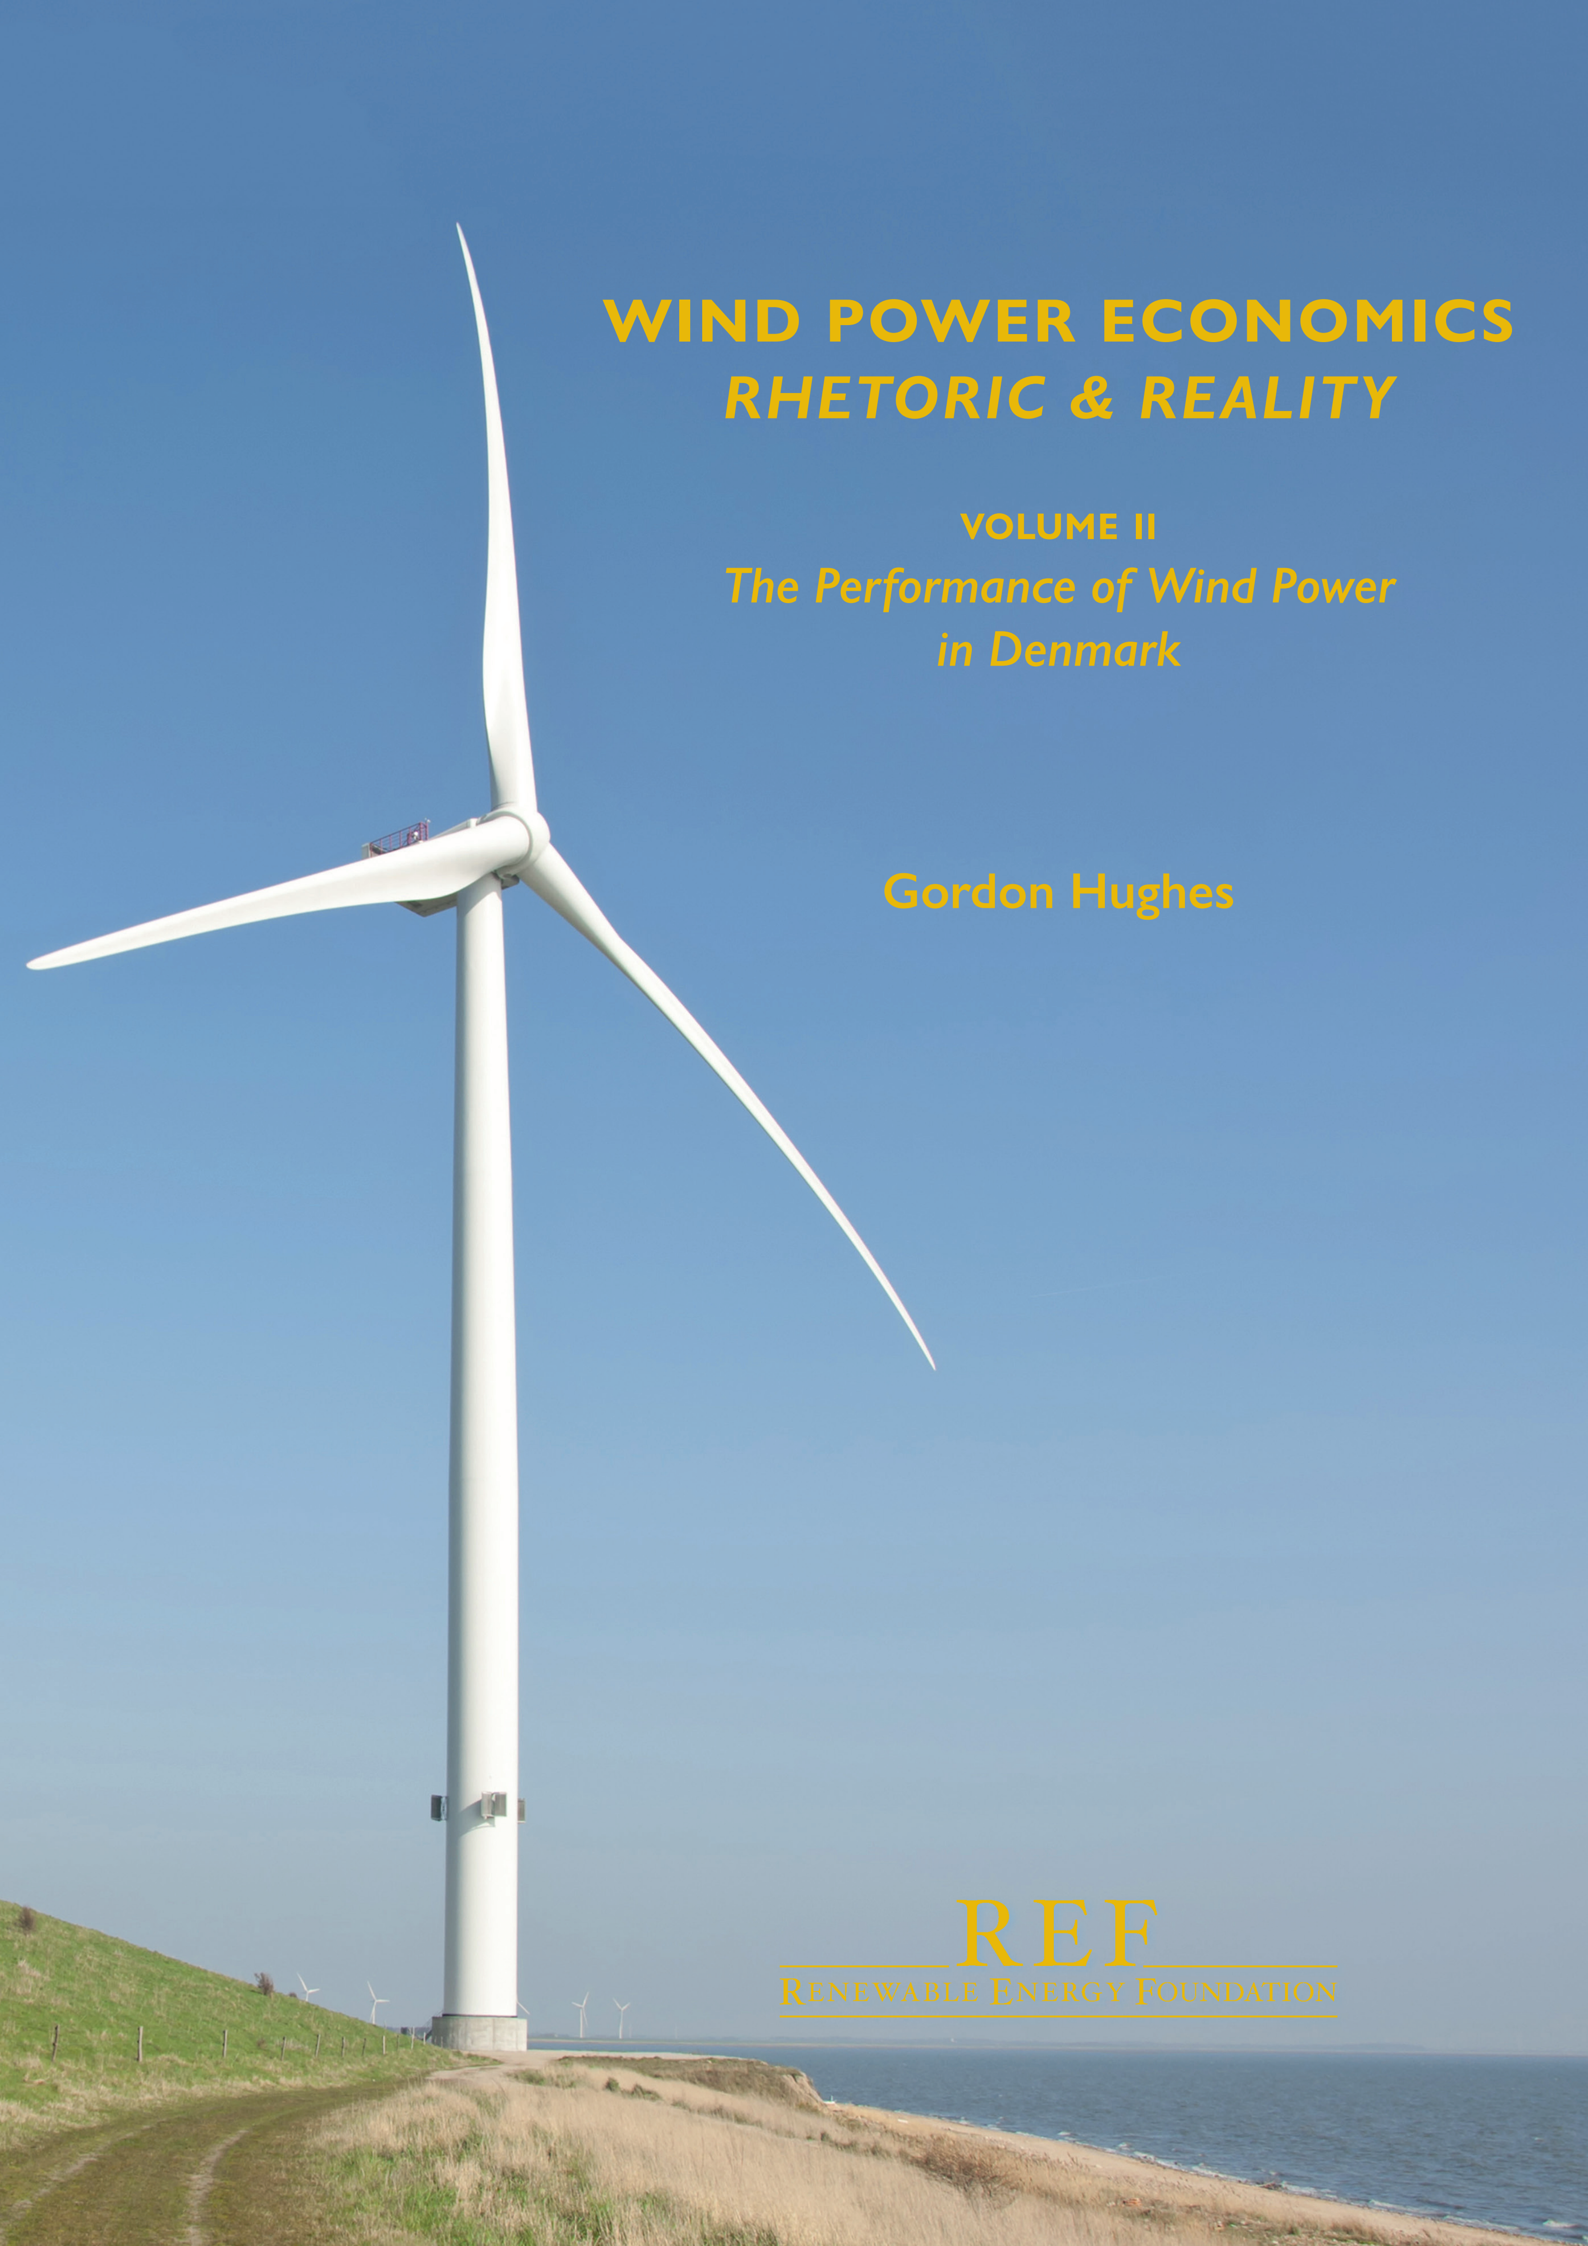

REF
RENEWABLE ENERGY FOUNDATION

WIND POWER ECONOMICS
RHETORIC & REALITY

VOLUME II
The Performance of Wind Power in Denmark

Gordon Hughes
School of Economics, University of Edinburgh

REF
Renewable Energy Foundation

Published by Renewable Energy Foundation
Registered Office
Unit 9, Deans Farm
Stratford-sub-Castle
Salisbury
SP1 3YP

The cover image (Adobe Stock: 179479012) shows a coastal wind turbine in Esbjerg, Denmark.

www.ref.org.uk

The Renewable Energy Foundation is a registered charity in England and Wales (No. 1107360)

CONTENTS

THE PERFORMANCE OF WIND POWER IN DENMARK: SUMMARY

1. This paper re-examines and extends my 2012 analysis of the relationship between age and performance of wind turbines in Denmark using a much larger dataset for the period from 2002 to 2019. It focuses on differences between different generations of turbines, in particular between (a) the small turbines with a capacity of less than 1 MW that were standard until the early 2000s, (b) larger turbines, especially ones with a capacity of 2+ MW, that became standard from 2005 onwards, and (c) offshore turbines – initially with a capacity of up to 3.6 MW but more recently with a capacity of 6+ MW.

2. Ageing may affect the performance of wind turbines in two ways. First, they may experience major breakdowns or less serious equipment failures which cause a loss of output until they are repaired. Electrical failures at substations or in transmission lines may also lead to a partial or total loss of output from a wind farm. Offshore facilities are notably vulnerable to such failures. Second, the yield from turbines may decline gradually as a result of blade erosion and other factors which reduce the aerodynamic or mechanical performance of turbines.

3. Major breakdowns and equipment failures have been analysed using standard reliability models to estimate the time to first breakdown or failure and then recurrences later in the life of the turbines. The results of the models show that small onshore turbines of < 1MW were significantly more reliable than onshore turbines of 2+ MW. Most of the Danish turbines in the 2012 study were small turbines; it is now clear that this difference in composition accounted for much of the difference between Danish and UK turbines in that study. Offshore wind farms with turbines from 2 to 8 MW are likely to experience a much higher rate of breakdown and failure than onshore turbines. About 80% of offshore turbines experience a major breakdown or equipment failure in their first 8 years of operation as compared with only 20% for small onshore turbines.

4. Problems of reliability combined with more gradual loss of output lead to a significant decline in the average load factor for large onshore and offshore turbines as they age. Identifying the effect of ageing is complicated by technical progress in operational and maintenance practices that increases load factors gradually over time. In the case of small turbines, the combined but offsetting effects of ageing and technical progress lead to a small decline in average load factors as turbines get older. In this respect the current results are quite similar to those reported in the 2012 study.

5. However, for larger 2+ MW onshore and offshore turbines, the age-related decline in performance is much more significant, amounting to a decline of about 3% per year for onshore turbines and a decline of as much as 4.5% per year for offshore turbines. At age 16, an onshore wind turbine or farm will only produce 63% of the output that it produced at age 1, while for offshore turbines the output will be down to 50% of peak output. In addition, operating and maintenance costs per MWh of output are likely to rise at similar rates, both to fix breakdowns and failures and because fixed costs have to be allocated across a declining output. These factors mean that the economics of keeping a turbine or wind farm in operation look increasingly unattractive once it passes an age of 16 years.

6. The results suggest that new generations of turbines seem to experience an extended period of low reliability and poor performance following their introduction. This period may last for a decade or more. Experience suggests that the causes of such initial problems can eventually be dealt with but there is clearly a lengthy period of learning for each new generation that should be taken into account. This is important for offshore developments because both manufacturers and operators switched to a new generation of 6+ MW turbines in the mid-2010s. There is a strong possibility that these will experience significant problems of low reliability and poor performance into the latter half of the new decade. In addition, manufacturers are contemplating a shift to yet another generation of 12+ MW offshore turbines with all of the likely problems associated with such transitions. It is far from clear that the manufacturing industry, which does not have a strong financial record or reserves, has the resources to cope with generational changes at this rate.

7. In recent years there have been a number of auctions linked to the development of offshore wind farms which have received extensive publicity, including the Kriegers Flak project in the Danish sector of the Baltic Sea. In this case the power purchase contract price *excluding transmission costs* is €49.9 per MWh (fixed in nominal terms) for 11-12 years. Many commentators assume that the headline prices for such auctions provide reliable evidence for trends in the cost of offshore wind generation. Such claims do not take account of the "winner's curse" that is often an important factor in such auctions. A financial model of the Kriegers Flak project is used to examine the economics of the project and the implications of the decline in the reliability and performance of offshore turbines due to ageing.

8. The assumptions required to justify the auction price set are extremely optimistic, even without allowance for the decline in performance with age. As an illustration, one break-even scenario for the project assumes that the market price of power for the DK2 (East Denmark) region region will rise at more than 5% per year in real terms for 25 years along with eliminating the discount for wind generation. If reasonable allowance is made for forecast risks, the net present value of the project is a loss of about €400 million. If both forecast risk and performance decline are taken into account the break-even contract price is about €85 per MWh, or 70% higher than the auction price.

9. In effect, the winner's curse associated with the auction for Kriegers Flak has transferred a large part of the subsidy required by the project from Danish electricity consumers to Swedish taxpayers who own Vattenfall. While residents of Denmark may welcome the transfer, the outcome is not a realistic way of funding long term investment in offshore generation. Further, Danish investors should not bail out Vattenfall by buying into the project, either during development or after completion, unless this is on terms that the developer is likely to find very unattractive.

10. More generally, financial investors such as PensionDanmark need to exercise extreme caution when evaluating investments in offshore wind projects for which development rights have been awarded by some form of auction. The risks and potential size of the winner's curse are so large as to mean such investments will be unlikely to earn a satisfactory return even in the best of circumstances. There is a high probability of losing a large portion of the money invested in such projects. This is because the developers will be unwilling to accept a valuation of the projects that involve substantial write-downs to reflect a reasonable assessment of the risks and probable revenues from the projects.

11. Headline prices from auctions for offshore wind do not provide a reliable guide to the future cost of offshore wind generation. Nonetheless, it is reasonable to conclude that the cost of offshore wind generation is falling, though not as rapidly as those who rely on auction prices as evidence may believe. Allowing for the risks concerning future performance and power market prices, a reasonable estimate of the cost of offshore wind power would be €70-80 per MWh excluding transmission costs. Even then, operators seem to be assuming higher levels of reliability and performance than can reasonably be expected on the basis of evidence from the last 10-15 years. There is a long way to go to reach "grid parity", even when no allowance is made for intermittency.

THE PERFORMANCE OF WIND POWER IN DENMARK

Gordon Hughes
School of Economics, University of Edinburgh

1. Background

In 2012 the Renewable Energy Foundation published my paper on the performance of wind farms in Denmark and the UK – Hughes (2012). It argued that the economic life of many wind farms would be much shorter than claimed because the evidence which we had collected indicated that the performance of wind farms declined over time, meaning that it would be uneconomic to operate them after 15-20 years. The rate of decline in performance was significantly higher in the UK than in Denmark but it was economically significant in both countries. There was some limited evidence which suggested that newer and larger turbines were less reliable than the earlier and smaller ones used particularly in Denmark. In addition, at that time it seemed that the time profile of the performance of offshore wind farms (all in Denmark) was especially poor.

My paper generated a large number of comments, especially from those in favour of renewable modes of power generation. There were two main lines of comment. The first and most common was that "we don't see this performance decline in …" [name your favourite wind farm]. Even when correct, this type of claim is no more reliable than purported rebuttals of the link between smoking and early mortality based on the claim that "my father has smoked 30 cigarettes a day for 60 years and is still fit and well". Of course there are exceptions: a relationship that may apply to a population need not apply to every member of that population. The difference between wind farms in the UK and Denmark suggests that there are many factors that may play a role in affecting performance. Wind farms in Denmark had lower average load factors, operated in less windy locations and the turbines were typically smaller. Such differences affect the stresses on turbine components and may be reinforced by maintenance practices.[1]

A second line of commentary suggested that the statistical methods used were mistaken or could not be relied upon. A central argument was that it is not possible to separate the effects of ageing from the variations in performance over time due to differences in wind conditions. This issue was

1 There were a number of comments challenging my results for Denmark, which seemed to hit a sore spot even though the average decline in performance with age was lower than that for the UK. In January 2013 I wrote a technical note, published on the REF website, which responded to these comments by providing some detailed analysis of the residuals for the statistical models estimated for Denmark. It is easy to point to outliers because the data is noisy but there is no evidence of systematic bias which would undermine the conclusions.

discussed in Appendices C and D of the original paper, which showed that the method used was more efficient than the option of using a monthly index of wind speeds. In an interesting follow-up to the paper, Staffell & Green (2014) used an approach which analysed our UK data using NASA data on meteorological conditions. In my view their model is statistically biased because of what is known in the literature as the "errors in variables" problem, but in practical terms they too found a decline in performance with age. The difference between the two sets of estimates for the UK is simply a matter of the rate of decline. However, the NASA data is unquestionably useful, so I have used it to extend the original analysis using a different modelling approach.

I have been asked frequently as to whether I have updated or extended my original analysis. Until now, the answer has been that I have not had the time or energy to do this. There was another element to my reluctance to return to the issue. Views about renewable energy have become so polarised that few want to put up with the abuse that follows suggestions that the most optimistic view of renewable generation might not be correct. All of the evidence that I examined was – and is – available to manufacturers and lobbyists but they have little interest in promoting anything other than the most glossy and unrealistic views of wind generation. All technologies involve trade-offs. The pretence that these trade-offs do not exist does little credit to the industry and means that the inevitable failures provoke a loss of credibility. The implications of issues that are regularly discussed inside the wind industry are routinely denied by lobbyists. I do not expect an outbreak of harsh reality, but admitting the engineering trade-offs which have to be made would be a first step.

One reason to return to the issue is a series of developments that were beginning to be apparent in 2012 but which have become steadily more important. The first is the impact of generational changes in the scale of turbines. Most of the turbines in operation in Denmark in 2012 were from the first major generation of commercial-scale turbines with a generating capacity of less than 1 MW. However, from the early 2000s turbines of 1-2 MW and then 2 MW gradually came to dominate new wind development. Since 2012 a new generation of turbines of 5+ MW, primarily for offshore use, has been introduced. Such changes in power technology are often accompanied by major problems of performance and reliability – this has been very clear in the case of Combined Cycle Gas Turbines. It is, therefore, interesting to examine whether there are significant differences in performance between the different turbine generations.

The second development is the shift to offshore sites for large scale wind generation. Wind power is land-intensive; it is clear that the future of wind power is offshore in countries with lengthy coastlines, high population densities and heavy use of non-urban land for agriculture. Offshore wind resources are more reliable, developments can be larger in scale and much bigger turbines can be deployed. Even so, the offshore environment is hostile and expensive for operators. Denmark has been a pioneer in offshore wind generation, so it is particularly interesting to examine the performance of Danish offshore turbines since the early 2000s. In 2012 the data on the performance of offshore turbines was very limited but it suggested that their reliability might be rather poor. This is something that deserves further investigation.

A third issue concerns the role of technical progress, not just in the capacity of new turbines but in the operation of the existing capital stock. In the case of, for example, aircraft manufacturing, it is well known that the combination of economies of scale and learning by doing have brought down the cost of producing standard aircraft – costs of building Boeing 707s have been studied in detail – quite substantially over time. At the same time, improvements in aircraft operating and maintenance practices have improved reliability and reduced the costs of providing passenger and freight services. Economists refer to such developments as embodied (in capital equipment) and disembodied technical progress. An example of embodied technical progress is the change in the shape of turbine power curves, so that turbines reach their rated (maximum) capacity at lower wind speeds and, thus, achieve a higher load factor for the same distribution of wind speeds.

The statistical model used for the original analysis allows for embodied technical progress since it is captured by the fixed effect for each wind farm that is allowed for by subtracting the mean values for each wind farm from each of the observations for that wind farm. However, it is more difficult to separate the effects of disembodied technical progress from those of ageing, and to achieve this I have used a variety of specifications and methods of estimation. A significant part of the difference between the effects of ageing estimated by Staffell & Green and myself arises because the specification of their model does not allow for disembodied technical progress. For technical statistical reasons, the identification of the separate effects relied upon the presence in the dataset of a sufficient number of wind farms that started operation well before 2002, the start date for the turbine register used as the source of our data. The closure or repowering of older wind farms and the growth in the number of new wind farms means that the original approach is no longer satisfactory and a new approach is required to deal with the issue of technical progress.

The other side of the coin is how to deal with major breakdowns and turbine decommissioning – not technical progress but failures. All types of equipment fail sooner or later. The analysis of failure curves is an essential part of performance analysis. In addition to examining data on operating turbines or wind farms, it is important to consider the ones that dropped out of the dataset because they are no longer reporting any output.

For these reasons the scope of this paper is considerably wider than my original paper. Because the data on wind power in Denmark is so extensive, the paper focuses exclusively on the performance of onshore and offshore wind turbines in Denmark. Denmark has been a pioneer of wind power, so the findings are relevant to any assessment of wind power in NW Europe. This is particularly relevant for Germany, which has much to learn from the experience in Denmark and has similar conditions in the main areas of wind development. There will be a parallel paper dealing with the performance of wind and solar power in the UK, but the current paper stands on its own as an analysis of what can be learned from the experience of wind generation in Denmark.[2]

With the goal of ensuring that the results of my analysis are relatively accessible to those without a substantial background in econometrics/statistics, I have put much of the more technical material in Appendices. The main text is designed to provide a (relatively) non-technical presentation of the analysis with a heavy emphasis on summary graphs. Still, the technical details matter; those wishing

2 This paper and the parallel one for the UK are part of a series of papers analyzing aspects of renewable power generation. In particular, I am finalising papers on: (a) the impact of intermittent renewables on power prices and the incentive to invest in dispatchable power generation in Europe; and (b) the costs of balancing electricity systems with significant reliance on intermittent renewables. I have drawn on material from both of these studies for the analysis presented here.

to comment on the analysis should refer to the Appendices for fuller explanations of how the analysis was carried out and some of the more detailed results.

Finally, let me add a comment on the relationship between empirical analysis and policy evaluation when dealing with the economics of renewable energy. Policies to promote renewable energy are intensely controversial in many rich countries. Part of the reason is that they involve large amounts of (effectively) public money which is spent in ways that some economists believe is relatively inefficient. Another factor is that the development of wind power in some countries has occurred in ways that patently affect existing property rights. It is not surprising that these policies attract criticism. Too many lobbyists and (sadly) official bodies respond to such criticism by adopting the line that everything is or will soon be perfect, with limited or no costs associated with a reliance on renewable energy. In the classic phrase, this is not merely a mistake but a blunder – a deliberate act of self-harm.

Renewable energy – from biomass to windmills to hydro power – has been a significant part of energy use for millennia. Under favourable conditions – for example the wind belt from Texas through the Mid-West, or semi-arid areas of China, or offshore in the shallow coastal waters of NW Europe – modern forms of wind generation may compete in costs with gas generation, especially if the costs of integrating intermittent generation with system demand curves are set aside. However, the policy message that large-scale reliance on renewable power generation is nearly costless is patently not true and discredits those who follow that line. Accepting that large scale decarbonisation of electricity and energy supply is necessary does not mean we can ignore the costs of using renewable technologies.

If governments, official bodies and operators wish to build public trust in policies to promote renewable energy in rich countries, an essential requirement is to move away from reliance on PR and to engage in an honest discussion of the costs and benefits of relying on different forms of renewable generation. This means that the engineering and economic trade-offs required should be acknowledged and factored into any analysis. This study and my earlier work are one step in an attempt to shed some light on those trade-offs.

Eminent economists ranging from John Maynard Keynes to two recent Nobel Prize winners (Al Roth and Esther Duflo) have suggested that economists should aspire to emulate dentists, engineers and plumbers. I embrace the spirit of analysing the performance of renewable generation from the perspective of a practical engineer. Policy should be based on the analysis of real data, not on the empty puffery of lobbyists and political convenience. In this case, the conclusions are the usual mixture of shades of grey.

2. Data on Danish wind turbines

The starting point for the updated analysis is the Danish Energy Agency's register of wind turbines in Denmark. At the end of September 2019, this contained information on more than 9,500 wind turbines, with a generating capacity of at least 6 kW, that have been installed in Denmark over the last 40 years. The information includes data on turbine characteristics, manufacturer and model, location, dates of commissioning and decommissioning, and monthly output since the beginning of 2002.

I have excluded the smallest turbines – with a capacity of less than 500 kW – from most of the analyses, because these are no longer significant for large scale power generation. In 2018 the total capacity of new wind turbines added to the register was 654 MW, of which less than 1 MW was for wind turbines with a generating capacity of less than 500 kW. Two-thirds of the new capacity in 2018 consisted of very large offshore turbines, but even restricting consideration to onshore turbines, those smaller than 500 kW accounted for less than 0.5% of new capacity.

There is no monthly output data for turbines decommissioned before 2002 and there were very few turbines with a capacity of at least 500 kW installed up to 1985. Hence, the sample analysed here consists of turbines with a capacity of at least 500 kW installed in or after 1986 which continued to operate past the beginning of 2002. About 60% of the turbines in the sample (2,600 in total) are onshore turbines with capacity of less than 1 MW. The next biggest categories consist of onshore (18%) and offshore (13%) turbines with a capacity of more than 2 MW. Offshore turbines with a capacity of less than 2 MW have also been excluded. There were a very small number of them installed on an experimental basis in the early 2000s; no one has any interest in using small turbines offshore today.

In terms of the age distribution of operating performance, the number of commercial turbines that operate past the age of 25 years is small. They account for less than 0.1% of the monthly output records. Most turbines, even in Denmark, are decommissioned between the ages of 20 and 24 years. Only 5% of the monthly output records are for turbines aged 20 years or more. As a consequence of changes in policy, over 70% of the turbines in the sample were installed in or before the year 2002. The number of turbines installed between 2004 and 2008 was very small. There has been a pick-up in the number of new installations since 2009 but the rate of new installations is well below the peak period from 1996 to 2002.

The time profile of installations has affected the distribution of turbines by capacity because almost all turbines installed from 2009 onwards have had a capacity of at least 2 MW. Turbines with a capacity of 1 to 2 MW were, in effect, a transitional technology installed between 1998 and 2002. Even during that period, turbines with a capacity of less than 1 MW were still being installed in large numbers.

The relationship between turbines and wind farms is important for the analysis. The basic unit in the register is an individual wind turbine. Wind farms are not formally identified in the data. However, we may define wind farms as clusters of identical turbines (manufacturer, model, capacity, etc) in one local authority with the same start date. There are about 1,270 turbines that operate as single units, a further 770 that operate as a cluster of 2 turbines and a total of 2,280 turbines in clusters of 3 to 39 turbines onshore and 3 to 72 turbines offshore. Practice in reporting output for the turbines which make up wind farms varies. In some cases, turbines are individually metered and the output for each turbine is reported. In other cases, there is a single meter for a group of turbines and the average output per turbine is reported for all turbines in the group or wind farm. There are even changes in reporting practice within the period covered by data for particular wind farms.

It is clear that the sample data cannot be treated as having been generated by about 3,756 fully independent turbines. This is not only because of the averaging of output for some wind farms but also because geographical proximity means that turbines and wind farms may be affected by the same random variables as their neighbours. To allow for this I have taken account of a hierarchy of clustering in carrying out statistical analyses. At the top level, I have used a latitude/longitude grid of 0.5 degrees square: there are 52 grid cells covering Jutland, Funen, Zealand and other islands (including Bornholm in the Baltic Sea) plus Danish coastal waters with either onshore or offshore turbines. Within grid squares the turbines making up wind farms are also treated as separate clusters. In most of the models it is sufficient to cluster the errors using wind farms and there is little additional adjustment to the standard errors from using grid square clustering.

3. Failure analysis for Danish turbines

In this analysis I distinguish between two types of equipment failure:

(a) Major breakdowns. A turbine is classified as suffering a major breakdown if its output in a month is less than 2% of the maximum output that it could produce in that month. The use of a 2% threshold is arbitrary but the results are similar for thresholds of 1%, 2% and 4%. The 2% threshold is less than 10% of the expected output at normal load factors. The chance that such a low level of output would not be associated with a breakdown lasting several weeks is very small. Cases where no output has been recorded for a month are treated as major breakdowns. Operators have had ample opportunity and incentive to correct any mistakes in the register. Note that this definition is, in practice, quite generous. Consider a breakdown lasting for a month from the middle of March to the middle of April. In neither month is output likely to be less than 2% of maximum output, so it will be treated as an equipment failure (see below) rather than a major breakdown. [3]

(b) Equipment failures. A turbine is classified as suffering an equipment failure if its load factor in a month is less than 50% of the average load factor for its peers in the same month. Since the standard error of the mean load factor is less than 5% of the mean in all but one month, the chance that such a low load factor is due to purely random variation is extremely small. This indicator will capture both major breakdowns and shorter periods of equipment failure which reduce output for a significant period of time, the latter typically lasting at least two weeks, and often longer if the failure extends over a period spanning two calendar months.

For each of these measures of failure I have estimated the failure curves for time to first breakdown by broad category of turbine using standard methods for failure (or survival) analysis without smoothing or other assumptions.

3 A recent example from the UK illustrates the difficulty of classifying failures. The new (November 2018) Rampion offshore wind farm experienced a major failure in its transmission system on or about 26th October 2019. It was offline until early December 2019. In the monthly data, October 2019 would simply have a low load factor, while output in November 2019 was zero (treated as a major breakdown) and output in December 2019 may or may not meet the threshold for an equipment failure (depending on other factors). Small differences in timing affect the classification: had the original failure occurred on, say, October 10th and lasted for the same period, then the data would show two months with load factors less than 50% of peer load factors – an equipment failure – rather than a major breakdown, which is clearly what happened.

The failure curves for major breakdowns are shown in Figures 1 & 2. There are two ways of reporting failure curves.[4] One approach, shown in Figure 1, focuses on the time to the first major breakdown for each type of turbine. Any turbine that has experienced a major breakdown is then removed from the analysis for later periods. The motivation for this specification is obvious when considering the use of failure models in many biomedical trials, since failure might be a very serious or irreversible event. In the case of power generation, failed turbines can be repaired, though at considerable expense for offshore turbines, so we may be equally interested in recurrence after an initial breakdown has been rectified. The failure curves by turbine type for all major breakdowns, which includes recurrence after turbines have been repaired, are shown in Figure 2.

Figure 1 – Failure curves for time to first major breakdown of Danish turbines by turbine category

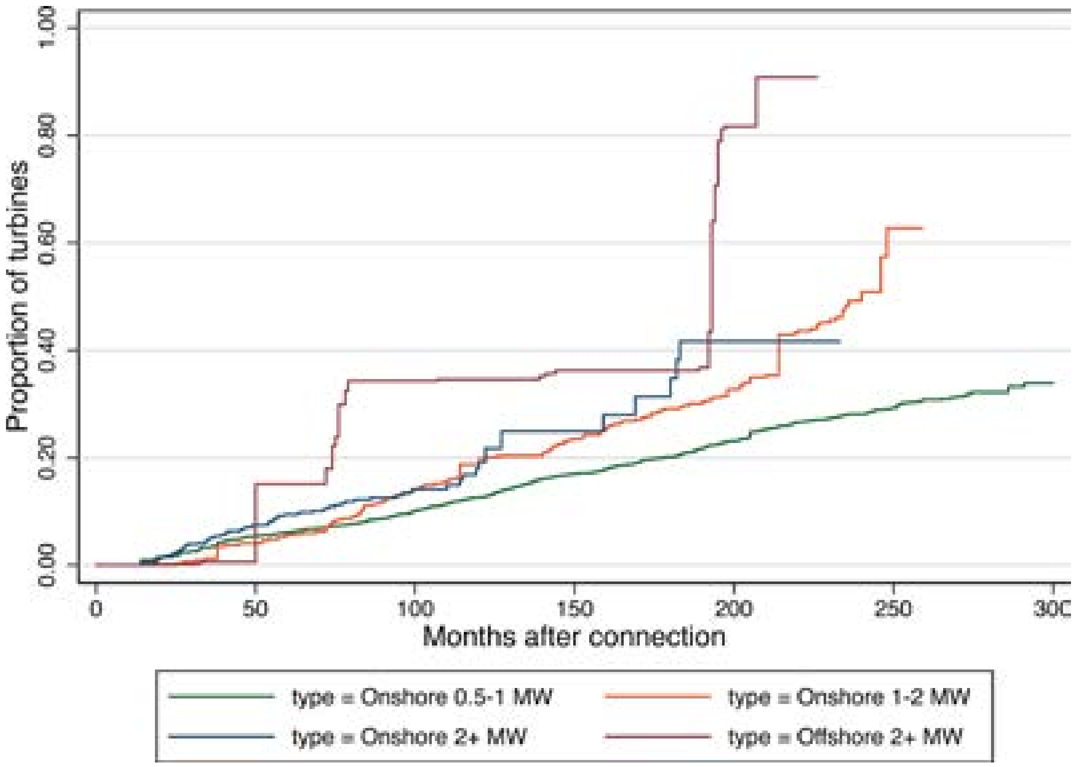

Source: Study estimates

The two figures are quite similar. Among onshore turbines, the medium and larger models – 1-2 MW and 2+ MW – experience major breakdowns at a substantially higher rate than smaller models in the size range 0.5-1 MW. There are indications that the larger 2+ MW turbines suffer higher rates of major breakdown than those in the medium category but the difference is relatively small. However, even more clearly offshore turbines break down at rates that are far higher than their onshore equivalents. Figure 1 shows that within 8 years more than 35% of offshore turbines had experienced their first major breakdown and this proportion rises to more than 80% after 15 years. The difference between large offshore and onshore turbines is somewhat reduced in Figure 2, which

4 It should be noted that the dataset for the failure analysis is unusually large and detailed. It covers 4,305 turbines with 740,000 monthly observations. Even relatively small differences in failure times are statistically significant.

takes account of recurrences, suggesting that offshore turbines may be more reliable after they have been repaired. This is of limited comfort since repairing a major breakdown of an offshore turbine is so expensive.

Figure 2 – Failure curves for all major breakdowns of Danish turbines by turbine category

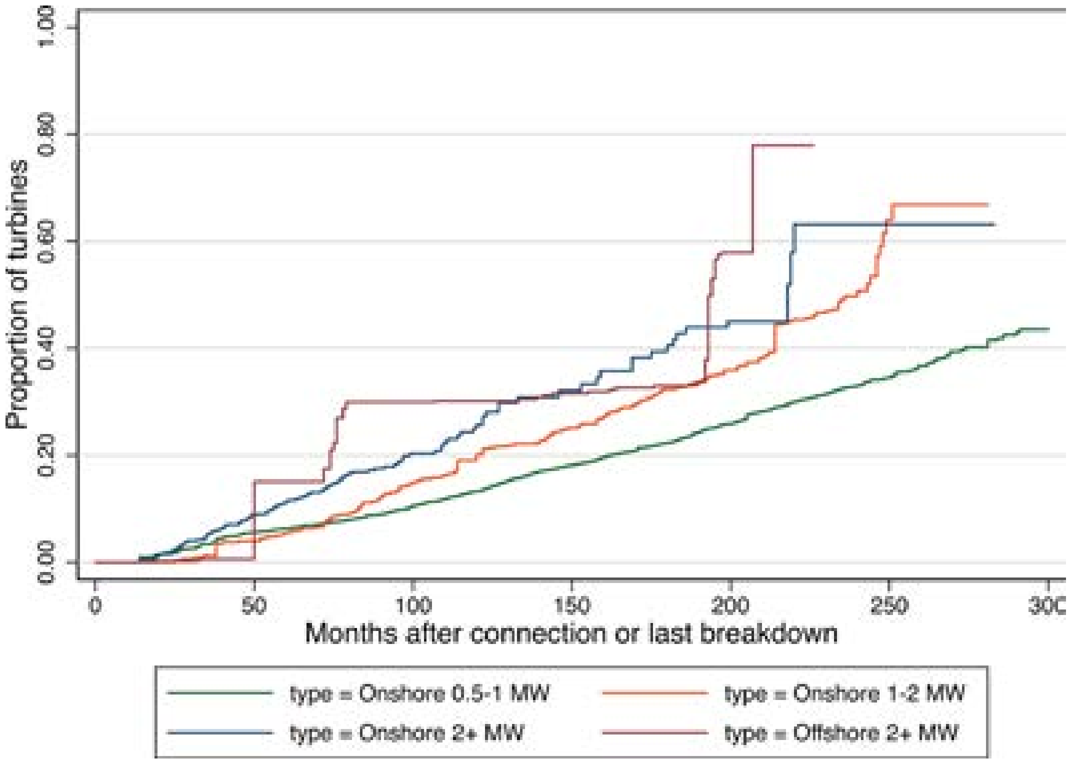

Source: Study estimates

Figures 3 & 4 show the equivalent failure curves for equipment failures by turbine category. Remember that these curves encompass major breakdowns. For the time to the first equipment failure in Figure 3, small onshore turbines have the lowest failure rates and large offshore turbines have the highest failure rates. There is slightly more difference between medium and large onshore turbines. Overall, the central feature is that nearly 60% of large offshore turbines experienced significant equipment failure within the first five years of operation, and almost 80% within their first eight years. Since most turbine manufacturers offer performance guarantees for between 5 and 8 years, the potential liability created by this rate of failure is worrying.

If we focus on all equipment failures – Figure 4 – there is less of a difference between large turbines onshore and offshore. This is because the time gap between onshore and offshore in time to first failure is offset by recurrent failures for onshore turbines. The shorter time between recurrences for onshore turbines may reflect differences in maintenance strategies. There is a large fixed cost in undertaking offshore repairs, so operators may spend more time and money in fixing potential sources of failure with offshore turbines when a failure occurs in order to prevent recurrence. The fixed costs of repair are less of a concern with onshore turbines so that prevention may be less important than getting turbines back into operation as quickly as possible. Whatever the reason, it seems that large (2+ MW) onshore wind turbines are seriously affected by recurrent failures which push up the overall risk of failure after their initial 5 years of operation.

Figure 3 – Failure curves for time to first equipment failure of Danish turbines by turbine category

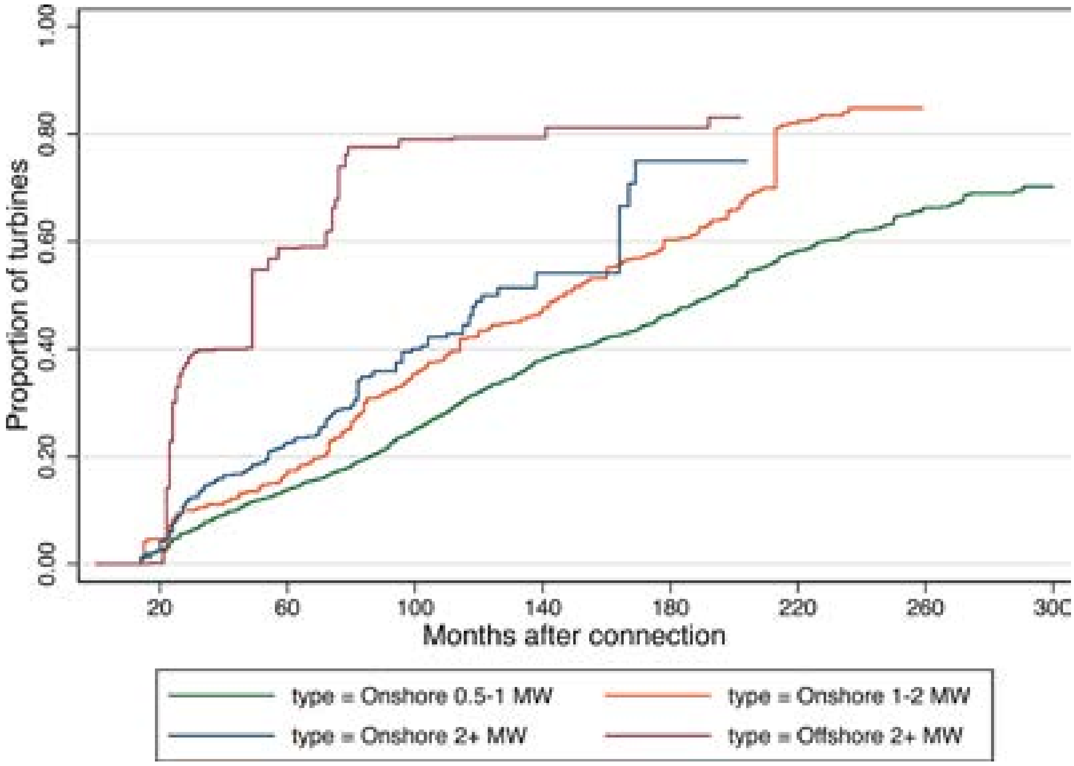

Source: Study estimates

Figure 4 – Failure curves for all equipment failures of Danish turbines by turbine category

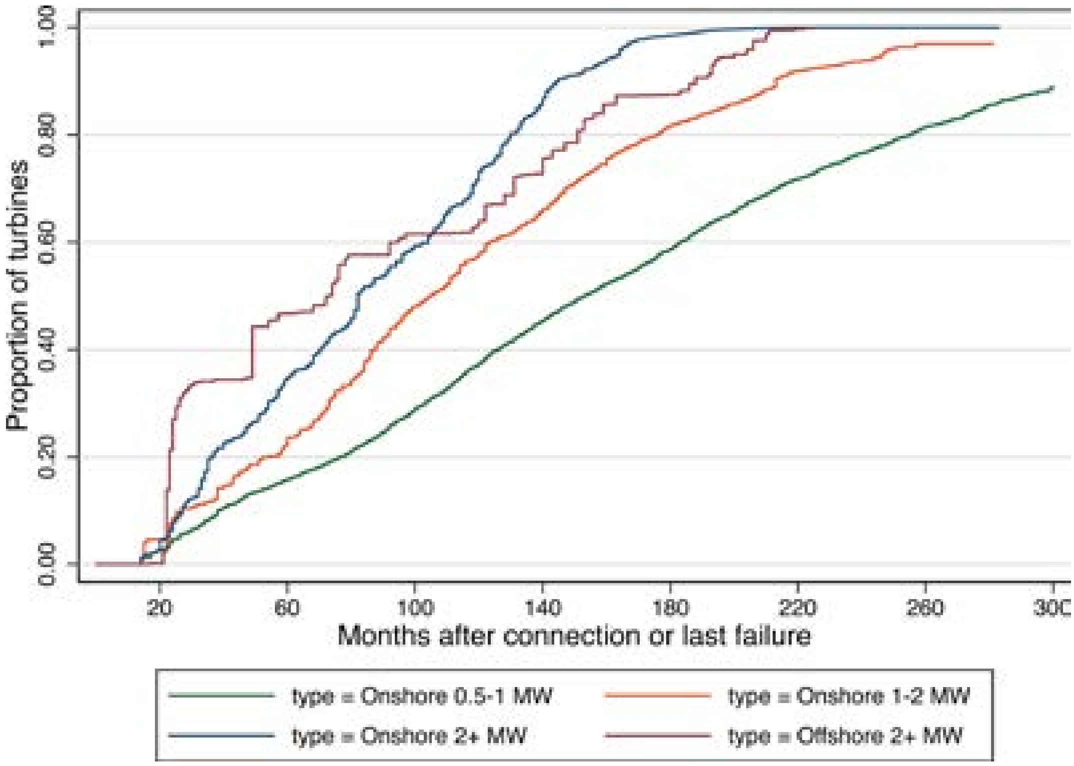

Source: Study estimates

Figure 5 examines the failure rates that underpin the curves in Figure 3 in more detail as it shows a crucial difference between onshore and offshore turbines.[5] The figure plots smoothed values of the hazard rate, i.e. the risk of failure, by time period after connection to the network. The values are smoothed over a window of 12 months to remove erratic monthly variations. The key feature is that the time profiles of the hazards are quite different for onshore and offshore turbines. In the case of onshore turbines the hazard rates are low initially and tend to increase (at varying rates) over time. As would be expected from the cumulative failure curves, the hazard rate for the smallest category of onshore turbines is consistently lower in each time period than that for medium and large turbines.

Figure 5 – Hazard rates for first equipment failures of Danish turbines by turbine category

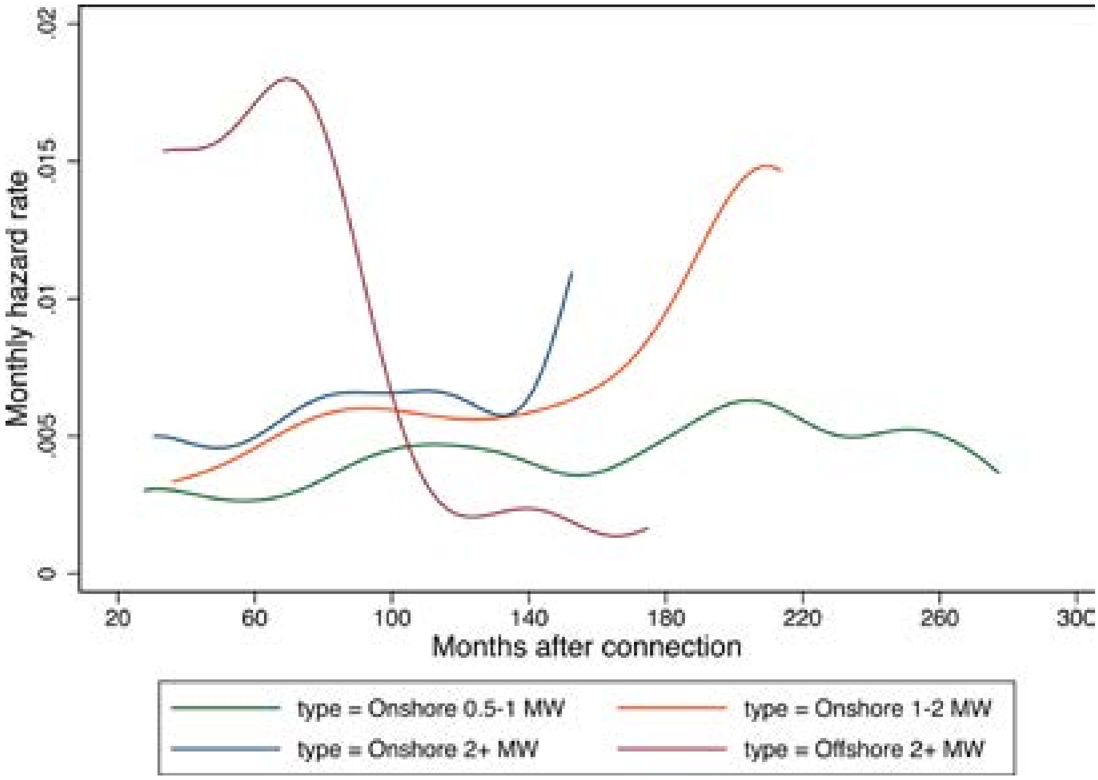

Source: Study estimates

In contrast, the hazard curve for offshore turbines shows high rates of failure in the first 8 years of operation, after which the offshore hazard rates fall below the rates for large onshore turbines. One way of interpreting this pattern is that offshore turbines or wind farms fall into two categories. One category consists of individual turbines or groups of turbines that are either inherently unreliable or particularly vulnerable to the stresses associated with offshore operation that lead to failures. These fail at a high initial rate and after 8 years all – or almost all – of them have experienced their initial failure. The second category consists of turbines that are more reliable or robust and, thus,

5 The calculation of the hazard rates shown in the figure is truncated for the initial months of operation for two reasons. First, the initial 12 months of turbine operation are excluded from the analysis of breakdowns and failures since commissioning may be progressive or involve temporary reductions in output. Second, the smoothing procedure means that a minimum number of observations (equal to the size of the window) is required in order to compute the smoothed hazard rate.

have consistently low rates of failure. The reasons that may cause offshore turbines to fall into the high and low hazard categories may include turbine design and manufacture as well as location and climate/weather conditions.

One clear conclusion from Figure 5 is that it is not reasonable to pool offshore and onshore turbines when examining the factors that affect both breakdown or failure rates and long term performance. On this basis, Cox and Weibull models of proportional hazards have been run for onshore turbines to identify the factors which increase or decrease the failure hazard rates for turbines.[6] The results of this analysis are shown in Table 1. The coefficients are reported as hazard ratios, so a coefficient value of 0.8 means that the hazard rate is reduced by 20% for a unit change in the value of the variable concerned while a coefficient value of 1.20 means that the hazard rate is increased by 20% for a unit change in the value of the variable.

The results confirm that small turbines (< 1 MW) have a much lower hazard rate than the default size category which is large turbines of 2+ MW for both first failure and all failures. The effect of turbine size is stronger for all failures, so that medium turbines (1-2 MW) have a hazard ratio that is significantly less than 1 as compared with large turbines for all failures. The main turbine manufacturers – Vestas, Siemens and NEG Micon – have hazard ratios that are significantly less than 1 when compared with the default category of "Other".[7] The relatively high failure rate of other manufacturers may account, in part, for their disappearance from the market.

The other variable which has a consistent and highly significant effect on both first and all failures is longitude. For each 1 degree further east the hazard rate is increased by about 20%. The impact is biggest for sites close to the Baltic Sea from the eastern coast of Zealand on a longitude about 12° E to Bornholm on a longitude of 15° E. The western coast of Jutland is on a longitude of about 8° E and has the lowest geographical failure rate.

There is some evidence that turbines installed before 2002 have a greater hazard for all failures than turbines installed later but the effect is not strongly significant in what is a very large sample of months at risk. There is no significant effect on the risk of first failure, so any impact must arise from maintenance or operating practices as the turbines age.

It is inevitable that large pieces of mechanical equipment such as wind turbines will fail, sooner or later. The evidence that smaller turbines are more reliable than, especially, large modern turbines presumably reflects a combination of factors, including design decisions and the stresses inherent in building bigger structures and blades to operate more powerful generators. Even so, the failure rates for large turbines are almost 3 times those for turbines of less than 1 MW when considering time to first failure and more than 10 times when considering all recurrent failures. The results

6 The Cox proportional hazards model is semi-parametric as it makes no assumption about the shape of the base hazard curve, but it assumes that the effect of explanatory variables is to shift the hazard curve proportionately up or down. The Weibull model is a parametric model which assumes that the base hazard curve corresponds to a Weibull distribution with parameters that are estimated. One consequence is that the base hazard curve is monotonically increasing or decreasing. The advantage of using the Weibull model is that it can identify weaker influences on hazard rates if its assumptions are warranted.

7 The categories of turbine manufacturers include companies that were absorbed by or are subsidiaries of the named manufacturer – e.g. Siemens includes Bonus. NEG Micon was acquired by Vestas in 2004 but sold a large number of turbines as an independent company before that date. It was originally formed as a merger of the companies selling turbines under the Nordtank and Micon brands and it acquired Wind World before Vestas took over the combined entity.

reinforce the hazard curves shown in Figure 5. The reliability of large turbines relative to small turbines deteriorates over time. As a consequence, the O&M cost for wind farms with large turbines must increase steadily over time as a proportion of potential revenue from power generation. On top of that there is the loss of production during periods when the turbines are unable to operate normally due to failures.

Table 1 – Proportional hazards models for equipment failures for onshore turbines

	Hazard ratio for first failure		*Hazard ratio for all failures*	
	Cox PH	*Weibull PH*	*Cox PH*	*Weibull PH*
Turbine size				
0.5-1 MW	0.338***	0.335***	0.094***	0.094***
	(-4.43)	(-4.36)	(-14.59)	(-14.62)
1-2 MW	0.738	0.738	0.215***	0.215***
	(-1.43)	(-1.39)	(-9.12)	(-9.15)
Manufacturer				
Vestas	0.461***	0.451***	0.326***	0.327***
	(-5.89)	(-6.00)	(-7.64)	(-7.56)
Siemens	0.322***	0.314***	0.238***	0.237***
	(-7.83)	(-7.90)	(-7.73)	(-7.69)
NEG Micon	0.431***	0.421***	0.360***	0.359***
	(-5.70)	(-5.77)	(-6.80)	(-6.77)
Rotor diameter	0.995	0.996	0.993	0.993
	(-1.30)	(-1.03)	(-1.79)	(-1.71)
Latitude	0.953	0.954	0.968	0.968
	(-0.77)	(-0.74)	(-0.64)	(-0.64)
Longitude	1.195***	1.200***	1.220***	1.220***
	(4.32)	(4.23)	(5.57)	(5.55)
Start year				
Before 2002			2.105*	2.004*
			(2.44)	(2.29)
2002–2010			1.320	1.257
			(1.05)	(0.87)
Observations	468421	468421	669698	669698
AIC	30695.4	6818.1	80866.1	7064.9
BIC	30783.8	6928.7	80980.2	7201.9
Log-likelihood	-15339.7	-3399.0	-40423.0	-3520.5
*Notes: (a) t statistics in parentheses; (b) * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$*				

Source: Study estimates

On this last point, it is interesting to examine the number of months affected by breakdowns or failures over the life of each turbine up to the most recent observation or to when it was decommissioned. This may be done by estimating a (generalised) negative binomial model using one observation per turbine with the number of failure or breakdown months as the dependent variable and operating age plus other variables as the explanatory factors.[8] Table 2 shows the estimates for onshore and offshore turbines with coefficients expressed as incidence rate ratios, similar to the hazard ratios in Table 1.

Table 2 – Negative binomial models for the number of months with equipment failures

	Onshore		*Offshore*
	Negative binomial	*Generalised NB*	*Negative binomial*
Turbine size			
0.5-1 MW	0.119***	0.141***	
	(-4.88)	(-4.89)	
1-2 MW	0.265***	0.285***	
	(-4.03)	(-3.82)	
Manufacturer			
Vestas	0.201***	0.204***	0.973
	(-6.62)	(-6.32)	(-0.06)
Siemens	0.203***	0.220***	0.750
	(-6.06)	(-5.74)	(-0.54)
NEG Micon	0.268***	0.261***	
	(-5.24)	(-5.21)	
Rotor diameter	0.983	0.982*	0.934***
	(-1.66)	(-2.15)	(-7.02)
Latitude	0.835**	0.853*	2.594***
	(-2.67)	(-2.51)	(3.52)
Longitude	1.185***	1.184***	1.170***
	(3.53)	(3.55)	(3.78)
Start year			
Before 2002	2.405*	1.850	
	(2.48)	(1.66)	
2002–2010	1.623	1.209	
	(1.61)	(0.57)	
Observations	3756	3756	549
AIC	13555.7	13453.2	1865.2
BIC	13630.5	13552.9	1891.1
Log-likelihood	-6765.9	-6710.6	-926.6

*Notes: (a) t-ratios in parentheses; (b) * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$*

Source: Study estimates

8 The negative binomial model is a generalisation of the Poisson model when the dependent variable is the number of occurrences of an infrequent event. The model allows for fatter and longer tails than the Poisson and statistically it performs much better than the Poisson model for the turbine failure data. The generalised negative binomial model is a further generalisation which allows some of the parameters of the model to be functions of the independent variables.

The results reinforce what has already been reported. The overall incidence of failures for small turbines is between 10% and 15% of that for large turbines, holding other factors constant. This is partly offset by the fact that turbines installed before 2002 have a higher incidence of failures than those installed more recently, though the effect is statistically marginal. Again, turbines in the east of the country have a higher failure incidence than those in the west.

Table 3 translates these equations into estimates of average incident rates for equipment failures and, thus, the average number of months over an operating life that turbine performance will be affected by equipment failures. The estimates are constructed for turbines manufactured by Vestas installed between 2002 and 2010 at notional locations of latitude of 56° N and longitudes of: 8° E, 10° E, and 12° E. Since there are important differences between the estimated failure models for onshore and offshore turbines, the estimates are reported separately for these two groups. Only large offshore turbines have been examined, while estimates for the three size categories of onshore turbines are shown.

For small onshore turbines in the west of the country the average number of months affected by equipment failure is less than 1 month per turbine per 20 years. This rises to nearly 12 months per turbine per 20 years for large turbines in the east of the country – and would be even higher for turbines located on the island of Bornholm at 15° E. For offshore turbines, the average number of months affected by equipment failure varies from 8 per turbine per 20 years in the west to 15 per turbine per 20 years in the east – and again even higher in the Baltic Sea further east.

Table 3 – Incidence of equipment failures over 20 year operating life

	Failure incidence per month			*Expected number of months affected by failures*		
Turbine size	*0.5 – 1 MW*	*1 – 2 MW*	*2+ MW*	*0.5 – 1 MW*	*1 – 2 MW*	*2+ MW*
A. Onshore turbines						
Longitude						
8° E	0.0035	0.0071	0.0248	0.8	1.7	5.9
	(5.26)	(4.04)	(2.43)			
10° E	0.0049	0.0099	0.0347	1.2	2.4	8.3
	(5.94)	(4.27)	(2.57)			
12° E	0.0069	0.0139	0.0487	1.6	3.3	11.7
	(5.10)	(3.88)	(2.57)			
B. Offshore turbines						
Longitude						
8° E			0.0335			8.0
			(8.61)			
10° E			0.0459			11.0
			(8.49)			
12° E			0.0628			15.1
			(5.98)			
Notes: (a) t-ratios in parentheses.						

Source: Study estimates

There are some important lessons that follow from this analysis of turbine breakdowns and equipment failures. The first is a warning about relying upon anecdotal evidence from personal experience or small samples. Turbine failures are infrequent events, especially for the older models of (relatively) small turbines that represent 60% of the turbines in the sample. However, in terms of the contribution of wind power to electricity generation, it is the performance of the 30% of turbines with a size of 2+ MW – both onshore and offshore – that matters. The analysis shows that the experience of small individual turbines is not a good basis for assessing the performance of the forms of wind generation on which policymakers hope to rely in future.

Second, geography appears to be surprisingly important as an influence on failure incidence rates. A difference between onshore and offshore turbines might be expected because of the challenge of designing equipment to cope with the hostile maritime environment. On the other hand, the large difference between failure rates from west to east in a country that is relatively small and has little variation in altitude is more surprising. Since the effect of longitude is both strong and consistent across measures and model specifications, it is unlikely to be due to random factors of location, and the underlying cause merits further investigation. Unfortunately, this is not possible with the existing data.

Third, breakdowns and equipment failures – especially offshore – have high financial and economic costs due to the loss of output and the cost of repairs. The definition of an equipment failure implies a minimum loss of 50% of output for a month and in practice the average loss is much higher. For an offshore turbine in the east of Denmark equipment failures at the rate shown in Table 3 mean an average loss of 5% of gross output. The cost of repairing equipment failures at offshore turbines is high – a minimum cost of £500,000 per incident is plausible because of the costs of mobilising repair vessels. The UK Government assumes opex costs for offshore wind of about £125 per kW of capacity per year at 2018 prices or roughly £1 million per year for an 8 MW turbine. That estimate covers grid connection costs, management and other items, in addition to maintenance and repairs.

Current bidders appear to be gambling that offshore opex costs can be reduced dramatically. On the other hand, the evidence from Denmark translates to a **minimum** cost of £300,000 per turbine per year for dealing with equipment failures with no provision for regular maintenance, etc. It should also be noted that maritime conditions in coastal areas of Denmark or in the Baltic Sea are much less hostile for offshore development than, for example, the Moray Firth in the northern North Sea.

Breakdowns and equipment failures are the consequence of economic choices made by turbine designers, manufacturers and operators. All types of complex machinery experience intermittent failures and engineers pay a lot of attention to the trade-off between cost and reliability. It is possible that there has been a collective judgement that the older generation of small wind turbines was, in economic terms, "too reliable" – i.e. that the trade-off has changed and all parties are aware of that change. Equally, however, anyone who deals with policymakers, lobbyists and financial investors is likely to get the impression that non-engineers and those looking at the industry as outsiders are not aware of such a change. The public is subject to a barrage of claims that cannot be true if there has been a shift towards lower costs and lower reliability.

Policy decisions that rely upon the most optimistic assumptions should be recalibrated to take realistic account of the trade-offs that have to be made. This is particularly clear when considering recent rounds of bids for offshore wind contracts in several European countries. On the basis of public information on capital costs and financial arrangements, the owners and operators of the new projects will suffer huge losses if (a) they receive no more than the contracted prices for their output, and (b) operating costs and performance reflect the experience of offshore turbines over the last decade.[9] If the bidders are not relying upon finding ways of receiving higher than contracted prices, they must be placing a huge bet on large improvements in operating performance and reliability.

Well, maybe. Still, there is little evidence for such an improvement, and the analysis in this section should give rise to serious thought about the exposure of financial investors to the risks of offshore wind.[10] If new wind farms experience anything close to the failure rates observed over the last decade, not only will equity investors lose their stakes but there may be serious write-downs for junior and even senior lenders to the projects. This issue is discussed in Section 8 below.

It is also possible that the primary manufacturers of offshore turbines – MHI-Vestas and Siemens-Gamesa – will take a large hit for liabilities associated with performance guarantees. The initial profits on supplying turbines can easily be overstated by inadequate provision for future liabilities leading to "one-off" write-offs as the true liabilities emerge. For an industry that has been heavily subsidised (indirectly) and is supposedly at the heart of future electricity generation, the manufacture of wind turbines has been remarkably unprofitable over the last 25 years, when proper allowance is made for risk and the occurrence of bankruptcies, mergers and write-offs.

4. Age and turbine performance in Denmark

In this study I have attempted to distinguish between the effects on turbine of (a) ageing, and (b) disembodied technical progress. In this context, disembodied technical progress reflects improvements in performance due to better operating and maintenance (O&M) practices that do not involve the replacement of existing turbines. These may include improvements in the management of stresses in mechanical equipment or blade erosion, as well as exploiting available wind resources better. The role of disembodied technical progress in increasing the productivity of capital equipment is well recognised and is often associated with what is called "learning by doing". It is separate from (embodied) technical progress that is linked to major investment in new capital equipment such as repowering wind farms.

The combination of ageing and technical progress presents a problem for statistical analysis as both are associated with the passage of time. This means that regression models using a pooled set

9 See Hughes (2019).

10 Following up the example of Rampion offshore wind farm given earlier, one of the equity investors in the wind farm is USS – the pension scheme for UK university staff (of which I am a member) – which has a huge deficit of liabilities over assets. Do either the managers of USS or its members have the slightest understanding of the risks of such an investment? Or, are they just following fashion in the expectation that they will be bailed out if things go wrong? Similarly, PensionDanmark is a major investor in both Anholt and Nysted offshore wind farms in Denmark as well Beatrice (GBR), Veja Mate (DEU), Star of the South (AUS) and Vineyard (USA) offshore wind farms in other countries.

of monthly data for turbines (or wind farms) monitored over a number of years can yield unstable estimates of the coefficients on both age and year (for technical progress). The method used in the original study yielded stable estimates of the effect of ageing because the sample was dominated by turbines installed between 1996 and 2002, while the effect of technical progress was combined with the time-fixed effects used to control for variations in wind conditions. However, this approach is no longer feasible for a larger sample of turbines and a longer run of years.

To deal with what is now a very large dataset – over 666,000 observations for more than 3,750 turbines for almost 18 years (January 2002 to September 2019) – I have used additional data and a more elaborate statistical procedure. The additional data consists of average monthly wind speeds at 10 and 50 metres above ground level by grid squares of 0.5° latitude and longitude. This is used to estimate average monthly wind speed at hub height for each turbine, so that monthly variations in wind resources can be included in the model. This new data is used to estimate a cross-section – i.e. across turbines – relationship between age and performance for each year after controlling for wind conditions, so that technical progress shifts the location of the underlying relationship. By pooling the estimates of the slope of the cross-section relationship and allowing for the fact that there are repeated observations on a population of turbines, the procedure calculates a pooled estimate of the effect of age on performance with appropriate standard errors.

The method is discussed in more detail in Appendix B. In fact, it is not strictly necessary to follow this procedure because a standard (OLS) regression should yield consistent, although inefficient, estimates of the coefficients of interest.[11] However, it is better to follow an approach that focuses on the influence of ageing on its own, so that it can be compared with separate estimates of the combined effect of ageing plus technical progress.

The results of using this approach of pooling annual cross-sections are shown in Table 4. As will be seen, there are some important differences arising from the way in which missing and zero observations are treated. The dependent variable is the log of the load factor, which is not defined for months in which production is zero or when no output is recorded (i.e. the value is missing). One approach is to omit such observations, but this introduces a bias because zero or missing output is likely to be a symptom of a turbine failure and is, thus, an indicator of a period of poor performance. On the other hand, to include these observations, it is necessary to replace the missing or zero value for output with some positive but small number. The replacement value chosen can introduce a bias. I have set the minimum load factor used in such cases at 0.1% but there is no "right" way of dealing with such degeneracy and different assumptions yield different results. Still, the results show that the inclusion of zero or missing values is important to get a good picture of the age-performance relationship.

11 This is a piece of regression-speak but it has a technical meaning that is important. Estimators that are "consistent" converge to the true value as the sample size increases, whereas ones that are "unbiased" have a mean that is equal to the true value in small samples. In this study, the sample sizes for each year on its own and for all years together are so large that, for practical purposes, consistent estimates may be treated as being very close to the true value. Efficiency is another statistical term that refers to the expected size of the standard error of the parameter of interest. An "inefficient" estimator is one whose standard error is larger than would result by using some other estimator, generally one making more assumptions about the statistical properties of the data. Again, using an inefficient estimator is not a major concern if the sample size is large enough. In this case the sample covers a large number of turbines for a relatively large number of time periods.

Table 4 – Decline in onshore turbine performance due to ageing

	All turbines	*Turbines by size category*		
		0.5 – 1 MW	*1 – 2 MW*	*2+ MW*
A. Zero and missing observations omitted				
2002 – 19				
Coefficient	-0.007	-0.003	-0.024*	-0.056***
Standard error	0.004	0.005	0.009	0.009
No of observations	666,751	534,515	74,681	57,555
2002–2010				
Coefficient	-0.008	-0.007	-0.022*	-0.087***
Standard error	0.007	0.008	0.010	0.019
No of observations	320,455	275,184	37,537	7,734
2011 – 19				
Coefficient	-0.007**	0.001	-0.026*	-0.024***
Standard error	0.003	0.003	0.010	0.005
No of observations	346,296	259,331	37,144	49,821
B. Zero and missing observations included				
2002 – 19				
Coefficient	-0.010	-0.006	-0.035**	-0.107***
Standard error	0.006	0.006	0.012	0.013
No of observations	669,640	536,636	75,116	57,888
2002–2010				
Coefficient	-0.010	-0.008	-0.039*	-0.184***
Standard error	0.009	0.010	0.018	0.028
No of observations	321,472	275,897	37,727	7,848
2011 – 19				
Coefficient	-0.010**	-0.003	-0.031**	-0.029***
Standard error	0.004	0.004	0.010	0.005
No of observations	348,168	260,739	37,389	50,040
*Notes: (a) * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.*				

Source: Study estimates

A separate issue is whether it is appropriate to pool turbines of all sizes or to estimate separate age-performance curves for each size class. The results in Appendix B reinforce what is shown in Table 4. There are substantial and statistically significant differences between the age-performance curves for small turbines (0.5 – 1 MW) and those for large turbines (2+ MW). The number of observations for large turbines prior to 2011 is relatively small, but even for the early period the hypothesis of the same age-performance curve is strongly rejected. Medium turbines (1 – 2 MW) tend to be more similar to large turbines. Most of them were installed in a brief period from 1998 to 2002 and they were, in effect, an intermediate technology in the transition to large turbines which have accounted for most turbines installed from 2008 onwards.[12]

12 There was a group of early model 2+ MW turbines installed in 2002 which seem to have been notably unreliable and which dominate the age-performance curves estimated for large turbines using data for the years 2002–2010.

Comparing the estimation results across size categories and time periods, it is clear that any decline in the performance of small turbines with age is small and not statistically significant. This mirrors the views expressed by Danish commentators on the earlier paper who focused almost entirely on the large number of wind farms installed up to 2002 using small turbines. On the other hand, there is a clear and strongly significant decline in the performance of large turbines with age. The rate of decline in the performance of large turbines was much higher during the 2002–2010 period than during the 2011–2019 period. It seems likely that this is an indicator of problems with the early versions of 2+ MW turbines installed in the early 2000s that were partly fixed by the time that the later wave of installations occurred from 2008 onwards.

Even so, the rate of decline in performance for large onshore turbines in Denmark from 2011 onwards is 2.4% per year if we omit zero and missing values, or 2.9% per year if those are included. The performance of medium turbines is very similar in the later period with a decline of 3.1% per year when zero and missing values are included. Since large turbines account for about 93% of the monthly observations on turbines installed from 2008 onwards, it seems reasonable to focus on their age-performance relationship as that is the best indicator of the performance of the kind of modern turbines which are relevant today. This evidence reinforces the findings of the failure analysis, suggesting that turbine manufacturers changed the trade-off between reliability/performance and cost when they introduced the new models of first medium and then large turbines. The performance of the older models of small turbines is not a good guide to how the more modern generations of turbines have been performing.

The results in Table 4 are based on an analysis in which each turbine is treated as a separate observation after allowing for the correlation of performance within wind farms. Table 5 provides a more direct comparison with the earlier study by treating wind farms as the unit of observation. One of the changes that occurred after the hiatus in installations from 2003 was a shift away from single or small groups of two or three turbines, with developers preferring larger clusters. New projects developed since 2008 have been dominated by large turbines, the shift is likely to be a response to much greater capital investment per turbine and, perhaps, the greater technical demands of managing large turbines.

The estimates of age decline in Table 5 are slightly larger than those for Table 4. Aggregation to wind farms reinforces the difference between (a) small turbines of less than 1 MW, and (b) medium and large turbines of greater than 1 MW. For both medium and large turbines the rate of age decline is much higher in the earlier period 2002–2010 than from 2011 onwards. The difference is especially marked for large turbines and the conclusion that large turbines experienced an initial but extended period of poor reliability and performance seems inescapable. There is a crucial lesson here that the introduction of a new generation of plant or machinery is often accompanied by a prolonged period during which the new units suffer what might be called “teething problems” that will eventually be sorted out but which are likely to cost manufacturers and investors large amounts of money and angst.

Even after this initial period, medium and large turbines simply do not perform as well over time as the earlier generation of small turbines. Across a range of models the rate of decline in performance of onshore turbines with age is about 3% per year. This means that after 15 years, a wind farm can be expected to produce only 63% of its output in the first full year under identical wind and weather conditions. At the same time, it is widely reported that O&M costs increase substantially with age so that the net revenue earned by a wind farm at age 16 will be much lower than at age 1. Depending on the length of any original subsidy arrangement or power purchase agreement, there will be a strong economic incentive to scrap turbines when they reach the age of 16+ years and to repower the wind farm site.

Table 5 – Decline in onshore wind farm performance due to ageing

(Missing and zero observations included)

	All turbines	*Turbines by size category*		
		0.5 – 1 MW	*1 – 2 MW*	*2+ MW*
2002 – 19				
Coefficient	-0.011**	-0.006	-0.047***	-0.118***
Standard error	0.003	0.004	0.013	0.013
No of observations	288,799	243,572	27,337	17,890
2002–2010				
Coefficient	-0.011*	-0.006	-0.065**	-0.205***
Standard error	0.005	0.005	0.022	0.027
No of observations	141,354	124,076	13,671	3,607
2011 – 19				
Coefficient	-0.011***	-0.006	-0.029***	-0.031***
Standard error	0.003	0.003	0.008	0.006
No of observations	147,445	119,496	13,666	14,283
*Notes: (a) * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.*				

Source: Study estimates

5. The performance of offshore turbines

The results reported in the earlier study suggested that the performance of offshore wind turbines in Denmark was very poor. However, the sample size was small and overall experience with offshore wind was very limited in 2012. A major element in the poor performance was a major failure of the electrical system at the Nysted offshore wind farm from July to October 2007. It was suggested that this was a very unusual occurrence. Unfortunately, that is not the case. The Horns Rev 1 wind farm, with 80 turbines, experienced a major outage from November 2018 to February 2019, while the Horns Rev 2 wind farm experienced a shorter outage from October to December 2015. In the United Kingdom such failures are also not unusual. Mostly recently the Rampion wind farm off the south coast of England was out of operation from late October to December 2019 as a result of a failure in its electrical system.

There is a systematic pattern here that needs to be understood. Even if offshore turbines operate well, there is little benefit if the associated subsea transmission and electrical systems are prone to fail. And such systems are certainly prone to major failures. Subsea transmission lines are notorious for the severity and length of their outages – the endless troubles that have affected the Moyle interconnector from Ireland to the UK, or the more recent West Coast interconnector from Scotland to England, are illustrations of the large number of problems with subsea transmission lines around the world over the last two decades.

Table 6 shows the estimated coefficients for the performance of offshore wind turbines. The number of wind farms (26) is too small to yield reliable results while the sample of turbines for 2002–2010 is dominated by two wind farms (Nysted and Horns Rev 1) which have experienced collective failures. Hence, I will focus on the coefficients for 2011–2019 using either OLS or pooled estimation. The rate of age decline is estimated to be 4.5% per year with a high level of statistical significance and no evidence of significant technical progress for offshore turbines.

Table 6 – OLS and pooled models of offshore turbine performance

(Missing and zero observations included)

	OLS 2002–2019		*OLS 2011–2019*		*Pooled 2011–2019*
	Model 1	*Model 2*	*Model 1*	*Model 2*	
Turbine age	-0.015	-0.063**	-0.045*	-0.046**	-0.045***
	(0.013)	(0.018)	(0.021)	(0.013)	(0.016)
Year		0.054**		0.001	
		(0.017)		(0.026)	
Ln(Wind speed)	1.566***	1.568***	1.296***	1.296***	1.256***
	(0.316)	(0.291)	(0.249)	(0.245)	(0.264)
% time WS ≤ 4 m/s	0.677	0.766	0.675	0.676	0.463
	(0.610)	(0.579)	(0.873)	(0.846)	(0.279)
Observations	70,402	70,402	49,585	49,585	49,585
No of turbines	548	548	548	548	548
No of wind farms	26	26	26	26	26
R-square	0.130	0.144	0.139	0.139	
Root MSE	0.720	0.714	0.641	0.641	
*Notes: (a) Standard errors in parentheses; (b) * p < 0.05, ** p < 0.01, *** p < 0.001.*					

Source: Study estimates

The greater height and size of modern offshore turbines means that the load factor for offshore turbines is less sensitive to average wind speed than that for onshore turbines – the elasticity is about 1.3 for offshore turbines but 1.9 for large onshore turbines and 2.4 for small onshore turbines. This should mean that the power generation from offshore wind farms should be less variable and easier to predict than from onshore wind farms, which has clear advantages for the management of electricity systems. On the other hand, a decline in performance of 4.5% per year is very severe and means that the expected output from a turbine aged 16 years will be only 50% of the expected output for the same turbine at an age of 1 year under the same weather conditions.

It is important to stress that the decline in performance is not simply a matter of an operating turbine producing a lower amount of output at identical wind speeds and other weather conditions. Much of the decline is a consequence of more frequent equipment breakdowns and other operational outages. Still, the practical consequence is the same, a reduction in the expected revenue from generation as the turbine ages.

There are two dimensions to the results on the performance of offshore wind turbines reported here that are important for both policy and the operation of electricity systems. The first is that offshore wind farms have a probability of experiencing a major outage lasting anything from 1 month upwards of at least 10% for each year of operation. That has major consequences for the calculation of the probability of the loss of load for an electricity system, because it means that the amount of reserve capacity required to maintain system supplies has to be much higher than was the case for systems with large amounts of dispersed generation. In Denmark that reserve capacity is provided by access to the Scandinavian and German grids so that the loss of either the Nysted or the Horns Rev wind farms do not jeopardise the stability of the grid. However, the situation may be different once the new Kriegers Flak wind farm (about 605 MW) becomes operational. In the UK, recent offshore wind farms are still bigger so they should be treated in security of supply assessments as being equivalent to but rather less reliable than nuclear plants or (in the past) large coal-fired plants. Losing either of the Hornsea wind farms or Triton Knoll or Moray East in the middle of winter could have a large impact on system stability.

The second dimension concerns the risks to lenders and investors implied by the decline in performance due to ageing. Not only is there a loss of generating revenue but the more frequent breakdowns require higher expenditures on repairs and preventative maintenance. The consequences of lower revenues and higher O&M costs are examined in the next section.

It is important to assess the results from the analysis against the context of alternative forms of power generation. Problems with the performance or reliability of early units of a new generation of machinery or plant are a normal feature of the industry. Power equipment manufacturers such as Alstom, General Electric and Siemens have run into major difficulties in meeting promised performance levels for each new generation of gas turbines. Such episodes often result in the manufacturer incurring large liabilities because of performance guarantees and repair costs, but they would expect to fix any design problems and to achieve much higher levels of reliability after the initial problems have been sorted out. The substantial decline in performance with age observed for offshore and large onshore turbines in the period from 2002 to 2010 is consistent with an initial period of poor reliability for a new generation of turbines.

However, for conventional power generation equipment, one would not expect to observe a continuous decline in performance at a rate from 3% to 4.5% over an extended period of a decade or more. Such a decline must also be associated with higher O&M costs and a substantial reduction in the expected economic life of the equipment. In effect it means that the equipment is what a valuer would describe as a "wasting asset". Physically it might last for 25 or 30 years but the net yield will fall to the point that expected revenues are little more than expected O&M costs. Technically, all plant and machinery is classified as wasting assets for valuation purposes, but there is a crucial difference between assets, such as gas turbines, that can operate quite reliably over an expected lifetime and those whose net yield declines at a significant rate.

There is a further lesson that can be drawn from this analysis. The initial period of poor reliability and performance following the introduction of 2+ MW turbines lasted for 8 to 10 years. The onshore wind industry then standardised on turbines of this period with a capacity of up to 3.5 MW. However, the offshore industry has introduced a new generation of turbines with a capacity of 6 to 10 MW. In Denmark the first turbines of this new generation started to operate offshore in 2018. If the early experience with the previous generation of onshore turbines is a reasonable guide, there will be an extended period of poor reliability and performance lasting to the mid-2020s, or perhaps longer. Part of the cost of this learning will certainly fall on the turbine manufacturers. Some of them may not survive as the financial record of turbine manufacturers is poor and often they do not have the reserves to cover their liabilities under performance guarantees.

At the same time, turbine manufacturers have made claims about moving to yet another generation of offshore turbines of 15-20 MW by 2030 – hoping to bring down costs by exploiting economies of scale. Setting aside the usual technical hype, the real question will be whether manufacturers have the resources required – not only to fund the initial development of the turbines but also to bear the inevitable losses that they will incur during the initial period of poor performance and reliability.

The ratio of puffery to solid technical reporting is unusually high in the wind generation business. Current claims and expectation are typical of boom periods when (a) technical advances are over-stated, and (b) there is a large pool of capital looking for projects to invest in, even on terms that offer very low returns after adjusting for the risks being taken. Such are the consequences of the extremely lax monetary and financial policies being followed by governments and central banks. The current situation in the wind industry is no different from a conventional property bubble with all the usual features of overselling and the exploitation of naïve investors. The eventual bursting of the bubble with all of the usual pain, is equally inevitable. Unfortunately, the failure to learn from past bubbles across a range of countries and businesses is an endemic feature of periods when interest rates are low and PR is rated more highly than technical competence.

6. Auctions and the winner's curse

The results of auctions for contracts to purchase power from new offshore wind farms produce a high level of (often) overheated rhetoric and speculation about the underlying economics of offshore wind generation. In the UK the auctions cover the strike price for Contracts for Differences (CfDs), while in other countries the auctions are used to set the offtake price for Power Purchase Agreements (PPAs). In each case the arrangement is that the auction effectively determines a guaranteed price for either a fixed volume of sales (period uncertain) or a fixed period of time (volume uncertain). There are detailed differences between the contracts which may be important in determining the incentive for early termination if the market power price exceeds the contract price.

Commentators with limited knowledge or interest in the economics of auctions tend to make claims about trends in the cost of offshore wind generation based on the results of small numbers of auctions with significant differences in contract conditions and operating circumstances. Those with a somewhat broader knowledge are likely to conclude that the winner's curse is thriving. To explain what this is, and how it might apply to offshore wind, it is useful to consider the lessons of the 3G spectrum auctions between 2000 and 2002.

At the end of the 20th century, 3G (UMTS) was seen as the route to a large increase in revenues from mobile telephony, particularly via the provision of high speed data services. Between 1999 and 2002, some 32 countries awarded licenses for the use of radio spectrum to operate 3G services to a mixture of incumbent operators and new entrants. The number of licenses varied from country to country. In some countries, such as the UK, there was a requirement that at least one license had to be awarded to a new entrant. In 19 countries licences were awarded by auction, while administrative mechanisms – beauty contests and other arrangements – were used in 13 countries where the cost of licenses was determined by a pre-set formula.

In two countries – Germany and the UK – which held auctions early in the process, the auction led to spectrum prices that were much higher than anticipated, with the result that very large sums of revenue were raised for the governments – €50.8 billion in Germany, €36.9 billion in the UK. These sums were equivalent to about €620 per person. Smaller but still significant sums were generated by auctions in 2000 in the Netherlands, Italy and Austria. However, by the end of 2000, the original enthusiasm for – and the money raised by – the auctions had waned, so that auctions in 2001 and after typically yielded payments of €50-60 per person. In two countries – South Korea and Poland – beauty contests with licenses awarded to incumbents raised similar sums per person, but elsewhere administrative procedures set much lower prices for spectrum.

While many factors – such as license conditions, amount of competition in the telecoms market, national income and expenditure levels, and size of the market – affected the outcomes, it is reasonable to conclude that €50-60 per person was a fair benchmark for the value of a 3G spectrum license in large OECD countries in 2000. The benchmark value was lower in small and medium countries where the costs of developing 3G networks were high relative to the revenue that might be expected. Overall, it was clear – and became even more obvious subsequently – that incumbent operators in the first wave of auctions had vastly overpaid for their spectrum licenses. This was the winner's curse.

More formally, the winner's curse is a vivid term used to describe the likelihood that the winning bidder in certain types of auction will have overpaid for the asset being sold. Consider an auction for an asset whose value is unknown but where information relevant to assessing that value is public and the ex-post value of the asset is the same for each bidder – these are known as common value auctions. Each bidder has to estimate a value from a distribution that depends on the public information available. The highest bidder wins and thus the winner is likely to be the one whose bid is furthest towards the "optimistic" end of the distribution, i.e. the one who makes the highest payment relative to the ex-post value of the asset. In mathematical terms, the net expected value of the asset conditional on winning the auction is negative and it is more negative the larger the number of bidders.[13] While the assumption of a common value for the asset may not be strictly true, the intuition that the winning bidder in an auction is likely to have over-paid in general unless bidders are experienced in auction bidding and take steps to offset the bias towards over-payment.

There were also other factors that affected the early 3G spectrum auctions. They occurred towards the end of the Dotcom bubble. Valuations of internet and telecom assets were very high and the

13 The amount of the over-payment is a function of the n-th order statistic of the distribution of the value of the asset, where n is the number of bidders, which increases as the value of n increases.

cost of capital for investment in 3G spectrum was low, reflecting a belief that demand for telecom data services would grow very rapidly in the next decade. Incumbents were extremely concerned either to exclude new entrants or to minimise their impact by increasing the cost of the spectrum that they had to acquire. They were also flush with cash from the profits being generated by their existing 2G businesses.

After participating in the initial auctions, the key bidders had run down their cash reserves and had learned how to take account of the winner's curse in preparing their bids in later auctions. In later auctions the sellers of spectrum, often national government, put more weight on other objectives (such as accelerating the rollout of 3G services) rather than maximising the revenue from selling spectrum.

An important question is whether over-payment for auctioned assets matters, except to the two parties to the transaction. The seller – the government in the case of 3G spectrum auctions – receives a higher price than it might have expected at the expense of the owners of the winning bidders. As pure redistribution, there might seem to be few consequences, and it can be argued that the outcome of the auction should not affect the behaviour of the winning bidders as they would want to make best use of the assets that they had acquired.

In practice, that is not what happened in the case of 3G auctions. Governments assumed that other spectrum was more valuable than they had previously thought and, as a consequence, changed their fiscal plans and the management of their spectrum assets. The winning bidders found that they had depleted their capital and lacked the financial resources to develop 3G services as quickly as had been expected. Since all of the bidders had to overpay in order to obtain spectrum, competition to roll-out 3G services as rapidly as possible was reduced and some of the winners decided that a strategy of a more gradual roll out combined with premium prices for 3G services might offer a better return on their investment in acquiring exclusive spectrum rights. Finally, the behaviour of bidders at later auctions for other segments of spectrum has clearly been affected by the history of extreme over-payment in some of the 3G auctions. New competitors have been reluctant to enter markets with well-established incumbents, so that the perception that the winner's curse is the price of new entry can be a significant barrier to entry.

The winner's curse is a particular manifestation of a phenomenon which affects large, one-off, projects in many sectors – transport, energy, information and communications, construction, etc. The euphemistic term is "optimism bias", but a more down-to-earth description is "delusion and deception". The core problem is that when such projects are assessed at the outset, their benefits tend to be substantially over-estimated and their costs are equally heavily under-stated. There is a growing literature which documents the problem by studying samples of completed projects and identifying the factors associated with more or less severe errors in *ex-ante* evaluations – see, for example, Flyvbjerg et al (2002), Flyvbjerg (2009), and Ansar et al (2014).

It is easier to document mistakes in *ex-ante* cost estimates, because public bodies usually have to report the *ex-post* cost of large projects. However, in many cases it is the effect of optimism bias on estimates of the benefits of large projects that is the more serious source of misallocation of resources. Since net benefits – revenues and operating costs – arise over an extended period into the future, it is easy for project developers to claim that errors are the consequence of bad luck that could not have been foreseen. While an element of bad luck cannot be ruled out, most risks could and should have been taken into account at the outset as a way of correcting optimism bias.

In the next section I examine the effects of optimism bias on estimates of future revenues and costs as manifested in auction bids for offshore wind projects. This is important because such auction bids are being interpreted as reliable indicators of the cost of offshore wind power. If, on the other hand, allowance is made for optimism bias on the basis that it gives rise to the winner's curse, the financial underpinnings of such auction bids must be examined in order to construct more reliable cost estimates. As an example of this process, I have constructed a detailed financial analysis of the largest offshore wind project in Denmark – Kriegers Flak – using the public information that was available at the time of the auction for the right to develop the project.

The analysis has two objectives:

- The first goal is to understand what the "true" cost of offshore wind from a large project of this kind might be if we strip out the over-bidding associated with the winner's curse.

- The second goal is to examine the impact of turbine breakdowns, and the associated decline in their performance over time, on the cost of offshore wind.

Both questions take us back to the issue of whether the winner's curse matters. If analysts are content to treat the results of offshore wind auctions as random draws from a very wide distribution, there would be no reason to assign any particular weight to the outcome of the Kriegers Flak auction or a small number of comparable auctions elsewhere in NW Europe. However, it is claimed that there is a strong downward trend in the real cost of generating electricity from offshore wind, i.e. that the mean or median of the distribution is falling rapidly. Such claims would, of course, be questionable on statistical grounds even if the auctions were not affected by over-bidding. Hence, examining the winner's curse provides a way of identifying the factors that must be taken into account when making inferences about the base cost of offshore generation.

While it is unwise to pay too much attention to auction outcomes, there are good reasons to believe that the cost of offshore generation has fallen, though prospects for the future are rather more uncertain. First, the overall capital cost of offshore projects per MW of capacity has been falling since the middle of the past decade after adjusting for the depth of water and distance from shore. This is linked to the economies of scale associated with the introduction of a new generation of larger (6 – 10 MW) turbines. This trend has to be set against the need to develop sites in still deeper water and further from shore. Second, as a highly capital-intensive technology, offshore wind has benefitted from the reduction in the cost of capital due to an extended period of low interest rates. The perceived risks of investment in offshore wind have also fallen, so the risk premiums are lower, which has reduced the cost of capital even further.

Neither of these factors guarantees a continued reduction in offshore wind costs. As the cost of turbines falls as a share of total capital costs, the gain from adopting larger turbines decline and the risks increase. Instead, it will be the cost of civil works and offshore structures that will be critical. These are mature technologies and there is little reason to expect rapid reductions in future. The real cost of capital in rich countries is well below historic average levels. This is reflected in bubbles in asset prices across the investment spectrum. Assuming that this indicates a "new normal", let alone that it presages further reductions in the real cost of capital, is the kind of classic error that occurs late in all speculative episodes.

7. Kriegers Flak and the economics of offshore wind generation

Preliminary preparation of the Kriegers Flak project was carried out by the Danish authorities. The project covers the construction and operation of a 605 MW wind farm in the Baltic Sea. In parallel, two transmission companies – Energinet and 50Hertz – are building a new transmission line that will act as part of an interconnector from Denmark to Germany via the new wind farm and two existing wind farms in the German sector of the Baltic Sea (EnBW Baltic 1 & 2). Because the grids in Denmark and Germany do not have the same synchronisation, the transmission line involves a back-to-back HVDC converter and other technical constraints. As a consequence, responsibility for and the costs of transmission of power from the wind farm fall to Energinet. The auction bids were for the offtake price at the wind farm under a power purchase agreement for a fixed volume of 30 TWh that was expected to be reached in about 11 years. After that the project operator would expect to sell the power at the power market price for the DK2 region less the projected transmission cost. The winning bid was submitted by Vattenfall with a fixed nominal price of €49.9 per MWh. Construction of the wind farm started in 2019 and is due for completion by the end of 2021. Vattenfall has contracted with Siemens-Gamesa to supply 72 SG-8.0-167 8.4 MW turbines with a hub height of 107m.

Details of the financial model of the Kriegers Flak project are given in Appendix C. The headline figures that are quoted for the capital cost seem relatively low. However, once the cost of the transmission line is included and allowance is made for the relatively shallow water (maximum depth 25m) and short distance to land (15 km), the capital cost of the project is very similar to the projected cost for offshore wind farms in Europe – see Hughes, Aris & Constable (2017).

I have used the NASA data to calculate the distribution of wind speeds at hub height at the Kriegers Flak site and combined this with statistical data on the performance of large offshore turbines similar to the Siemens-Gamesa model. The predicted average load factor in the early years of operation is only 42.8%, without making any allowance for a decline in performance due to ageing.[14] This is well below the reported assumption of an average 50% that underpins the expected duration of the power purchase contract. Clearly, Vattenfall and its advisers have convinced themselves that the wind farm will outperform what might be expected on the basis of the data relating to the project that is available. As we shall see, there are other parameters affecting the financial outcome of the project for which Vattenfall must have made rather optimistic assumptions to justify its bid.

The starting point of the analysis is a standard financial model using certainty-equivalent values for the key parameters.[15] This was used to identify the key sources of uncertainty that affect the

14 This estimate is based on a Monte Carlo simulation of (a) the distribution of monthly wind speeds at hub height based on historical data for the site, and (b) an econometric model of the monthly load factor as a function of month and average monthly wind speed at hub height for large offshore turbines with a capacity of at least 7 MW. The econometric model fits the data very well with an $R^2 > 90\%$.

15 A certainty-equivalent value compresses the distribution of potential values of a parameter to a single value that yields the expected value of the cash flow of the project adjusted by the risk premium. In the special case of risk-neutrality (no risk premium) and a symmetric distribution, the certainty-equivalent value is equal to the mean

value of the project. I have not examined the risk of costs overruns in constructing the project because this is easily understood and largely within the control of the developer. More important are the risks concerning three parameters that cannot easily be controlled:

- The average load factor without any allowance for performance decline, as discussed above.
- The levels of variable O&M costs, which is linked to the performance of the turbines.
- The path of real offtake prices, calculated as the NordPool market price for region DK2 minus the transmission cost, in the period after the end of the power purchase contract.

The base version of the financial model assumes an average load factor of 50%, which is generous given the statistical analysis, plus a level of O&M costs based on estimates for the early 2020s reported by UK and US official agencies. It is also assumed that there is no trend in the real NordPool power price relative to its average level from 2011 to 2019. On these assumptions the project has a negative present value of -€387 million, given the auction power price of €49.9 per MWh. In plain terms, equity investors will lose almost all their money, and this is on a very optimistic assumption about the load factor over the life of the project.

It is possible that Vattenfall believes that it can turn an underlying project loss into a profit by some kind of financial engineering. The financial structure that I have modelled is quite simple: a debt to equity split of 80-20 with a grace period for debt repayment of 2 years, which is rather generous to the project in view of its low return and high risks. On the other hand, more complex financial structures have a tendency to fail dramatically if the assumptions on which they are based are not realised. In practice, it seems more likely that Vattenfall is relying on being a state-owned company with a large balance sheet to override the risks that would sink this investment when viewed as a standalone project.

Still, even companies with large state shareholdings, such as Vattenfall, Orsted, and EdF, usually want to improve their balance sheets by bringing in long term financial investors to refinance major projects once they have started operation. Part of the reason for the timing is that the risks of construction and initial operation are past. Potential financial partners will see an asset that will earn revenues for a period from a guaranteed power purchase contract plus an option on future power prices for the remaining years of the project life. On this basis, they may be willing to accept a lower rate of return on their investment than would have been expected by an equity investor at the outset. The expected rate of return on such a buy-in might be as low as 5-6% in real terms, since the risks are far from negligible.

Even so, Vattenfall would face a large write-down were it to sell a substantial share in the Kriegers Flak project after two years of operation on the base assumptions. The reason is that the cash flow in the first 10 years of the investment is fully committed to debt repayments and the lenders would not wish to extend the debt period beyond the expiry of the initial power purchase contract. The value of the equity in the project would have to be written down by nearly 80% in order to sell a share in the project to investors with a hurdle rate of return of 6%.

of the distribution. For skewed distributions with no risk premium, the median value of the distribution may be a better measure of the certainty-equivalent value.

I have examined what operating assumptions Vattenfall and potential investors would have to make for the project to just earn a positive present value over 25 years. The best scenario seems to rest on a combination of: (a) an average load factor of 52% over the life of the project (vs 50% in the baseline); (b) operating costs that are 20% lower than in the baseline; and (c) a market power price that increases in real terms at an average 4% per year over 25 years, so that the offtake price (the market price minus transmission costs) at the end of the power purchase contract is nearly €49 per MWh at 2018 prices and increases to €85 per MWh in year 25. Under these assumptions the net present value of the project is €17 million, which is still a very low return for a project that relies upon very risky assumptions.

However, when the details are examined, the net present value gives a somewhat misleading view of the project from the perspective of potential financial investors. The cash flows are fully committed to debt repayments for 10 years and the value of the equity share rests entirely on the prospects for the market price and, thus, free cash flow more than 12 years into the future. In practice, therefore, Vattenfall will have to arrange a swap of equity for debt in the project in order to persuade buy-in investors to take a stake in the project. By allocating nearly 75% of a higher buy-in price to reduce debt to reach a 50:50 debt-equity ratio, Vattenfall would be able to offer buy-in investors a modest return on their investment in the early years combined with the prospect of a much higher return in later years.[16] By reducing the financial leverage of the project, Vattenfall could increase the net present value of its stake to about €47 million, which while still low would be accompanied by significantly less risk.

Having identified the key parameters that affect the value of the project the next step in understanding the impact of the winner's curse is to capture the uncertainty that gives rise to the distribution of project values. It is essential to differentiate between two sources of uncertainty:

(a) Structural uncertainty reflecting the inherent variability of, for example, monthly wind conditions and market prices. This extends to random equipment failures which will determine the level of O&M expenditures.

(b) Forecast uncertainty concerning the key parameters discussed above. No-one knows what the future growth in real market prices will be but it is an essential element in valuing the project. Relying on a single certainty-equivalent value – or standard scenario analysis – is very crude because they give no insight into the distribution of potential project values.

The distribution of the value of the Kriegers Flak project is quite concentrated when structural uncertainty alone is incorporated in the breakeven scenario – see Figure 6. The standard deviation of the net present value is €17.3 million for Monte Carlo simulations of a variant whose mean present value matches that of the certainty-equivalent model discussed initially. The range from the 5th to the 95th percentiles of the net present value is from €4 million to €60 million.

There is one important lesson from the analysis of structural uncertainty which relates to the interpretation of observed load factors. Over the full set of Monte Carlo simulations the 5th

16 While the press release would put a different gloss on the matter, the combination of an equity-for-debt swap and an investor buy-in would be equivalent to an equity write-off as it would reduce the cash flow drain on the company while also reducing its expected net earnings from the project in future years.

percentile of the average load factor for the first full year of operation is 46.9%, while the 95th percentile is 57.4%, when the mean load factor is calibrated to 52%. Now, suppose that a potential investor is told that the expected load factor for a wind farm is 52% averaged over a year. How could the investor test that claim, given the known year-to-year variability in wind conditions and other factors affecting total output? In statistics, this is formulated as a calculation to estimate the size of the sample – the number of annual observations – required to identify the difference between the sample mean and a known value.

Figure 6 – Distribution of Kriegers Flak NPV: structural risk for the breakeven scenario

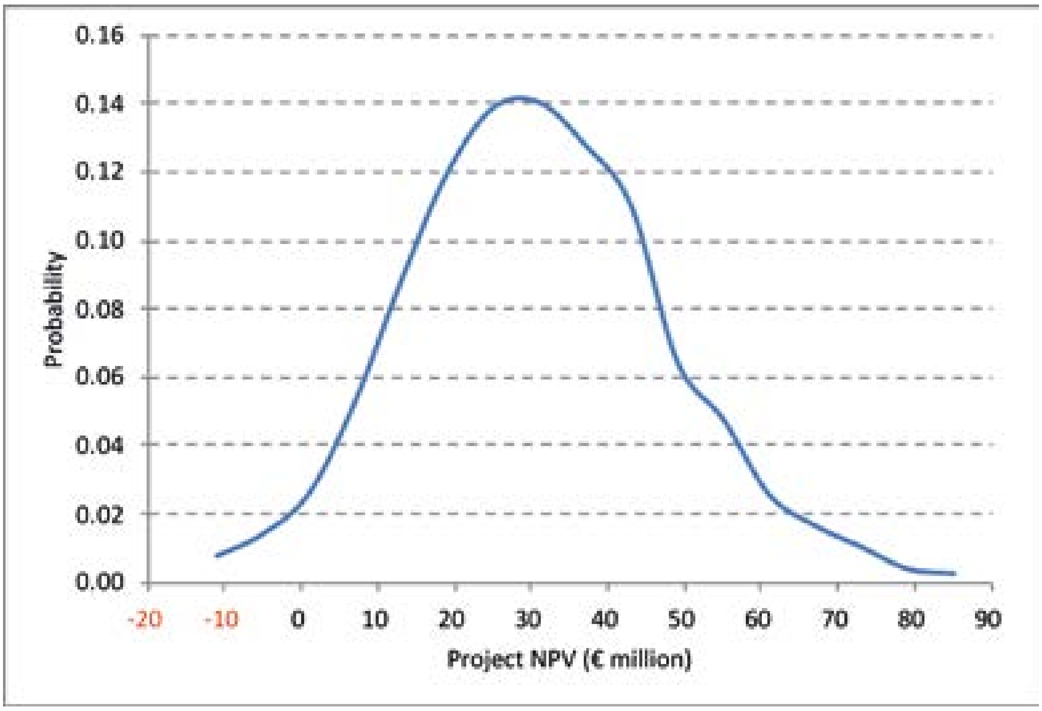

Source: Study estimates

Assume that we will only reject the claim if there is only a 5% chance of it being correct. Using the data for Kriegers Flak, we would require a sample of (at least) 6 years if the true mean load factor is 45% and 11 years if the mean load factor is 48%.[17] Even these numbers understate the problem, since the observations must be for years in which there are no breakdowns or failures at the wind farm. Figure 3 shows that nearly 80% of offshore turbines experience their first breakdown within 8 years of commencing operation. There is effectively a zero probability that the 72 turbines at Kriegers Flak will operate without any breakdowns for 8 years. The consequence is that there is almost no chance of verifying that the wind farm can achieve an average load factor of 52% unless the alternative hypothesis is that the load factor is no greater than 40%, which would only require 4 annual observations.

Importantly, the expected average load factor, a parameter that is fundamental to any assessment of the economics of offshore wind, is not capable of being verified by any well-founded statistical

17 These sample size calculations are based on a standard deviation in the estimated annual load factor of 3.19 percentage points derived from wind data for 39 years from 1981 to 2019 for the Kriegers Flak grid square.

procedure, given the annual and monthly variability of wind speeds at a site such as Kriegers Flak. There is no reason to believe that the variability at Kriegers Flak is significantly different from that at other offshore wind farms, so the conclusion is likely to apply to all offshore wind farms.[18]

The implication for the winner's curse is that the average load factor for Kriegers Flak – and other offshore projects – is an assumed parameter which cannot, for practical purposes, be verified by either the operator or potential investors. When auction bids were submitted in 2016, the most advanced Siemens model that had been certified was the SWT-7.0-154 model. That model had only been certified in January 2016, so there were none in operation. Hence, the assumed load factor used by Vattenfall in bidding for the contract had to be based purely on theoretical calculations.

Even in 2018, when construction of the project started, the main experience of operating turbines from this generation on a large scale were for 3 offshore wind farms:

- Thorntonbank in Belgium with 48 Senvion 6M126 turbines (rated capacity 6.15 MW, hub height 95m) which started operation in 2013. The operator reports a full offtake power purchase agreement based on an expected load factor of 35-39%.[19]
- Westermost Rough in the UK with 35 SWT-6.0-154 turbines (rated capacity 6.0 MW, hub height 100m) which started operation in 2015. In the first full year of operation, Westermost Rough achieved an annual load factor of 45% at a site with an average wind speed at hub height that is 8% higher than for Kriegers Flak.
- Burbo Bank Extension in the UK with 32 V164/8000 turbines (rated capacity 8.0 MW, hub height 108m) which started operation in 2017. The average load factor for the first two full years of operation was 39%.

Based on the experience of these wind farms and the analysis of the statistical data on the distribution of wind speeds at the Kriegers Flak, the best estimate of the expected average load factor for Kriegers Flak is about 43%. Clearly, a bidder might believe that they can outperform this estimate, but neither can a worse outcome be ruled out. The reported assumption of an average load factor of 50% gives a strong impression of adopting a convenient round number. To achieve an average load factor of 50% implies an out-performance of 16%, while an average load factor of 52% implies an out- performance of 21%. In modelling forecast uncertainty, I have assumed that the forecast range for the average load factor is from 0.95 to 1.25 times the base value of 43%.

Prospects for future power prices are similarly uncertain and forecasts are often strongly influenced by short term trends. An important, but often under-appreciated, consideration is the extent to which offshore wind production is inversely correlated with market prices on an hourly or

18 Wind measurements for offshore sites can be collected by lidar units installed on fixed or floating platforms, which is much cheaper than using a meteorological mast. This data is then fed into a model of turbine performance to calculate estimated load factors. While models of turbine performance may not always reflect in-service performance, the real problem is the amount of year-to-year variability in wind conditions which cannot be adequately captured by the relatively short period of on-site measurement prior to developing a site.

19 The rated wind speed (i.e. the wind speed at which output reaches its rated capacity) for the Senvion 6M126 turbine is 14 m/s, whereas it is 13 m/s for the Siemens SWT-6.0-154 and the V164/600 and it is 12 m/s for the SG-8.0-167 turbine. The Siemens-Gamesa turbine also has a slightly lower cut-in speed than the other turbines. These factors should mean that the SG-8.0-167 turbine can achieve a higher load factor in identical wind conditions.

daily basis. For the DK2 and the German markets, the annual average price weighted by offshore wind output divided by the demand-weighted annual average price varied in the range 0.83 to 0.90 during the period 2015–2019. This means that offshore wind output sold into these markets received less than 90% of the average market price.[20] The margins in Belgium, Netherlands and the UK were smaller but may be increasing over time. As the volume of offshore wind production increases it is likely that the discount for offshore wind output relative to the average market price will also increase.

Market power prices in NW Europe are both cyclical and strongly linked to the market price of gas, which is the marginal fuel for much of the time. Between 2011 and 2019 the annual average power price at 2018 prices for region DK2 varied in a range from €25.3 to €53.6 per MWh. There was a peak in 2011 with a trough in 2015 and another peak in 2018. Many commentators tend to interpret cyclical increases – e.g. from 2015 to 2018 – as if they signify longer term trends. On the data available, though this is limited, there is no convincing evidence of any long term trend in real power prices.

On the evidence available, the most defensible baseline assumption for the future market price of power in NW Europe in future is that there will be no long run trend, either up or down. Heavy subsidies for renewable energy – both solar and wind – in Germany and other countries have created markets that are highly volatile. Market prices are determined by the variable costs of operating marginal plants. During periods of moderate or high solar or wind production, market prices are low and occasionally negative. At other times, gas plants are the marginal suppliers but in the last decade they have not earned a sufficient marginal price to keep open existing gas plants, let alone build new ones.

The consequence is that electricity systems have to rely on capacity contracts for backup generators or providers of storage services in order to maintain electricity supplies. The costs of such contracts plus subsidies for renewable generators are funded by fixed consumer levies, with the result that the gap between market prices and the prices paid by users has grown and will continue to increase. In effect, the power market is becoming a market that determines short term dispatch, but it is substantially separate from the energy market for industrial and other consumers.

In this regime it is hard to see why the real market price should increase in the longer term, especially since governments seem determined to support new investment in generation with low or zero marginal operating costs. Indeed, one can plausibly make the case that the long run trend is likely to be for a reduction in the real market price, driven by the increasing dominance of low variable cost generation. For these reasons, I have assumed that the range of forecasts for the trend in the real market price is from -2% to +5% per year with 0% growth as the baseline assumption. This range is not symmetric so as to accommodate the possibility that the decommissioning of older, unprofitable, conventional (dispatchable) power plants will push up market prices at peak periods and thus the average market price –see Pollitt & Chyong (2018).

Finally, there is uncertainty about the future level of O&M costs. Some commentators claim that wind farm operators are adopting practices that will lead to substantial reductions in the variable

20 It is worth noting that the discounts for onshore wind are even larger in the DK2 and German markets – nearly 15% in Denmark and 20% in Germany.

and fixed O&M costs incurred for new offshore wind farms.[21] UK and US projections of levelised costs agree on a total operating cost of about $20 per kWh at 2018 prices, excluding transmission costs for projects commissioned in the early 2020s. Variable O&M costs (per MWh of output) account for 40-50% of total O&M costs. In the case of Kriegers Flak, the wind farm may also incur lower O&M costs because it is located in relatively shallow water with a depth of 16-25m and is only 15 km from land. Both of these factors should reduce the capital cost of constructing the wind farm but, again, there is little evidence concerning the potential reduction in O&M costs. In the forecast risk analysis, I have assumed that O&M costs will be between 0% and 20% lower than the estimates of O&M costs made for the UK and the US. Any claims for a substantially larger reduction in unit O&M costs would have to be supported by new independent evidence.

Figure 7 – Distribution of Kriegers Flak NPV for breakeven scenario: structural and forecast risk

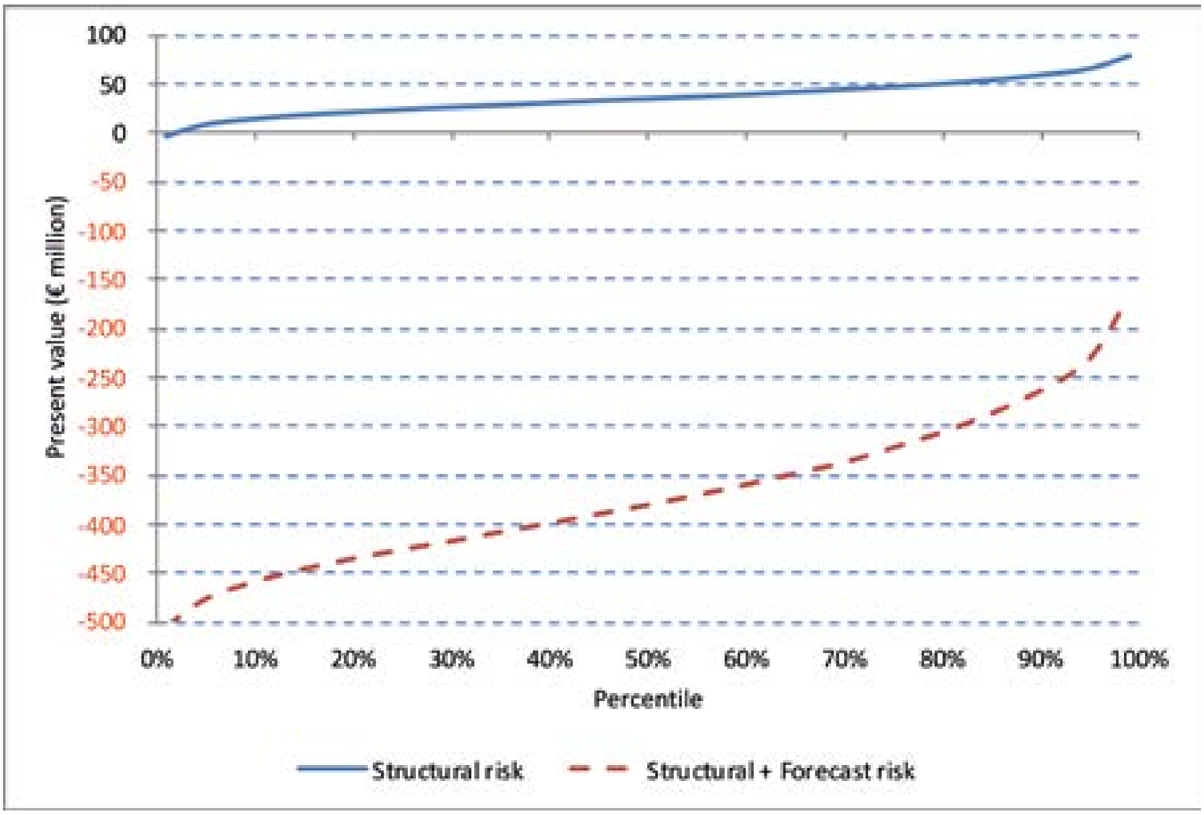

Source: Study estimates

Figures 7 and 8 show the distributions of the net present value of the Kriegers Flak project on different assumptions about risk. Figure 7 contrasts the distributions for (a) the breakeven scenario with structural risk alone, and (b) the combination of structural plus forecast risk. In (b) the breakeven scenario is simply one (very optimistic) draw from the full range of alternative forecasts.

21 Offsetting this view is the fact that Kriegers Flak and other projects use a new generation of large offshore turbines. Past experience indicates that each new generation experiences above-normal rates of breakdowns and other operating problems for up to a decade after its introduction. Thus, it is rather unlikely that either manufacturers or service providers will rely upon substantial reductions in O&M costs until the reliability of the turbines has been proved in practice.

The gap between the two distributions highlights the scale of the winner's curse. If the breakeven assumptions are used, the chance of earning less than the hurdle rate of return is less than 5%. In contrast, the median NPV is about -€380 million when the full range of forecast values is taken in account and there is a 20% chance that Vattenfall will lose its entire investment in the project. The probability of breaking even is practically zero because the probability of the required forecast outcomes is so small.

Figure 8 – Impact of forecast variables on Kriegers Flak NPV

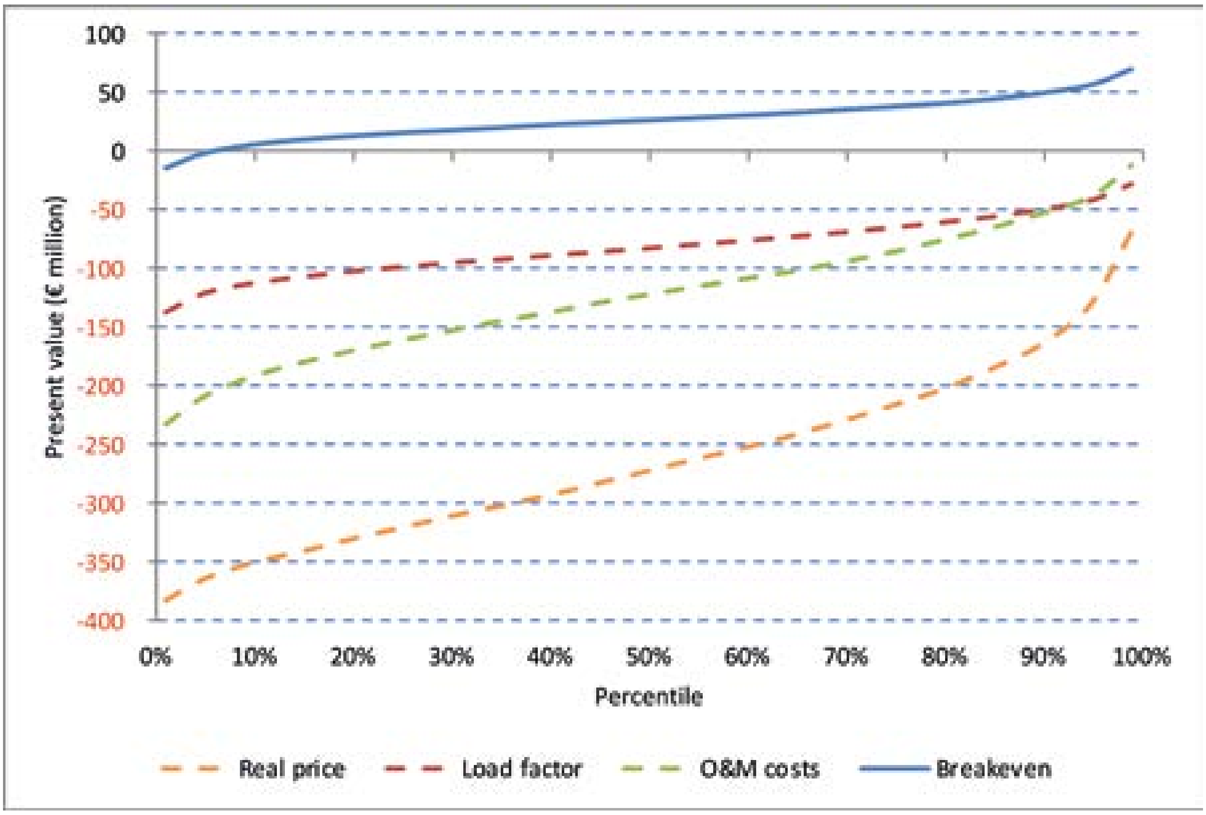

Source: Study estimates

To highlight which assumption has the largest impact on the project NPV, Figure 8 shows the distributions when uncertainty over one of the breakeven assumptions is added to the breakeven model with structural risk. The solid blue line is the same as the solid blue line in Figure 7, while the dashed orange line shows the NPV distributions when uncertainty over the real market price is added. The figure shows that uncertainty over the achieved load factor has the lowest impact on the distribution of NPVs while the uncertainty over the real market price after the expiry of the power purchase contract is clearly the most important source of project risk. This is not a surprise because (i) the range of possible outcomes for the real market price is particularly large, and (ii) the risk is leveraged by fixed transmission costs.

When seen in this light, Vattenfall's auction bid is primarily a bet on the level of market prices from 2033 onwards. It could make a lot of money if power prices in the mid- and late-2030s are high enough, but it will lose a lot of money if market prices for wind generation do not rise sufficiently over the next 10-15 years. However, there is little point in receiving high prices if the decline in turbine performance largely offsets the real increase in market prices. In fact, the bid is a two-pronged gamble that (i) real power prices will increase substantially, and (ii) any age-related decline in the performance of the turbines is small enough not to counteract the benefits of higher power prices. On both counts, the evidence suggests that the project could turn out to be a very expensive bet.

Table 7 shows the values of the guaranteed contract price which ensures that the median value of the distribution of net present values is zero when allowance is made for both structural and forecast risk under a variety of scenarios. The second column of the table indicates the scale of the winner's curse, so that under the main assumptions the contract price would have had to be 49% higher for the winner to have a 50% chance of earning a satisfactory return on its investment.

Table 7 – Impact of alternative scenarios on the break-even contract price

	Required contract price (€ per MWh)	*% increase relative to actual contract price*
Baseline assumptions	74.5	49%
Baseline with no wind discount	72.3	45%
Baseline with PPA volume of 45 TWh	65.8	32%
Baseline + decline in average load factor at 3% per year	79.6	60%
Baseline + decline in average load factor at 4.5% per year	82.5	65%
Baseline + O&M costs increase at 3% per year	77.6	56%

Source: Study estimates

Other scenarios in the table show the sensitivity of the break-even price to assumptions or contract terms. For example, increasing the contract volume from 30 TWh to 45 TWh – in effect increasing the length of the contract from 12+ years to 18+ years – reduces the break-even price by €8.7 per MWh or about 12%. This is the value of the hedge offered by an extended guarantee when the future path of market prices is very uncertain. If the performance of the turbines were to decline over time, the break-even price would be between €5 (3% per year decline) and €8 (4.5% per year decline) higher. A gradual increase in variable O&M costs has a smaller effect but would reinforce the requirement for a higher contract price to cover costs.

It is unclear why Vattenfall bid such a low price for the Kriegers Flak contract. A company like Vattenfall, with large revenues from a mixture of established distribution, heat and power generation businesses can easily lose several hundred million Euros. Danish electricity consumers may be delighted to benefit from support from Swedish taxpayers, the ultimate owners of Vattenfall, but they might look with concern at the propensity of Orsted to bid at auctions in the same manner.

It is difficult as an outsider to know how to view the propensity of large power companies to risk large amounts of money on offshore wind projects such as Kriegers Flak. Clearly, incentives for optimism bias are very strong and, in consequence, no one should treat auction bids as a reliable guide to the long run cost of offshore generation. It may be a regular occurrence for companies to lose money on large projects but that cannot be a basis for the sustainable operation of a sector. In this case, it is taxpayers and energy customers in Sweden who will bear the cost of Vattenfall's optimism bias, though they may not be aware of the fact. Elsewhere the losses are likely to fall on investors, workers and customers as companies with money to burn are displaced by more cautious and efficient operators.

Finally, I should emphasise again that the long term costs of offshore wind generation are falling, if not at the rate implied by absurd auction prices. The break-even contract prices for Kriegers Flak are well below the equivalent prices for early offshore projects, even when we add back the cost of transmission. Changes in the design of new turbines mean that the expected load factor for a given distribution of wind speeds will be higher. This increases the utilisation of the capital invested and may eventually bring down O&M costs.

All of this means that there could be a reasonable future for offshore wind generation, but not one – at least for the medium term – that is "subsidy-free". At the moment a fair portion of the subsidies is coming from unwise investors but that will not last. Bringing costs sufficiently down to match, for example, gas generation after allowing for transmission and storage costs, remains a large challenge. Even if countries in Western Europe are willing to forego the low cost and other advantages of gas generation, that is not true in rest of the world.

The financial model has highlighted both (a) the scale of the winner's curse in offshore power auctions, and (b) the extent to which problems with the reliability and performance of wind turbines push up the long term cost of both onshore and offshore wind generation. A combination of the decline in performance due to age and an increase in operating costs adds more than €10 per MWh – equivalent to 14% – to the break-even contract price for Kriegers Flak. The arrangement is a little artificial because the structure of the power purchase contract implies a large drop in the power price that is received after the initial contract period, but that is not so unusual with other power purchase agreements.

The larger point is that a decline in the performance of offshore turbines due to ageing at a rate of 4.5% per year means that the effective economic life of such turbines is little more than 15 years. Power prices and operational arrangements after that period will have a very small effect on the decision to invest or not to invest in offshore generation.

8. Financial investors and Kriegers Flak

In the last section I focused on the decision to bid for and invest in the Kriegers Flak project from the perspective of the project sponsor, Vattenfall. Clearly it is – at least publicly – committed to a set of assumptions which underpin a belief that the project has a reasonable chance of breaking even, with the possibility of a significant profit if performance and market conditions are favourable.

There is a different perspective that may be equally important in the longer term. Even large power companies like Vattenfall cannot afford the cost and risks of financing a large number of projects on the scale of Kriegers Flak. As discussed earlier, almost all wind developers rely on selling majority stakes in their projects to financial investors who take on a large share of the risks and rewards of future operations. It is usual for the original developer to enter into a contract to operate and maintain the project, as well as retaining a significant minority stake. Partly this is a matter of aligning incentives and risks, but the developer may hope to offset a part of any losses with profits from its operating business.[22]

Now, let us consider the project from the perspective of PensionDanmark or any similar large financial investor. Such institutions sometimes come under significant pressure to invest in flagship projects under the rubric of "alternative investments". This may be a direct investment or one made through a notionally arms-length infrastructure fund such as Copenhagen Infrastructure Partners. In either case, the problem is how the financial investor should address the risk of being on the wrong side of the winner's curse. This can be best approached by considering the valuation of the 75% stake in the project on the assumption that the acquisition will be accompanied by a reduction in the project debt to no more than 50% of the original investment.

Figure 9 – Distributions of the investment return for structural risk alone

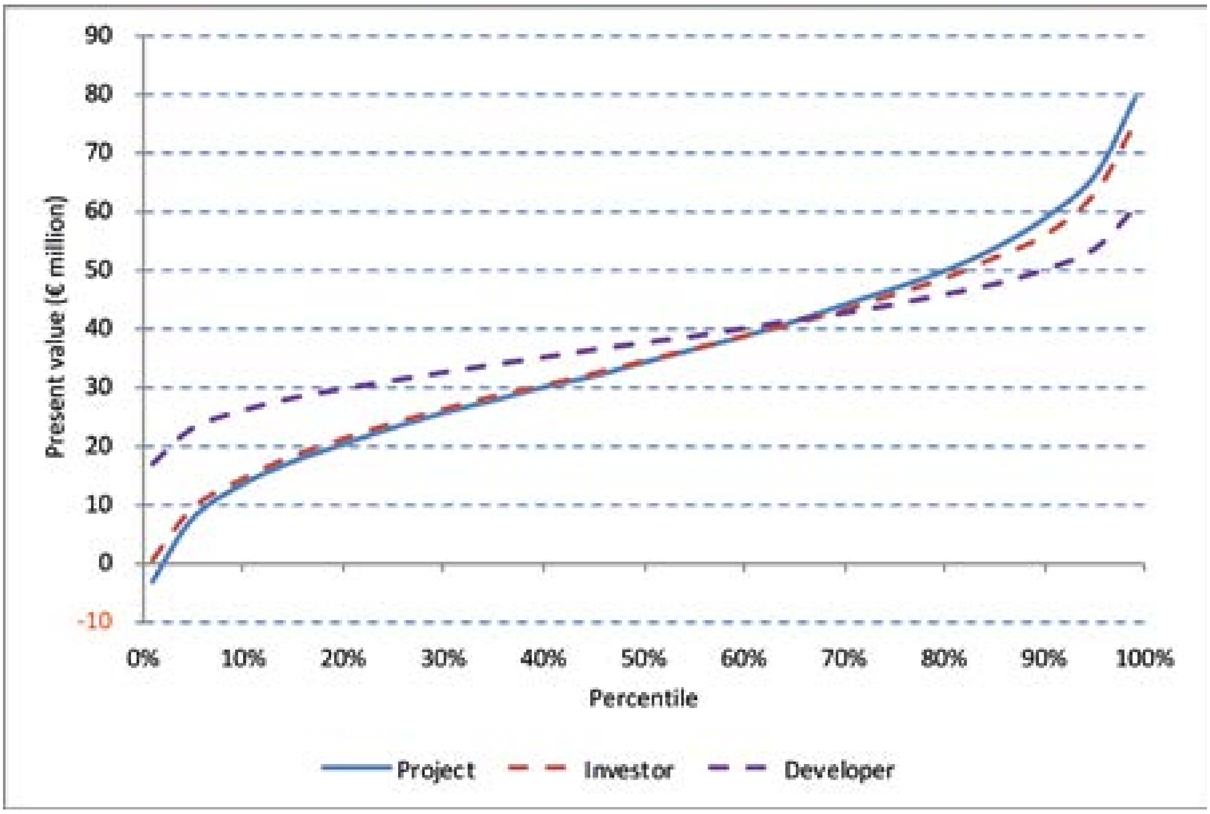

Source: Study estimates

Based on its breakeven assumptions, the project developer is likely to estimate that a reasonable valuation of the 75% stake would be about €480 million. Remember that the developer has to believe that its breakeven assumptions will be validated, since otherwise it would have to acknowledge the prospect of incurring a very large loss on the project. On that basis, Figure 9 shows the distributions of present values for the project, the developer and the financial investor for a valuation of €480 million. The financial investor has a better than 95% probability of earning at least 6% on its investment. The developer de-risks the project by accepting a lower return at the top of the

22 The cost of the operating contract is included in the O&M cost discussed above.

distribution in exchange for a higher return at the bottom of the distribution. Both parties would benefit from the transaction under these assumptions.

However, now consider the same distributions in Figure 10 which allows for the winner's curse, i.e. incorporating forecast risk in addition to structural risk. The transaction would lead to a very large transfer of risk from the developer to the financial investor. The developer will still make a loss on the project but most of the consequences of its error in bidding at the auction will fall on the financial investor.

Figure 10 – Distributions of the investment return for structural risk alone

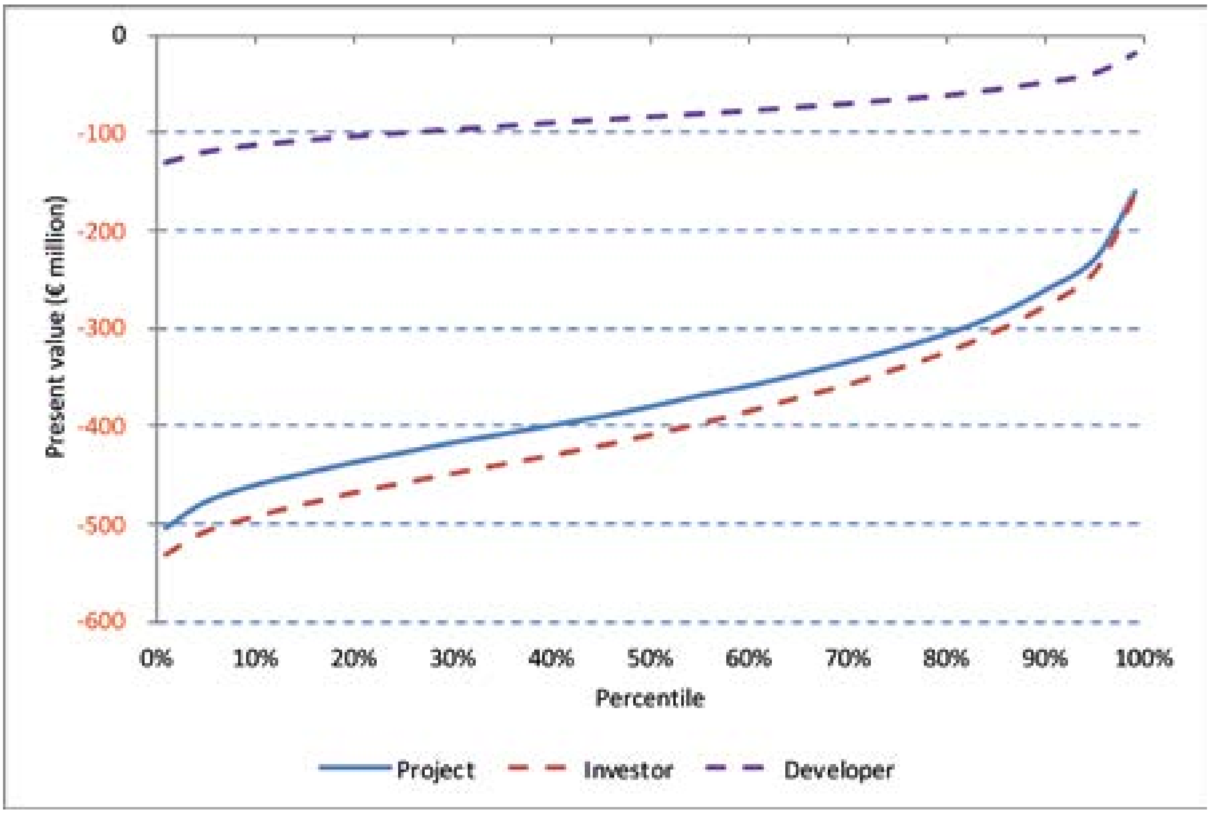

Source: Study estimates

To understand and assess the options available to both parties, it is helpful to construct what is known as a pay-off matrix, often associated with analysing strategic choices such as the Prisoner's Dilemma. The matrix – shown as Table 8 – is larger in this case because we have to take account of uncertainty about the state of the world – i.e. will the breakeven assumptions turn out to be valid or not – as well as the choices available to each party. In each combination the pay-off is taken as the median of the distribution of the present value for each party. The financial investor can choose to do a deal or not; if it does a deal, it would pay amounts for the 75% stake in the project varying from €480 million (if it is convinced that the breakeven scenario is correct) down to €60 million (if it assumes that the forecast scenario is correct). The mixed strategy assumes an outcome that is equivalent to 50% Breakeven and 50% Forecast. This is a way of testing whether some combination of strategies might be better than either of the pure strategies. The row marked "Financial investor minimum" shows the worst outcome for each of the three strategies under different states of the world and choices by the developer. The least bad (maxi-min) outcome is not to do a deal at all. If it were to do a deal, this should only be on the basis of the forecast scenario, i.e. by paying as little as possible for the stake. Any other strategy will run the risk of incurring a large loss.

Table 8 – Pay-off matrix for financial investment under alternative strategies

			Financial investor strategy		
Developer strategy		*State of the world*	*Breakeven*	*Forecast*	*Mixed*
	Deal price (€ mln)		480	60	270
FI pay-off (€ mln)	Deal	Breakeven	40	460	250
		Forecast	-400	18	-193
	No Deal	Breakeven	0	0	0
		Forecast	0	0	0
Developer pay-off (€ mln)	Deal	Breakeven	32	-342	-155
		Forecast	-89	-463	-276
	No Deal	Breakeven	34	34	34
		Forecast	-379	-379	-379
Financial investor minimum			-400	0	-193
Developer minimum			-379	-463	-379

Source: Study estimates

On the other side, the developer has already committed itself to the breakeven assumptions through its bid for the project and the decision to go ahead with the investment. If it sticks by those assumptions, then the no deal strategy is its best option as there is no advantage to accepting a price that is less than €480 million for the 75% share in the project. Retaining the whole project is risky, but the worst outcome in this scenario is better than accepting the price that the financial investor is likely to offer.

This is only part of the story, because the financial investor must also consider whether the assumption of no decline in performance and no increase in O&M costs due to ageing is likely to be correct. If the conclusions of the analysis earlier in this paper were to apply in future, the pay-offs to the financial investor will fall by about €380 million in the breakeven scenario and by about €190 million in the forecast scenario. In this case, the chance of earning a satisfactory return on the investment, even when a low price is paid for the stake in the project, is low. Any financial investor willing to put money into such a project is either incompetent or entirely careless of the interests of those whose money is at risk.

Conclusions

There is nothing special about the incidence and nature of the turbine failures and the decline in operating performance discussed in this paper. All complex pieces of engineering design involve economic and physical trade-offs between cost and reliability. Currently, much attention is focused on blade erosion causing loss of aerodynamic performance as wind turbines age. In the past there have been problems with gears, generators, etc. In future, who knows but the trade-offs will always be present. It would be possible to manufacture turbine blades with the type of materials that are used for modern airplane wings, but the cost of a large passenger jet is a minimum of $250 million whereas even a large offshore wind turbine costs less than 5% of that sum.

As in every other industry, turbine manufacturers and operators have made an economic trade-off between capital cost, O&M costs, reliability and performance over time. There is nothing improper or unreasonable about such trade-offs. It is the job of engineers to try to optimise design to get the combination of costs and other characteristics that best suits the requirements of their customers. On the other hand, what is not excusable is the pretence by some that such trade-offs do not exist or that they do not affect the long term performance and economics of wind generation.

The main conclusions from the empirical and financial analyses are:

- The generation of wind turbines with a capacity of less than 1 MW are significantly more reliable and have experienced a lower decline in performance with age than later generations of turbines with a capacity of 1-2 MW and 2+ MW. This difference largely accounts for the difference in age effects between Denmark and the UK in the original study because a much higher proportion of data for Denmark in 2012 was for turbines of less than 1 MW.

- New generations of turbines seem to experience an extended period of low reliability and poor performance following their introduction. This period may last for a decade or more. Experience suggests that the causes of such initial problems can eventually be dealt with but there is clearly a lengthy period of learning for each new generation that should be taken into account.

- Offshore wind farms experience higher rates of breakdowns and declines in performance with age than onshore wind farms. Since a new generation of very large (6+ MW) turbines came into use in the mid-2010s for offshore wind farms, it is likely that the combination of new generation effects and offshore operation will have a major impact on the overall performance of new offshore wind farms completed in the last 2-3 years or still under development.

- Headline prices from auctions for offshore wind do not provide a reliable guide to the future cost of offshore wind generation. Costs are falling but not at anything like the rate implied by simple comparisons of auction bid prices, which have been seriously affected by the winner's curse. Allowing for the risks concerning future performance and power market prices, a reasonable estimate of the cost of offshore wind power would be €70-80 per MWh plus transmission costs. Even then, operators seem to be assuming higher levels of reliability and performance than can reasonably be expected on the basis of evidence from the last 10-15 years. There is a long way to go to reach "grid parity", even when no allowance is made for intermittency.

- Any competent financial institution should avoid investments in offshore wind projects for which development rights have been awarded by some form of auction. The risks and potential size of the winner's curse are so large as to mean such investments will be unlikely to earn a satisfactory return even in the best of circumstances. There is a high probability of losing a large portion of the money invested in such projects. This is because the developers will be unwilling to accept a valuation of the projects that involve substantial write-downs to reflect a reasonable assessment of the risks and probable revenues from the projects.

None of these conclusions should surprise any analyst with a reasonable understanding of the performance of new technologies in a wide range of industries. The tendency to over-sell novelty and to underestimate the operating problems that have to be worked out for all new technologies is routine in a world with short attention spans and a desire to believe in simple solutions to difficult problems.

What is more troubling is the ignorance and naivety of policymakers, academics and commentators. Optimism bias in claims about the cost and performance of infrastructure and other projects has been endemic for millennia. Currently, the offshore wind business is trapped in a speculative bubble akin to property and financial bubbles in which claims and expectations lose their base in reality. Genuine but modest progress in bringing down capital and operating costs is translated through endless press releases into a belief that the basic economics of the business has changed. Of course it has not, and over time a lot of expectations will be disappointed, but there will be little accountability for those who promote the bubble.

Looking back to the canal boom in the 1700s and the railway boom in the 1800s, during the early stages the losses caused by over–optimism were borne by private investors who lost much of their capital. Gradually, they became more wary, prompting two responses. First, profitable operators acquired some of the marginal or failed projects with the goal of exploiting network economies of scale. Second, governments and public companies started to promote and fund projects on the grounds that 'either the market had failed or the projects were of "strategic importance".' Large amounts of money were lost in this way but eventually modern networks began to emerge.

People who favour public intervention draw on such examples to advocate the allocation of large sums of public money – usually through indirect forms of support – to underwrite the rapid expansion of renewable energy. There is little basis for such policies. Any network economies from wind generation are tiny, if they exist at all. It is unlikely that any economic clusters will develop near to offshore wind farms, which is the core of the argument for supporting transport infrastructure. Even more important, the economic life of a wind farm is no more than 20 years, while railway lines have an economic life of 100 years or more, so that the combination of low discount rates and delayed benefits does not arise.

What remains is the familiar but usually misunderstood argument concerning the need to support technological development. Wind turbines, even offshore turbines, are a relatively mature technology. There is no justification to support the introduction of new generations of turbines other than via loans on commercial terms. Indeed, the evidence in this paper shows that there are good reasons not to provide such support, because each new generation gives rises to issues with cost, performance and reliability that can only be addressed by the unromantic method of making gradual engineering improvements. Providing public support for investment in new wind farms worsens the situation by encouraging the adoption of immature technologies on a large scale. This is a classic case where making haste slowly is the right approach. The endpoint in 20 years will be no different but the cost of reaching it will be significantly lower.

At the beginning of this paper I said that its conclusions would be the usual mixture of shades of grey. The economics of wind generation, especially offshore generation, are improving, but important issues of performance and reliability, as well as of integration in electricity systems, have not yet been solved. Wind generation is likely to play an increasing role in electricity systems, provided that land and wind resources are favourable. On the other hand, it is not clear whether, in system terms, the economics of wind power will be better than the economics of large scale solar generation, especially in latitudes from 40°N to 40°S. Addressing the problems of performance and reliability for wind generation highlighted in this analysis will be crucial in determining the balance between wind and solar.

Acknowledgements

A. The data in the Danish register of wind turbines and output was obtained from the Danish Energy Agency (ENS) repository at **www.ens.dk**. I am particularly grateful to Emil Axelsen & Iben Spliid of ENS for providing me with up-to-date data on output in 2018 & 2019.

B. The NASA daily weather data by grid square were obtained from the NASA Langley Research Center (LaRC) POWER Project funded through the NASA Earth Science/Applied Science Program – see **https://power.larc.nasa.gov/**. I am grateful to members of the POWER project for facilitating access to the detailed data that they have compiled.

References

J. Aldersey-Williams, I.D. Broadbent & P.A. Strachan (2019) "Better estimates of LCOE from audited accounts – a new methodology with examples from UK offshore wind and CCGT", *Energy Policy*, Vol 128, May, pp. 25-35.

A. Ansar, B. Flyvbjerg, A. Budzier & D. Lunn (2014) "Should we build more large dams? The actual costs of hydropower megaproject development", *Energy Policy*, Vol 69, June, pp. 43-56.

R.J. Barthelmie & S.C. Pryor (2003) "Can satellite sampling of offshore wind speed realistically represent wind speed distributions?", *Journal of Applied Meteorology*, Vol 42, January, pp. 83-94.

E. Berndt, C. Kolstad & C-K Lee (1993) "Measuring the energy efficiency and productivity impacts of embodied technical change", *Energy Journal*, Vol 14, No 1, pp. 33-55.

C. Carrillo, J. Cidras, E. Diaz-Dorado & A.F. Obando-Montano (2014) "An approach to determine the Weibull parameters for wind energy analysis: the case of Galicia (Spain), *Energies*, Vol 7, 2676-2700.

B. Flyvbjerg (2009) "Survival of the unfittest: why the worst infrastructure gets built and what we can do about it", *Oxford Review of Economic Policy*, Vol 25, Autumn, pp. 344-367.

B. Flyvbjerg, M.S. Holm & S. Buhl (2002) "Underestimating costs in public works projects", *APA Journal*, Vol 68, No 3, pp. 279-296.

G.A. Hughes (2012) *The performance of wind farms in the United Kingdom and Denmark*, London: Renewable Energy Foundation.

G.A. Hughes (2019) *Who's the Patsy? Offshore wind's high-stakes poker game*. GWPF Note 18, London: Global Warming Policy Foundation.

G. A. Hughes, C. Aris & J. Constable (2017) *Offshore wind strike prices: Behind the headlines*, GWPF Briefing 26, London: Global Warming Policy Foundation.

G.N Petersen & H. Olafsson (2016) "Frequency of winds in various terrain", Conference paper: EGU2016-10548, EGU General Assembly 2016, Vienna.

M. Pollitt & C.K. Chyong (2018) *Europe's Electricity Market Design: 2030 and beyond*, Brussels: Centre on Regulation in Europe.

S.C. Pryor, M. Nielsen, R.J. Barthelmie & J.Mann (2004) "Can satellite sample of offshore wind speeds realistically represent wind speed distributions? Part II: Quantifying uncertainties associated with distribution fitting methods", *Journal of Applied Meteorology*, Vol 43, May, pp. 739-750.

S&P Global (2018) *Vattenfall AB*, S&P Global RatingsDirect, 8th November 2018.

I. Staffell & R. Green (2014) "How does wind farm performance decline with age?", *Renewable Energy*, Vol 66, pp. 775-786.

APPENDIX A
DATA ON TURBINES, GENERATION AND WIND SPEEDS

Turbine register

The Danish Energy Agency maintains a register of all wind turbines with a generating capacity of at least 6 kW that have been installed in Denmark since the 1970s. This register includes information on the characteristics of each turbine including capacity, hub height, rotor diameter, manufacturer, model type, location, date of initial operation, date of decommissioning (if applicable), etc. The information in the register is relatively complete and the main omissions concern smaller turbines, many of them with a capacity of less than 100 kW.

Data on turbine output is recorded by year from 1977 to 2001 and monthly thereafter. Again, the data is relatively complete but with months in which either zero output is recorded or output is missing. It is possible that missing values represent a failure to record output but this is unlikely as most operators have a strong incentive to record positive output when the turbines are producing because of the way in which subsidies are paid. In addition, there are sequences of missing values followed by restarted production that indicate outages. The date on which decommissioned turbines are taken out of production is recorded and data for subsequent months are excluded from the dataset.

Output for turbines that form part of wind farms – defined as groups of turbines from the same manufacturer, of the same type, located in the same local authority and commencing operation within a short period of each other - is recorded in different ways. In the case of small wind farms the turbines may be individually metered, but it is more common for them to report from a single meter or a small number of meters covering the wind farm. In such cases, the metered output is divided between all turbines, so that the monthly output figures are averaged across all turbines connected to the meter. A failure at an individual turbine will reduce the apparent output for all turbines, until the failed turbine is either repaired or marked as having been decommissioned.

People commenting on the original paper have tended to assume that the decline in performance that was reported should be observed as a (more or less) steady decline in the output from individual turbines after controlling for the amount of wind. Such a decline may occur but it is not the primary mechanism by which performance deteriorates. Instead, it is the frequency and length of outages at individual turbines due to failures of mechanical, electrical or other equipment which leads to a reduction in expected output from wind farms or groups of turbines over time.

In a few cases, particularly at offshore wind farms, there may be a major failure which results in the loss of output from an entire wind farm lasting for a period from a week to months. Often these are linked to an electrical failure in the transmission system or main sub-station serving the wind farm. Repairing such failures for offshore installations may require considerable time to mobilise ships, people and equipment since repair work offshore is much more difficult to organise and carry out than equivalent work onshore. This is the trade-off for taking advantage of the stronger winds and higher capacity of offshore sites. Hence, one would expect that a larger portion of the loss in performance at offshore wind farms will be due to extended major breakdowns than for onshore wind farms. This is, indeed, what the Danish data shows.

There are a very small number of cases in which the load factor calculated from the reported output data exceeds 100%. In most of these cases the error was clearly due to misreporting that had not been picked up by Danish Energy Agency staff, where all of the output from a small group of turbines has been reported as if it had been produced by a single turbine. Such errors were corrected with output divided equally across each of the turbines in the group. For 3 monthly observations (out of nearly 670,000) the load factors were extreme outliers in the overall monthly distribution. Rather than exclude these observations, the load factors were capped at 85% for the month concerned, a value higher than that observed for any other cases with similar average monthly wind speed.

In the analysis I have examined the difference between the results obtained when (a) cases, for which output is zero or missing are omitted because log(Load Factor) is not defined, and (b) cases, for which output is zero or missing are treated as having a load factor of 0.1%. Dropping observations with output that is zero or missing tends to reduce the rate of decline because such observations are more frequent for older turbines. For the main discussion I focus on the results obtained when such cases are not excluded, since no output is one symptom of the major breakdowns and equipment failures that occur as turbines age.

Table A1 shows the distribution of turbines in my working dataset by age of installation and type, which combines turbine capacity offshore and onshore. Offshore turbines with a capacity of less than 2 MW were few in number and were an experiment that was not followed up. The early offshore wind farms used turbines with a capacity of 2.0 to 2.3 MW that were, in effect, variants of onshore turbines. Most offshore wind farms commissioned since 2009 have used turbines designed for offshore use with a capacity of at least 3.0 MW, increasing to 8.3 MW at Horns Rev 3 which was commissioned in 2019.

Table A2 provides information on turbines that have been decommissioned. One point worth noting is that 13% of the medium (1 – 2 MW) onshore turbines and 23% of the large (2+ MW) onshore turbines commissioned in the period 2000-04 have been decommissioned. The average age at decommissioning was 16 years for medium turbines and 6 years for large turbines. It seems fairly clear that the early units of the larger turbines were troubled with a variety of problems. These failures continued into the later periods, though at a lower rate, with 21 turbines commissioned between 2005 and 2014 having been decommissioned with average ages at decommissioning of less than 5 years. It is not unusual for the manufacturers of power generation equipment to encounter such problems when introducing a new range of models marking a generational shift in design and capacity. One would expect that turbines commissioned from 2010 onwards should prove more reliable, but the early problems are reflected in the performance of large turbines in the period 2002-10.

Wind speeds

The data on wind speeds at 10 and 50 metres above surface level have been extracted from the NASA POWER database which provides daily estimates of mean, minimum and maximum wind speeds for a 0.5° grid – about 55 km N-S and 32 km E-W at a latitude of 55°N – for the period from January 2002 to September 2019. The estimates of mean wind speed have been used to compute an estimate of the average wind shear coefficient. This was used with data on hub height for each turbine to estimate the daily wind speed at hub height.

Wind turbines produce little or no power when the wind speed at hub height is less than 4 m/s and most models reach their rated capacity at a wind speed of 12-14 m/s, so that output does not increase with higher wind speeds. Hence, it is useful to estimate the proportion of time that wind speeds are less than or equal to 4 m/s and greater than or equal to 12 m/s. This has been done by assuming that the distribution of wind speeds within each day can be approximated by a modified Rayleigh distribution. The Rayleigh distribution is a special case of the Weibull distribution in which the shape parameter is set equal to 2. The advantage of the Rayleigh distribution is that it has only one parameter – the scale parameter – which can easily be estimated from the mean. The modification is introduced because the Rayleigh distribution has a minimum of zero, whereas the reported minimum wind speeds are well above zero for many days and is over 20 m/s in the most extreme cases. To allow for this I have assumed that the Rayleigh distribution is shifted, so that the minimum value is equal to the minimum daily wind speed.

It is interesting to consider whether there are systematic differences in the average wind speeds at the sites where different types of turbines have been installed or for turbines commissioned in different periods. The hub height of turbines varies systematically across turbine types from 45m for small turbines to 53m for medium turbines and 80m for large turbines. There has also been a clear trend towards taller, large and offshore turbines. The average hub height for large turbines has gone from less than 70m up to 2004 to more than 80 metres after 2010. Similarly, for offshore turbines the average hub height has increased from 69m in 2000-04 to 101m in 2015-19. As a consequence, even with the same wind conditions the average wind speeds at hub height have increased as turbine manufacturers have designed bigger turbines to exploit the increase in wind speed at greater heights.

Table A3 shows various indicators of the distribution of wind speeds at 50m and at hub height by year of commissioning and turbine type. Cells with less than 10 turbines have been excluded from the comparisons. Overall, there is little evidence of significant differences in the average wind speed at 50m for different types of onshore turbines. Offshore wind farms do appear to enjoy higher wind speeds than onshore wind farms with the exception of those commissioned in the period 2010-14. However, once the adjustment to hub height is made, the average wind speed at hub height for large onshore turbines is the same as that for offshore turbines, because up to 2014 the average hub height for large onshore turbines was significantly greater than that for offshore turbines. This changed from 2015 onwards. The offshore turbines commissioned in 2015-19 are considerably taller and have a higher average wind speed at hub height than large onshore turbines commissioned in the same period.

The distribution of wind speeds is important for the total yield from a wind farm. Again, the picture is somewhat mixed. Offshore wind farms tend to experience winds of at least 12 m/s for a greater percentage of the year than large onshore turbines, but they also have low speeds (not more than 4 m/s) for more of the year. Overall, the variance of wind speeds is larger for offshore wind farms than for large onshore turbines. Again, it seems that developers and manufacturers are opting for taller offshore turbines in order to get best advantage from the variability in offshore wind conditions.

Table A1 – Distribution of turbines by start year and type

	Onshore			*Offshore*	*Total*
Start year	*0.5 - 1 MW*	*1 - 2 MW*	*2+ MW*	*2+ MW*	
1986-94	71	3	0	0	74
1995-99	1,683	71	2	0	1,756
2000-04	813	276	64	193	1,346
2005-09	26	20	72	100	218
2010-14	3	0	336	202	541
2015-19	4	4	308	54	370
Total	2,600	374	782	549	4,305

Source: Study estimates

Table A2 – Distribution of decommissioned turbines by start year and type

Start year		*Onshore*			*Offshore*	*Total*
		0.5 - 1 MW	*1 - 2 MW*	*2+ MW*	*2+ MW*	
1986-94	Number decommissioned	8	3			11
	Av age at decom (months)	283	172			253
1995-99	Number decommissioned	172	1	0		173
	Av age at decom (months)	212	226			212
2000-04	Number decommissioned	27	35	15	0	77
	Av age at decom (months)	180	194	74		166
2005-09	Number decommissioned	2	1	6	0	9
	Av age at decom (months)	67	17	49		49
2010-14	Number decommissioned	0		15	0	15
	Av age at decom (months)			41		41
2015-19	Number decommissioned	1	1	9	0	11
	Av age at decom (months)	12	17	17		17
Total	Number decommissioned	210	41	45	0	296
	Av age at decom (months)	208	185	48		181

Source: Study estimates

Table A3 – Distribution of wind speeds by start year and turbine type

		Onshore			*Offshore*
		0.5 - 1 MW	*1 - 2 MW*	*2+ MW*	*2+ MW*
1986-94	Av wind speed at 50 m	8.0			
	Av wind speed at hub height	7.7			
	Av % of winds ≤ 4 m/s	13.2			
	Av % of winds ≥ 12 m/s	9.2			
1995-99	Av wind speed at 50 m	7.9	7.9		
	Av wind speed at hub height	7.8	7.8		
	Av % of winds ≤ 4 m/s	12.5	11.9		
	Av % of winds ≥ 12 m/s	10.0	9.8		
2000-04	Av wind speed at 50 m	7.9	7.9	7.9	8.3
	Av wind speed at hub height	7.8	8.0	8.3	8.7
	Av % of winds ≤ 4 m/s	12.5	11.4	11.0	10.3
	Av % of winds ≥ 12 m/s	9.8	11.2	13.4	17.4
2005-09	Av wind speed at 50 m	7.9	8.2	8.1	8.7
	Av wind speed at hub height	7.9	8.5	8.7	8.9
	Av % of winds ≤ 4 m/s	11.7	10.5	9.5	8.9
	Av % of winds ≥ 12 m/s	10.2	15.2	16.6	18.6
2010-14	Av wind speed at 50 m			8.1	8.0
	Av wind speed at hub height			8.8	8.5
	Av % of winds ≤ 4 m/s			8.7	11.2
	Av % of winds ≥ 12 m/s			17.0	16.7
2015-19	Av wind speed at 50 m			8.1	8.8
	Av wind speed at hub height			8.8	9.6
	Av % of winds ≤ 4 m/s			8.8	6.1
	Av % of winds ≥ 12 m/s			17.1	23.0
All	Av wind speed at 50 m	7.9	7.9	8.0	8.3
	Av wind speed at hub height	7.8	8.0	8.7	8.7
	Av % of winds ≤ 4 m/s	12.5	11.5	9.3	10.3
	Av % of winds ≥ 12 m/s	9.9	11.1	16.2	17.5

Source: Study estimates

APPENDIX B
TECHNICAL PROGRESS AND MODELS OF WIND FARM PERFORMANCE

In many industries and activities the measured performance of components, machines or personnel over time reflects the combined effect of various changes. This may be illustrated by an example entirely separate from the controversy that may surround the performance of wind turbines. Consider a combination of a driver and a standardised racing car that are performing timed laps of the Le Mans 24-Hour racing circuit over a number of years. The achieved time for each lap is a random variable that depends upon weather, track and traffic conditions. With experience and fine tuning of the engine and bodywork, the driver and the team may learn how to improve the average lap time from one year to the next. This is the disembodied technical progress that contributes to an improvement in performance over time. The term "disembodied" is used because the car, engine, etc are assumed to be the same over the period examined. On the other hand, the driver, engine and car body may begin to show the effects of ageing over time as wear and tear, slower reactions, etc contribute to a gradual reduction in performance as time passes.

The example is a little stretched because the competition for success at Le Mans means that both drivers and racing cars are replaced if there is any suspicion that they are slowing down or performing less well than in the past. Still, it provides a framework for considering the separate influences of technical progress, ageing and random variables such as weather conditions on the achieved performance of wind turbines over a period of years.

The statistical issue is illustrated in Figure B1 in which turbine age is plotted on the x-axis and turbine performance on the y-axis for a set of turbines observed from 2010 to 2018 with different starting ages and a random component in the data. We assume that the influence of weather and other factors has been controlled. The data for the turbines observed in 2010 are shown as dark blue diamonds with the dark blue solid line showing the fitted relationship between age and turbine performance in 2010. Similarly, the data for the turbines observed in 2018 are shown as light blue squares with the light blue solid line showing the fitted relationship between age and turbine performance in 2018. The line for 2018 is shifted upwards as a consequence of disembodied technical progress, but the effect of age on performance remains. As time passes, new turbines are built and the existing ones age, so more turbines are observed in 2018 than in 2010.

The two dashed lines in the figure show the fitted relationship between age and turbine performance for two separate turbines. Turbine T1 was 2 years old in 2010 and reaches 10 years in 2018, while turbine T2 was 8 years old in 2010 and reaches 16 years in 2018. These fitted lines incorporate the composite impact of technical progress (shifting up the age-performance relationship as time passes) and ageing. As a consequence, the lines are much flatter than the pure age-performance lines and they do not tell us about the impact of age on performance independent of any technical progress affecting all turbines.

In principle, it is possible to separate the influences of age and year/technical progress with the fixed effects model which was used in the original study.[1] The estimation is equivalent to pooling all of the curves fitted for each turbine over time and then using differences in the age profiles of turbines in each year to separate the two effects. However, the method relies critically on a statistical assumption that there is no correlation between the year in which performance is measured and the fixed effects for individual wind farms in the sample as a whole.[2] That assumption is no longer satisfied because of a change in the composition of the dataset. The original dataset was dominated by small (< 1 MW) turbines installed between 1996 and 2002, most of them in groups dispersed around the country. The rate of new turbine installations was very low between 2003 and 2007 and the rate of new onshore installation since 2008 has been an average of about 70 per year and they are more geographically concentrated. As a consequence, what was a dataset in which 85% of the turbines had been installed in a short period with a large amount of variation across turbines in different locations has become one with a substantially higher correlation between the factors that determine the fixed effects and time.

The change in the composition of the dataset means that estimates of the effects of age and time (technical progress) are quite unstable when both variables are included in the fixed effects model. Small changes in the sample or other variables lead to large changes in the coefficients for age and time, though the difference between the two coefficients does not change. Hence, the fixed effects model cannot be used to obtain separate estimates for the influence of ageing and technical progress on performance.

An alternative approach which can be used with the extended dataset is to estimate the age-performance relationships for each year separately from 2002 to 2019 and use a method known as Seemingly Unrelated Estimation (SUE) to pool the coefficient estimates and thus obtain pooled estimates of the slope of the age-relationships for the full period or for sub-periods. The rate of technical progress can then be recovered by running a regression of the constant term in each equation against year. The point about SUE is that while each equation is distinct, it incorporates cross-equation correlations in the variance-covariance matrix. In this case, that can be constructed to allow for the fact that the equations are estimated using overlapping samples of turbines in each year.

This method will generate an unbiased but inefficient estimate of the coefficient on age in the relationship between turbine performance and age. An inefficient estimate is one whose standard error is larger than the minimum that could be achieved by using all of the information available,

1 The fixed effects model focuses on variation within panels, i.e. variation over time in the performance of each turbine. The fixed effects refer to the characteristics of turbines that do not change over time such as size, location, model type, etc. The influence of such unchanging factors is removed by re-formulating the equation that is estimated in terms of differences between the observed values for each year and the average values for each turbine for all years.

2 In more technical terms, the assumption is that the correlation between the exogenous regressors and the u_i panel fixed effects is small or zero. This condition is not satisfied for the extended dataset when time (year) is included among the regressors in order to estimate the rate of technical progress. As a consequence, this means that the statistical assumptions required to get unbiased estimates using the random effects model are not satisfied either.

e.g. including the turbine-specific variables whose influence is captured by the panel fixed effects. On the other hand, the sample sizes for each year are large – more than 30,000 per year – and the slope coefficients are well determined, so the loss of efficiency is not critical.

An important consideration in the original specification was the need to control for variations in wind conditions over time and location. This was achieved by estimating coefficients on monthly dummy variables on the assumption that wind conditions were relatively similar across Denmark. This time it is possible to adopt a better method by using data that is now published by NASA. This consists of monthly estimates of average weather conditions for grid squares of 0.5° latitude and longitude that cover most of the planet. The weather variables include average, maximum and minimum wind speeds at both 10 and 50 metres above ground level. These wind speeds can be adjusted to hub height for each turbine using a standard wind shear equation using average wind shear coefficients calculated from the wind speeds at 10 and 50 metres.[3] The resulting variable differs across turbines within grid squares as well as over time. It can capture some or much of the variation in performance over location and time, reducing the role of fixed effects in the models.

Staffell & Green (2014) use the same NASA data for the UK in an ingenious but ultimately, in my view, misguided manner. They interpolate average hourly wind speeds by grid square and apply a wind shear adjustment to estimate hourly wind speeds at hub height for each wind farm in the UK. Then, they use power curves for different types of turbine to calculate potential output in each month for each wind farm. Finally, they use the ratio of actual to potential output as their indicator of wind farm performance. The difficulty is that each stage in the calculation introduces new errors. In particular, wind speeds are never constant over an hour so the use of a highly nonlinear power curve to calculate potential output may introduce large errors that are correlated with average wind speed and turbine characteristics.[4] The consequence is that their model is vulnerable to the statistical problems associated with measurement error which will, in general, result in biased estimates of the main coefficients in their models.

It is important to be aware that we are dealing with statistical measurements rather than engineering data. Even with access to all of the SCADA data for each turbine in the register, it is not clear that it would be possible to construct a good predictive model of output as a function of some simple indicator of wind speed. Standard assessment methods tend to assume a specific distribution of wind speeds for a site that can be written as a function of 1 or 2 parameters. It is common to use a Weibull distribution or, as a special case, the Rayleigh distribution which is equivalent to the Weibull distribution with a shape parameter equal to 2 – see Carillo et al (2014). Even these models are no more than approximations because the actual best fit distribution depends upon the measurement interval and how variations in wind direction are taken into account. As an illustration, a

3 For the very sophisticated, this procedure is an approximation as the wind shear coefficient may not be constant at different heights above ground level. Still, it is much better than adjustment methods based on ground roughness assumptions that take no account of meteorological conditions.

4 This can be demonstrated by using wind speed measures taken at intervals of five or ten minutes to simulate the calculations using different time intervals. A crucial, but unknown, factor is the way in which the mechanical and blade designs of different models of turbine respond to variations in wind speeds and directions in fluctuating wind conditions. Design features will interact with locational characteristics that influence wind dynamics, so the nature of the errors will be specific to each turbine. This is part of what gives rise to the turbine or wind farm fixed effects in the models.

study carried out by meteorologists in Iceland – Petersen & Olafsson (2016) – highlighted the variability in Weibull parameters fitted to distributions of wind speeds across different monitoring sites in Iceland and pointed out that a unimodal Weibull distribution was quite unsuitable for some sites.

It is a mistake to treat the role of wind speed(s) in the model as the product of an engineering specification.[5] Instead, it – or they – act as proxies or instruments for some larger but unknown set of variables which determine the output from a turbine given its age and other characteristics.[6] In statistical terms, what matters is that the instruments are correlated – preferably, highly correlated – with the unknown regressors and uncorrelated with the error term in the model. Both conditions are clearly satisfied for our data, so that the within-year estimates of the relationship between age and performance should be unbiased.

The procedure used in this study is based on NASA's daily average wind speeds by grid square adjusted to hub height for each turbine using the average wind shear coefficient estimated from the average wind speeds at 10 and 50 metres above surface level. In addition, the proportions of time that the wind speed is less than 4 m/s and above 12 m/s are estimated assuming the within-day distribution of wind speeds is a modified Rayleigh distribution (see Appendix A) with a minimum equal to the minimum wind speed for each day. These daily values are then averaged to give monthly values to match the monthly output data. Since the three variables – average wind speed, % of time that the wind speed at hub height is less than or equal to 4 m/s (PWH-LOW), and % of time that the wind speed at hub height is greater than equal to 12 m/s (PWH-HIGH) – are inevitably highly correlated, the best statistical performance was achieved by using the log of average wind speed and the percentage of time with low wind speeds (PWH-LOW).

Table B1 shows the results of using different methods to estimate the base model with and without year as a regressor and with turbine size effects in addition to year. In each case, the dependent variable is the log of load factor excluding cases in which the reported output is zero or missing. The model is estimated using 666,753 monthly observations covering the period 2002-19 for 3,753 onshore wind turbines at 1,544 wind farms. As one would expect, many of the coefficients are well-determined with small standard errors. The Ordinary Least Square (OLS) coefficient of age on performance using the full dataset implies an age profile that declines at 0.6% per year, while the declines in the panel model with population averaged effects (XT-PA) is very similar at 0.7% per year.

5 There is another reason why hourly wind speeds based on satellite measurements using synthetic aperture radar (SAR) should not be relied upon to generate engineering estimates. Satellite remote sensing is extremely useful as a way of measuring meteorological data at locations for which direct measurement is too difficult or expensive. However, the range of uncertainty in measured wind speeds can be substantial – up to ±2 m/s at high resolution. There are also various sources of bias in the measurements – see Barthelmie & Pryor (2003) and Pryor et al (2004). Better instruments and more frequent measurements have reduced the range of uncertainty over time. This implies that the measurement errors in the model vary systematically with time, thus biasing any coefficient estimates for time and age. Any errors are much smaller when values are averaged over a month and a larger area. The estimation procedure based on pooling within-year equations is not affected by these measurement errors.

6 To avoid confusion, wind speed is not being used as what econometricians call an instrumental variable. These are important when a key variable – e.g. ability in estimates of the return to education - is not measured and is correlated with other variables such as the length of schooling, since more able students are likely to spend more time in education. Monthly wind speed has a low correlation with turbine age, because more windy sites were developed first. Omitting it may cause age to pick up a part of the influence of wind speed, thus biasing the coefficient on age. In addition, including monthly wind levels should improve the statistical performance of the model and may reduce the standard errors of the coefficients.

Adding year to the regressors pushes up the OLS estimate of age decline slightly but this is offset by the small positive coefficient on year. In contrast, there is a large increase in the rate of age decline for the panel model but the difference between the age and year coefficients is very close to the original coefficient on age when year is omitted. The coefficients on the two wind speed variables and on hub height in the OLS estimates are well determined and not affected by the inclusion or omission of year as a regressor. However, in the panel models, the coefficients on both hub height and longitude change substantially – even reversing sign for hub height – when year is included.

The reason for the wide differences in estimates of the coefficients on age and other variables in the panel models is clear – multicollinearity between year and the other variables. From 2010 onwards there was a clear shift in the overall composition of the turbine fleet with most new turbines being 2+ MW models with an average hub height of 80+ metres and located in the western half of Denmark. A standard test for the effect of multicollinearity in OLS is to calculate the Variance Inflation Factor (VIF). A conservative threshold for concern about multicollinearity is a VIF of 3, or less conservatively 5 or above. In the case of the models with both age and year, the VIFs for both age and year were over 5. If the model is estimated separately for each turbine size category, the VIF values are well above 5 for small and medium turbines, but the VIF is below 2 for large turbines.

The results of using OLS to estimate the standard model with both age and year separately for each size category are shown in the last 3 columns of Table B1. There is no evidence for either an age decline or technical progress for the smallest size category of < 1 MW. The medium category is clearly an intermediate one but with coefficients on age and year that are similar to those for the category of large turbines. For turbines of 2+ MW there is strong evidence for an age decline of about 2.1% per year, partially offset by disembodied technical progress at a rate of 1.2% per year. A rate of disembodied technical progress of about 1.2% per year is consistent with earlier evidence for rates of disembodied technical progress.[7]

There is, of course, an alternative way of interpreting the estimates derived from this data. This is that (a) there is, for practical purposes, no disembodied technical progress for wind turbines, and/or (b) if there is such technical progress, it cannot be separated from the effects of ageing. In that case, there is no reason to include time/year in any of the statistical models and all that we can observe is the net effect of ageing after allowing for efforts to offset its impact on performance. It follows that the dashed lines in Figure B1 are all we can expect to observe, and any upward shift in the solid lines is a consequence of changes in turbine specifications rather than technical progress.

In a sense, the choice between the alternative interpretations is a matter of philosophy rather than something that can be readily resolved by appeal to empirical evidence. Learning by doing, one of the classic explanations of disembodied technical progress, has been documented in many studies. Equally we know that engineers and operators do their best through O&M practices and by upgrading components to improve the performance of existing plant and machinery. But such learning and other practices are not universal, so that the performance of a population of wind turbines reflects an average over what may be very diverse approaches to managing the stock of installed equipment.

If the "no technical progress" interpretation is preferred, there is still a statistically significant decline in turbine performance with age. The rates of age decline vary from 0.5% per year for small turbines to 0.8% per year for medium turbines to 1.6% per year for large turbines. The hypothesis of no difference between the coefficients across size categories is decisively rejected.

One of the reasons for studying the effect of ageing on turbine performance is to examine the widespread view that wind generation is a "plant and play" technology that you can build and let operate for 25 years with low operating costs. This view underpins the attraction of building and selling new wind farms to financial investors, who expect a (relatively) steady stream of revenue for 20+ years. There can be no doubt that such a view is over-simplified at best and sometimes seriously wrong. Wind turbines are complicated – and not particularly robust – bits of machinery. In the search for bond-like investments in place of property, it is essential to understand the technical risks involved as well as the extent to which these can be either managed or offset. From such a perspective there is much to be learned from viewing ageing and technical progress as separate processes even if it is their combined effect that matters to a financial investor.

Figure B1 – Illustrative data on turbine performance vs turbine age

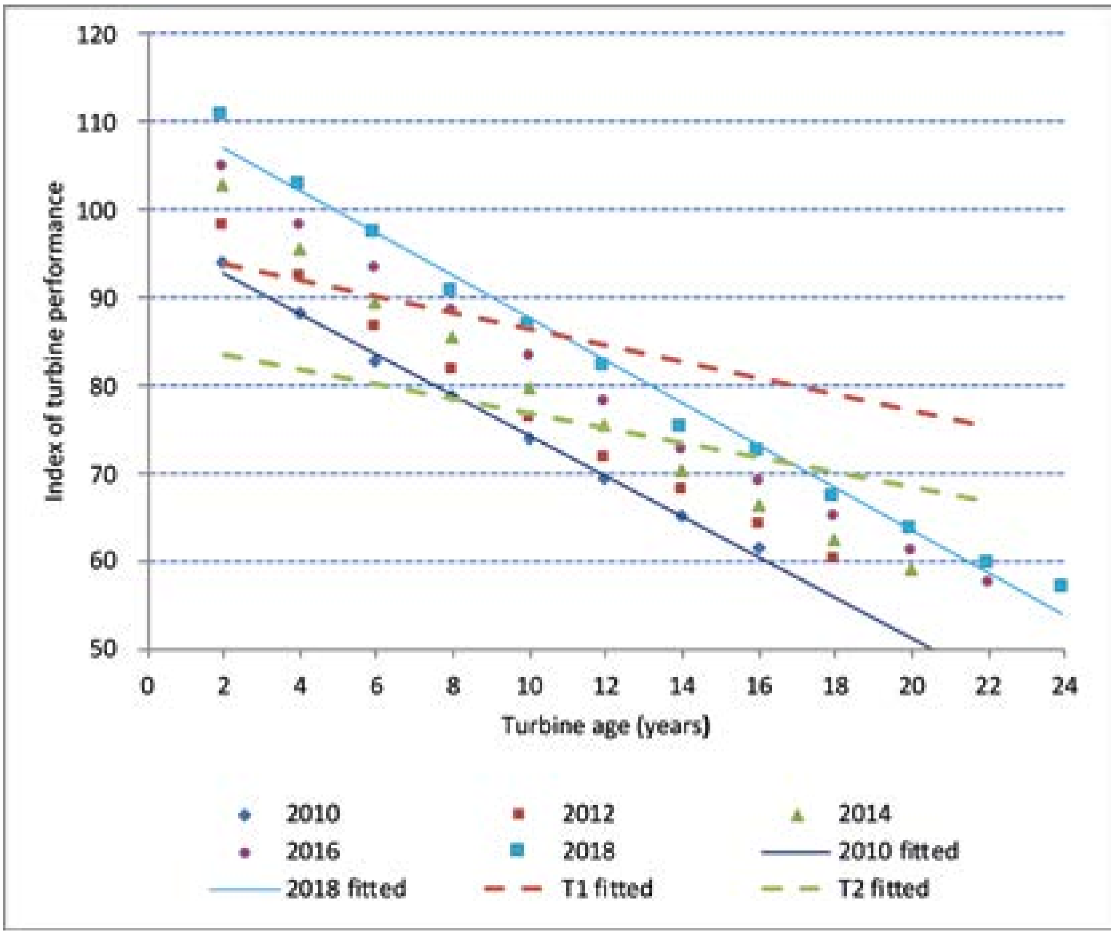

Table B1 – Regression models to identify technical progress

	Base model without year		*Base model with year*		*OLS model by turbine size*		
	OLS	*XT-PA*	*OLS*	*XT-PA*	*0.5 - 1 MW*	*1 - 2 MW*	*2+ MW*
Turbine age (years)	-0.006***	-0.007***	-0.009**	-0.053***	-0.003	-0.023*	-0.020***
	(0.001)	(0.000)	(0.003)	(0.001)	(0.005)	(0.010)	(0.005)
Year			0.004	0.047***	-0.002	0.016	0.011*
			(0.003)	(0.001)	(0.005)	(0.011)	(0.005)
Ln(Wind speed)	2.412***	2.413***	2.410***	2.405***	2.441***	2.544***	1.899***
	(0.023)	(0.007)	(0.022)	(0.007)	(0.024)	(0.049)	(0.109)
% time WS ≤ 4 m/s	0.755***	0.749***	0.751***	0.719***	0.755***	1.033***	0.418*
	(0.037)	(0.010)	(0.037)	(0.010)	(0.039)	(0.087)	(0.165)
Ln(Hub height)	0.087**	0.064***	0.026	-0.934***	0.013	-0.124	0.168
	(0.027)	(0.016)	(0.061)	(0.034)	(0.078)	(0.118)	(0.169)
Longitude	-0.013**	-0.007*	-0.014**	-0.025***	-0.024***	0.026	0.006
	(0.005)	(0.003)	(0.005)	(0.003)	(0.006)	(0.016)	(0.013)
No of observations	666,751	666,751	666,751	666,751	534,515	74,681	57,555
No of turbines	3,753	3,753	3,753	3,753	2,600	373	780
No of wind farms	1,544	1,544	1,544	1,544	1,174	139	231
R-square	0.538		0.538		0.605	0.548	0.209
Root MSE	0.332		0.331		0.283	0.327	0.606
Notes:	*(a) Standard errors in parentheses. * p < 0.05; ** p < 0.01; *** p < 0.001.*						
	(b) Wind speed: monthly average speed at hub height in m/s.						

Source: Study estimates

APPENDIX C
ASSUMPTIONS FOR THE FINANCIAL MODEL OF KRIEGERS FLAK OFFSHORE WIND FARM

As a way of assessing the potential impact of the estimated decline in wind farm performance due to ageing, I have constructed a basic financial model for the Kriegers Flak wind farm that is currently under construction by Vattenfall. It is located in the Baltic Sea at about 15 km from shore with a maximum sea depth of 25m. It will be connected to a neighbouring German wind farm, EnBW Baltic 2, with a further onward connection to Germany. The financial arrangements for the project are unusual because the AC transmission lines from the wind farm will operate as an interconnector between the eastern part of Denmark and Germany, which requires an onshore converter to deal with the different synchronisation of the Nordic and German grids. As a consequence, the price paid for power from the wind farm excludes transmission and grid costs.

The total capacity of wind farm on completion in 2021 will be 600 MW produced from 72 Siemens–Gamesa 8.3 MW turbines. As a result of a tender managed by the Danish Energy Agency, Vattenfall will receive a fixed offtake price (i.e. excluding the cost of transmission) of €49.90 per MWh for the first 50,000 MWh of generation per MW of capacity. It is expected that the contracted volume will be met in about 11 years, which implies an average load factor of 50%. The contract price is not indexed for inflation and the wind farm will sell its output at market prices after the contracted volume has been met. The average NordPool day-ahead market price for market region DK2 from 2016 to 2019 was €33.9 per MWh at 2018 prices without allowing adjustment for the negative correlation between wind output and market prices. In addition, it is necessary to deduct an estimate of the transmission cost from the market price. For a comparable offshore project, Borssele in the Netherlands, the transmission cost is over €14 per MWh. On this evidence, it is very optimistic to assume that the net market price for power from Kriegers Flak will exceed €25 per MWh at 2018 prices under current market conditions.

Public information suggests that the expected capital cost of Kriegers Flak, excluding the transmission investment, will be €1.1-1.3 billion or €2 million per MW, in the middle of this range.[8] This is well below the cost of comparable German projects in the North Sea. Even the Horns Rev 3 wind farm, completed in 2018, was estimated to have a cost of €2.6 million per MW and that was in

8 Excluding the cost of the associated investment in transmission greatly reduces the headline cost of the project. Announcements made by Energinet (the Danish grid company) and 50Hertz (the German grid company) indicate that the cost of the 220 kV transmission line from Kriegers Flak to Denmark will be in the range €450-500 million, while the cost of the work required to connect the wind farm to the Baltic 2 wind farm and upgrade transmission to Germany will be at least €200-250 million. The overall project will cost at least €3.25 million per MW of capacity at 2016 prices.

shallower water and benefited because much of the transmission infrastructure was already in place. I have used a capital cost of €2.17 million per MW, the top end of the quoted range, but I think that the eventual cost is likely to be at least €2.35 million per MW.[9] This is an illustration of the risks of this project that will be highlighted later.

The financial model assumes that the financing structure is 80% debt and 20% equity. The debt is assumed to carry an average interest rate of 2% in real terms, with a grace period of 2 years on debt repayments, and debt repayments over 10 years to match the expected length of the power purchase contract. These terms are extremely generous for this kind of project and, in effect, assume that Vattenfall, as a state owned company, is acting as guarantor of the debt. Stand-alone project finance would be more expensive and would be structured to require earlier repayments. The cost of equity capital is assumed to be 8% in real terms.[10] Using a lower cost of capital does not alter the conclusions because the cash flow profile is so poor.

The parameters of the model for which there is most uncertainty are the levels of variable and fixed operating costs. UK and US projections of levelised costs agree on a total operating cost of about $20 per kWh at 2018 prices, excluding transmission costs for projects commissioned in the early 2020s.[11] A review undertaken by NREL in 2016 quoted a range of $93 to $158 per kW for O&M expenses with the higher figures for fixed-bottom substructure requiring in-situ maintenance. Allowing for inflation and converting to Euros, the UK/US estimates translate to a total O&M cost of about €78 per kW, which is allocated as €39 per kW for fixed O&M and €39 per kWh for variable O&M.

There is limited experience of the average load factor that offshore wind farms will achieve in Denmark. The ones installed between 2005 and 2014 have had average load factors of about 45%, while two small offshore projects that commenced operation in 2018 achieved an average load factor of close to 50% in 2019. It seems reasonable to use this as a baseline assumption and then to examine how much higher the average load factor would have to be to ensure financial viability for the project. The statistical analysis suggests that the achieved load factor in the first 12 months after

9 Aldersley-Williams et al (2019) use data from audited accounts for project SPVs to estimate capex and opex costs for offshore wind projects in the UK in order to estimate the levelised cost of electricity. Following that approach, I have estimated the discrepancy between forecasts of capex costs and realized capex costs for the offshore wind projects discussed in Hughes et al (2017). The median discrepancy implies a typical cost overrun of 17%, taking the realized cost to €2.35 million per MW.

10 The weighted average cost of capital (WACC) is 3.2% in real terms and I have assumed inflation of 1.5% per year, so the nominal cost of capital is less than 5%, which is regarded as low for a project of this scale and risk. Information extracted from the S&P 2018 Rating Report for Vattenfall – S&P Global (2018) - suggests that the assumptions in the financial model are distinctly generous. For its Swedish distribution business, Vattenfall has a WACC of 5.85% but this is expected to fall in future. The company has an average cost of borrowing of 4.4% and applies an overall hurdle rate of 8% for its return on capital employed, while its return on equity in 2017 was 11.1%. These figures imply that the project WACC might reasonably be 100-200 basis points higher than that used in the model. There is no basis for assuming that the cost of finance should be significantly lower.

11 Aldersley-Williams et al (2019) Table 3 report an average estimate of £37 per MWh at 2012 prices for post-OFTO operating costs derived from the audited accounts of a sizeable sample of offshore wind farms in the UK. By post-OFTO, they are referring to the separation of transmission and wind farm assets so that operating costs include transmission charges. Even after transmission costs are deducted, these costs are nearly 50% higher than the estimates used for Kriegers Flak.

commissioning is at least 20% lower than the average for months 13 to 24. This has been taken into account in modelling gross revenues.

Using these assumptions in the financial model, the present value of the project to the equity investor is a loss of almost €400 million over 25 years on the baseline assumptions about performance, with no decline due to ageing and no associated increase in operating costs. One route to achieving a zero present value would be to increase the assumed load factor but doing that on its own requires a load factor of just over 80% for 25 years. That is physically not attainable given the distribution of wind speeds and the shape of the power curves for current turbines. An alternative is to reduce the level of both fixed and variable O&M costs. To reach a zero present value by that route alone requires a reduction of 85% in O&M costs, which would take these costs well below the equivalent costs for onshore wind farms.

The project value is very sensitive to both the level and the growth of the real market price after the expiry of the initial power purchase contract. Hence, Vattenfall may have based its bid on the assumption that real power prices will increase strongly, not just during the initial 10 years but throughout the life of the project. The project would break even if power market prices in real terms were to increase at an average of 5.5% per year for 25 years. In nominal terms, the assumption implies that the market power price in 2045 would be just over €200 per MWh. It is not clear how this could happen, because market prices in Denmark and Germany for wind and solar power have been strongly influenced by excess generation capacity during periods when wind and solar resources are high. An increase in the ETS carbon price will have a very limited effect on the expected market price for a wind operator.

Overall, the most plausible set of assumptions to achieve a zero present value would be based on a combination of these adjustments. The combination of an average load factor over 25 years of 52%, a 20% reduction in O&M costs and an average increase in market power prices of 4% per year in real terms would yield a very small positive present value. Whether these assumptions are reasonable is a matter for debate, but improvements in the baseline assumptions are required to cover the investor's cost of capital even under the best possible scenario for the performance of the wind farm. This breakeven scenario will be used as the basis for assessing the impact of a decline in the performance of the wind farm over time.

Table C1 provides a summary version of the financial model with intermediate calculations omitted and values expressed in €000 per MW of capacity, so as provide a metric independent of the size of the project. The table shows the Base Breakeven scenario using the combination of assumptions required to achieve a positive present value with no decline in the performance of the wind farm. The present value of the project is a total of €16.5 million. This is, of course, extremely low for a project requiring an initial investment well in excess of €1 billion.

Many wind operators seek to limit the demand on their capital resources and to reduce their exposure to market risks by selling a large stake in projects to financial investors after construction and the initial period of operation required to deal with early problems. In this case, I have assumed that Vattenfall may consider such a buy-in sale under the following arrangement:

- The company retains a 25% stake in the project, usually required by the outside investor to align the interests of operator and owners.

- A large portion of the buy-in price is allocated to reduce the debt used to finance the project. This is necessary to adjust the cash flow profile from the project. Without an equity-for-debt swap the net cash flow after debt service is strongly negative for at least 10 years, which is unlikely to be accepted by the financial investor.

- The financial investor will apply a cost of capital of 6% in valuing its buy-in investment. This is lower than the developer's cost of equity because the risks of construction and initial operation have been removed.

Table C1 – Summary version of financial model for Kriegers Flak

(Values per MW of capacity)

Year	*Annual output*	*Price*	*Variable O&M*	*Net revenue*	*Fixed costs*	*EBITDA*	*Depreciation*	*Taxable profit*	*Corp tax*	*Post-tax profit*	*Project Cash flow*	*Accum cash flow*
Unit	*MWh*	*€/MWh*	*€/MWh*	*€000*	*€000*	*€000*	*€000*	*€000*	*€000*	*€000*	*€000*	*€000*
-1											-433.2	
0	3,644	49.9	7.2	155.6	31.2	124.4	144.4	-54.7	0.0	-54.7	89.7	89.7
1	4,555	49.2	7.2	191.1	31.2	159.9	144.4	-19.1	0.0	-19.1	125.3	215.0
2	4,555	48.4	7.2	187.8	31.2	156.6	144.4	-22.4	0.0	-22.4	-51.3	163.7
3	4,555	47.7	7.2	184.6	31.2	153.4	144.4	-22.2	0.0	-22.2	-51.1	112.7
4	4,555	47.0	7.2	181.4	31.2	150.2	144.4	-22.0	0.0	-22.0	-50.8	61.8
5	4,555	46.3	7.2	178.2	31.2	147.0	144.4	-21.7	0.0	-21.7	-50.5	11.3
6	4,555	45.6	7.2	175.1	31.2	143.9	144.4	-21.3	0.0	-21.3	-50.2	-38.9
7	4,555	45.0	7.2	172.0	31.2	140.8	144.4	-20.9	0.0	-20.9	-49.8	-88.7
8	4,555	44.3	7.2	169.0	31.2	137.8	144.4	-20.5	0.0	-20.5	-49.4	-138.1
9	4,555	43.6	7.2	166.0	31.2	134.8	144.4	-20.0	0.0	-20.0	-48.9	-187.0
10	4,555	43.0	7.2	163.1	31.2	131.9	144.4	-19.5	0.0	-19.5	-48.3	-235.3
11	4,555	42.4	7.2	160.2	31.2	129.0	144.4	-18.9	0.0	-18.9	-47.8	-283.1
12	4,555	48.8	7.2	189.7	31.2	158.5	144.4	14.1	3.1	11.0	155.4	-127.7
13	4,555	51.3	7.2	200.8	31.2	169.6	144.4	25.2	5.5	19.6	164.0	36.3
14	4,555	53.8	7.2	212.3	31.2	181.1	144.4	36.7	8.1	28.6	173.0	209.3
15	4,555	56.4	7.2	224.3	31.2	193.1	0.0	193.1	42.5	150.6	150.6	359.9
16	4,555	59.2	7.2	236.7	31.2	205.5	0.0	205.5	45.2	160.3	160.3	520.3
17	4,555	62.0	7.2	249.7	31.2	218.5	0.0	218.5	48.1	170.4	170.4	690.7
18	4,555	65.0	7.2	263.2	31.2	232.0	0.0	232.0	51.0	181.0	181.0	871.7
19	4,555	68.1	7.2	277.2	31.2	246.0	0.0	246.0	54.1	191.9	191.9	1,063.6
20	4,555	71.3	7.2	291.8	31.2	260.6	0.0	260.6	57.3	203.3	203.3	1,266.9
21	4,555	74.6	7.2	307.0	31.2	275.8	0.0	275.8	60.7	215.1	215.1	1,482.0
22	4,555	78.1	7.2	322.8	31.2	291.6	0.0	291.6	64.1	227.4	227.4	1,709.4
23	4,555	81.7	7.2	339.2	31.2	308.0	0.0	308.0	67.8	240.2	240.2	1,949.6
24	4,555	85.4	7.2	356.2	31.2	325.0	0.0	325.0	71.5	253.5	253.5	2,203.2
Present Value per MW in €000											27.6	
Total PV for 600 MW project in € million											16.5	

Source: Study estimates

Accepting such a buy-in investment can enhance the overall value of the project for the developer, provided that the risk profile has not deteriorated during the early years of the project. The reason is that a buy-in generates cash early in the life of the project and thus reduces reliance on cash flow in later years which are heavily discounted and subject to greater uncertainty. On the other hand, an early buy-in may also crystallise a loss due to over-optimistic expectations at the outset while limiting the possibility of benefiting fully from changes in market conditions later in the life of the project.

In the financial model shown in Table C1, a buy-in at the end of the second year of operation increases the present value of the project to €47 million on the assumption that an equity-for-debt swap reduces the amount of outstanding debt by 37.5% to a 50:50 debt-equity ratio. This means that the investor pays €501 million for a 75% stake in the project, of which €390 million is allocated to a reduction in debt. Hence, Vattenfall receives the remaining €111 million and retains a 25% stake in the project. This kind of financial engineering is routine but it only works when the outcomes for risky variables are favourable.

Summary financial models of the kind shown in Table C1, which rely on annual certainty-equivalent values, provide a way of identifying the key parameters that affect the prospect of earning an adequate return on a project. However, they provide no insight into the variability of the return and the primary factors which determine such variability. It is instructive to go into more detail by taking account of the uncertainty about, for example, the amount of generation and the post-contract market price. Hence, I have constructed a Monte Carlo version of the model as a basis for carrying out a risk analysis focusing on three sets of assumptions: (a) average wind speeds and power generation; (b) market power prices; and (c) O&M costs.

The model takes account of two types of uncertainty and risk. The first type is what can be called **structural uncertainty** associated with the inherent variability of, in particular, wind speeds, power yields, market prices, and maintenance costs. Such uncertainty cannot be eliminated though it can be transferred via various insurance arrangements. However, insurance is costly – at least for a risk-neutral and well-capitalised investor – so the model allows for such uncertainty.

The second type of risk, which I will refer to as **forecast uncertainty**, arises from uncertainty about technical change and future market conditions. For example, an investor may believe that it can utilise turbines capable of reliably delivering a 20% higher power yield than current turbines for the same distribution of wind speeds. That is not an insurable risk, though the manufacturer of the turbines may be willing to share the risk either by offering performance guarantees or by taking an equity stake in the venture.

Structural uncertainty is based on public information. While some bidders for a project may collect better data or carry out more sophisticated analyses, there is only a limited amount they can do to manage the risks of weather and markets. They can hedge risks by investing in a portfolio designed to yield a satisfactory aggregate return. All such options are open to other investors and do not substantially differentiate the bidders.

In contrast, bidders may adopt very different positions when considering the forecast risks. Technically there is no private information as no one can know what the market price for power will be in 2030 or how a new generation of turbines will perform. The logic of the winner's curse applies: the winner of an auction with public information will be the company which takes the most “optimistic” view of forecast risks, i.e. the company which bases its bid on extreme values of the distributions of forecast risks.

Following that logic, we may define the range of plausible values for forecast risk items as being defined by (i) the value based on current information – i.e. the measured load factor for similar turbines after controlling for the distribution of wind speeds – and (ii) the value necessary to yield a positive present value for the project. On occasions, companies may realise they are bidding on basis that is almost certain to lose money, but this is rare. It is more usual for companies to convince themselves that a bid based on an extreme value for a forecast risk has at least a 50% chance of breaking even.

Wind speeds and load factors. The Kriegers Flak site is not especially windy. The NASA data implies average monthly wind speeds at hub height (107m) varying from 7.3 m/s in July to 10.2 m/s in January. The monthly coefficients of variation – the ratio of the monthly standard deviation to the monthly average – are 10-16%. Using a statistical model of the monthly load factor as a function of the average monthly wind speed calibrated with data for offshore turbines in Denmark of at least 7 MW, the expected load factor in the first year of full operation – i.e. after the initial proving period - for Kriegers Flak would be about 43%. This is well below the assumed value of 50% that is consistent with the expected time to use up the PPA contract volume, let alone the value of 52% in the breakeven version of the model.

For the forecast risk analysis, it has been assumed that the realised load factor from the project will be in the range from 0.95 to 1.25 times the stochastic predictions based on the historic monthly wind speeds at hub height, which gives a range from 41% to 54% for the achieved load factor.

Market prices. The assumptions about market power prices are based on day-ahead NordPool prices for region DK2 in € per MWh from 2011 to 2019, adjusted to 2018 prices using a standard Euro price index. On a monthly basis the average real price varies in the range €35–€41 per MWh but with large standard deviations. In daily data, the market power price is highly correlated with the amount of renewable (wind and solar) generation but demand factors are the primary influence on monthly prices. Hence, the monthly average prices are treated as normal variates with mean and standard deviations based on historical values. Allowing for a transmission cost of €12 per MWh, the stochastic model implies that the net real power price at the wind farm weighted by output will be about €26 per MWh, well below the breakeven level, if there is no long run trend in real power prices.

The critical question, then, is whether real power prices will increase in future and, if so, by how much. Because of the structure of the power purchase contract, the present value of the Kriegers Flak project is highly leveraged to the level and growth of market power prices at the end of the contract. If the NordPool market power price were to grow at 4% per year for the next 25 years, the project could have a positive present value, provided that other factors are favourable. Such an increase will mean that the real market price after 12 years will be 60% higher than now and the nominal price will be 90% higher. This cannot be ruled out but it would represent a fundamental departure from the pattern of the last decade.

For the forecast risk analysis it has been assumed that the real power price for the DK2 region will increase at a rate between -2% and 5% per year over the life of the project.

Wind discount. In all countries in NW Europe with significant amounts of offshore wind capacity – the UK, Germany, Denmark, Belgium and the Netherlands – there is a negative correlation between the expected volume of offshore wind generation and the day-ahead market price for power after controlling for other factors that influence prices. The effect of this negative correlation is to reduce the average market price received by offshore wind generators relative to the load-weighted average market price. The size of this **wind discount** varies across countries and over time but it tends to increase as the share of offshore wind in total generation increases. In NW Europe the wind discount is at least 5% and is over 10% in Germany. It is likely to increase as the current reliance on gas for marginal generation decreases; this is inevitable because there is no incentive to invest in new gas plants under the current market regime in Germany and countries linked to the German market.

For the forecast risk analysis it has been assumed that the wind discount for the DK2 region will be between 5% and 15% over the life of the project.

O&M costs. On the question of O&M costs, some commentators claim that wind farm operators are adopting practices that will lead to substantial reductions in the variable and fixed O&M costs incurred for new offshore wind farms. In the case of Kriegers Flak the wind farm may also incur lower O&M costs because it is located in relatively shallow water with a depth of 16-25m and it is only 15 km from land. Both of these factors should reduce the capital cost of constructing the wind farm but again there is little evidence of a reduction in O&M costs that is economically significant. Lowering the O&M costs by 10% in the risk model increases the present value of the project by about €45 million. This impact is larger the greater the output produced by the wind farm because variable O&M costs (per MWh of output) account for 40-50% of total O&M costs. Since there is no evidence as to how large the reduction in fixed and variable O&M costs for this project might be, I have made an assumption that reflects a general but informed view that O&M costs can be reduced.

For the forecast risk analysis it has been assumed that O&M costs will be between 0% and 20% lower than the standard estimates of O&M costs for offshore wind made for the UK and the US.

These assumptions, especially when taken together, reflect a wide range of the forecast risk which will determine the realised value of the project. An operator keen to develop its offshore business and has other projects in the region might regard the favourable assumptions as defensible, whereas those of a more sceptical inclination might believe that the outcomes may be closer to assumptions that reflect current performance and market conditions.

Figure C1 shows the results of the structural risk analysis which focuses on quantifiable variability in wind speeds, market prices, etc.[12] The results incorporate the forecast assumptions used to generate a positive present value for the project, as discussed above, and assume that the buy-in price

12 The Monte Carlo analysis is based on 5,000 simulations using Latin Hypercube Sampling and an identical seed for the random number generator. The software uses a pseudo-random number generator based on the Mersenne Twister algorithm.

is a fixed €480 million. The graph shows the percentile distributions of the project present value without a buy-in (solid blue line) and with a buy-in (dashed red line). On this basis, the chance of a negative project present value is less than 5% and the mean present value is €35 million, with a 20% chance of exceeding €50 million. The buy-in present value with a fixed buy-in price transfers risk from the developer to the buy-in investor. The mean project value falls to €32 million while the standard deviation of the project value falls from €17.7 million to €9.5 million. The probability of making a loss is removed but the return at the top end of the distribution is also reduced.

Figure C1 – Structural risk analysis: percentile distribution of present values

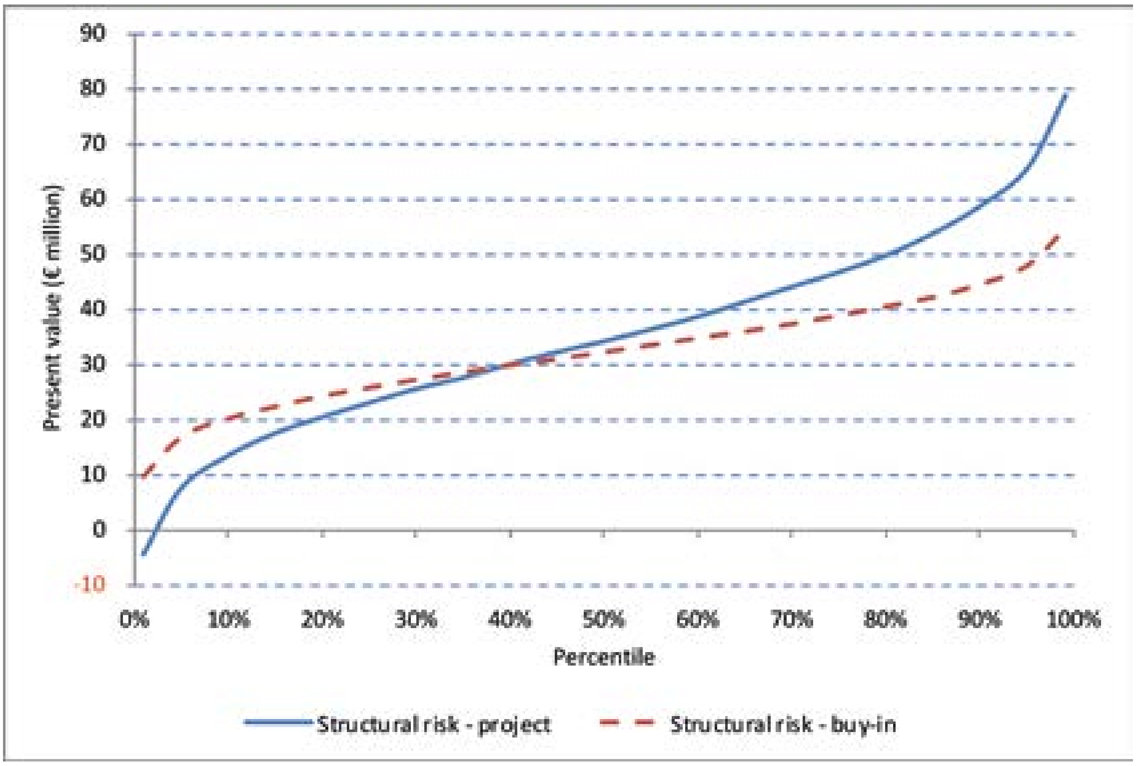

Source: Study estimates

The picture changes completely once we take account of forecast risk. Figure C2 shows the percentile distributions for the net present value of the project with structural risk alone and with structural plus forecast risk. The project values are consistently negative when forecast risk is added to the analysis. There is little or no chance that a buy-in investor would be willing to acquire a stake in the project at a price that would be acceptable to the developer, so the buy-in valuation isn't of much interest. The figure illustrates the importance of the winner's curse for this auction as the median present value of the project is -€381 million when forecast risk is taken into account as compared with +€34 million when the optimistic assumptions of the winner are adopted.

Put in a different way, the auction bid price that would yield the same expected present value as the structural risk model after allowing for forecast risk, i.e. for the winner's curse, is €76 per MWh rather than the winning auction price of €49.9 per MWh. That, of course, highlights the problem for any bidder who recognises the winner's curse. If it had bid something close to €76 per MWh it is likely that it would have lost the auction to another bidder who had made no or little allowance for the winner's curse.

Figure C2 – Structural + forecast risk analysis: percentile distribution of present values

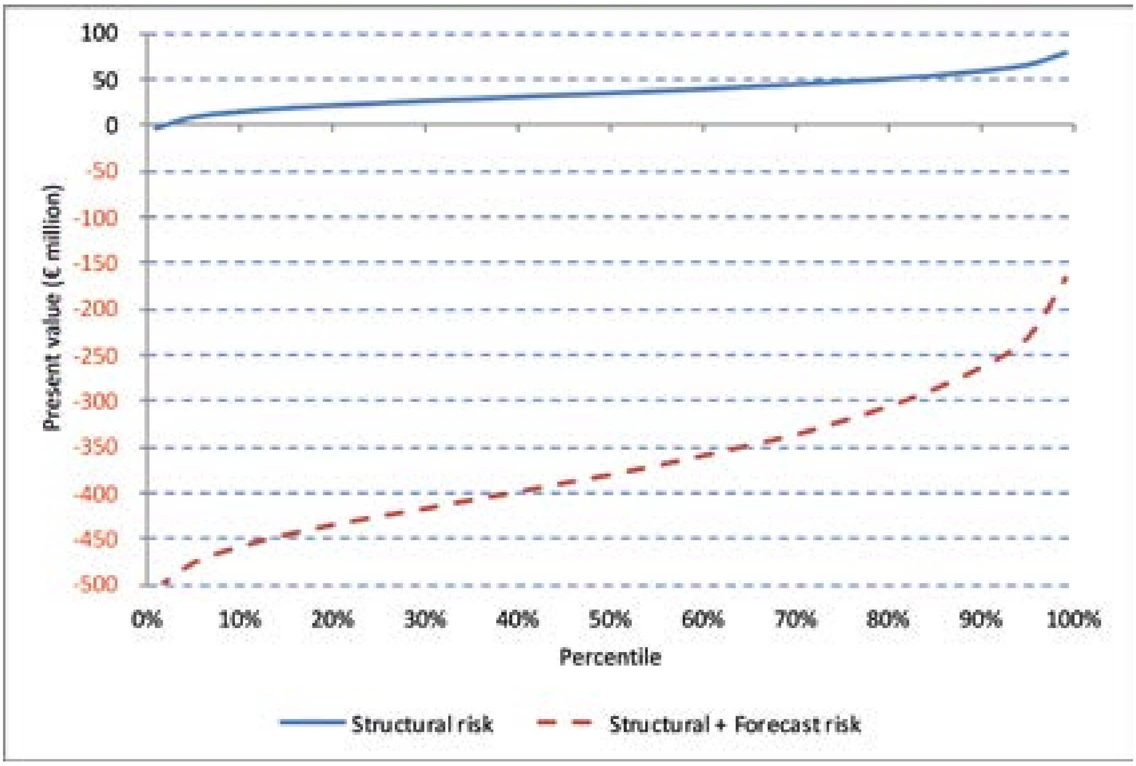

Source: Study estimates