How large should a portfolio of wind farms be?

Richard Green and Iain Staffell[1]

Imperial Business School, Imperial College London

South Kensington, London SW7 2AZ, +44 (0)207 594 2611

r.green@imperial.ac.uk i.staffell@imperial.ac.uk

Abstract

A portfolio of energy generators is likely to out-perform a single source in terms of a better trade-off between the level and stability of profits. We model the financial performance of portfolios of wind farms located around Great Britain in the early 2020s. We measure the expected annual profits and their variance over eighteen years of historic demand and weather conditions, as these measures are likely to be most relevant to investors.

We find the portfolios on the efficient frontier contain relatively few wind farms (no more than four across the country), implying there is no barrier for small investors. The optimal portfolio out-performs the average by a factor of three in terms of return-to-risk ratio, and there is poor correlation between the annual energy output and annual profits of a portfolio, implying that careful market analysis is needed if investors are to build optimal portfolios of wind farms.

[1] Research support from the Engineering and Physical Science Research Council, via Project EP/I031707/1, *Transforming the Top and Tail of the Electricity Networks*, and from the Alan Howard Charitable Trust is gratefully acknowledged. We thank participants at the 2012 IAEE European Conference in Venice, the 2012 BIEE Conference and the 2013 Toulouse Conference on Energy Markets for helpful comments.

1 Introduction

A portfolio of energy sources is likely to give better results, in terms of the trade-off between cost or profit and its variability, than relying on a single source (Awerbuch, 2000; Roques *et al*, 2006). Dispersing wind farms over a wide area can also reduce the impact of variations in wind speed and hence the intermittency of output (Nørgaard and Holttinen, 2005; Sinden, 2007; Roques *et al*, 2010). Hour-to-hour variations in wind output are critical for system operation, but are unlikely to have a significant impact on investors when measured over financially relevant timescales, such as a year. However, there can be significant year-to-year variations in wind conditions, which would have an impact on profitability, and these may differ between regions. There is also a systematic tendency for wind farms to receive prices below the time- or demand-weighted average electricity price, because the hours in which they generate are the hours in which their output depresses the price. In this context, a wind farm sited away from the bulk of a country's capacity, which therefore has different operational patterns, may receive a better average price. These are benefits from siting some stations away from the main area of wind generation, but they could be negated if this implies choosing a site with a lower average wind speed. We wish to model the trade-off to be made by choosing a less windy site and sacrificing some output, if it gives output patterns that differ from the bulk of farms, and thus reduces the inter-annual deviation in profits across a portfolio of farms.

This paper estimates the mean and variance of annual profits for portfolios of wind stations located around Great Britain, using an electricity market model calibrated to the 2020s. We do not study the operational problems caused by intermittency, but take into account the trade-offs between price and output discussed above. We are therefore seeking a set of portfolios that are optimal for the investor, rather than from the perspective of a system planner (who would want to take account of the externalities caused by intermittency).

We take 18 years of hourly wind speed and electricity demand data covering the period from 1994 to 2011, and scale these to anticipated 2020 levels. The wind speed data, for 730 sites around Britain and its waters, are used to predict the output of the wind fleet that is expected to be operating in 2020, totalling 26 GW of installed capacity. Thirty of these farms are then singled out, one in each of National Grid's transmission charging zones, and their individual results are used to assemble portfolios.

We use a merit order stack model to calculate the wholesale price of electricity in every hour, assuming that the stations with the cheapest full-load running costs are always able to meet the pattern of demand. This price is received by all the wind stations in our sample. We calculate annual profits per kW of capacity at each location, before taking the average annual profit per kW for portfolios of between one and 11 farms. Eighteen years of data for each portfolio allowed us to calculate the mean and standard deviation of these annual profits.

The highest annual profits are received by a "portfolio" of a single wind farm, in the region with the highest average wind speed, but this also had the highest standard deviation of annual profits. Generators can reduce the variance of their annual profits by investing in a small portfolio of farms. The lowest variance came from a portfolio of just two widely spaced farms (in the Scottish Hebrides and the south coast of England), and the most profitable 'portfolio' consisted of just a single farm in Scotland. The other portfolios on the frontier giving the best trade-offs between risk and return were also surprisingly small – never more than four stations.

We calculate the efficiency of every possible portfolio, based on the distance between that portfolio and the (unattainable) optimum point that combines the highest average profits and lowest variance. The efficiency of a portfolio is equal to the ratio of the distance from that optimal point to the frontier, relative to its distance to the portfolio, which is analogous to Farrell's measurement of productive efficiency. We find that the efficiency of a portfolio is positively correlated with its size, but very weakly so. Furthermore, the average efficiency is just 0.339, implying a well-chosen portfolio of 2–4 farms will yield almost a three-times better trade-off between risk and return.

A developer may not want (or be able) to build a full market model in order to predict the profit advantages of a diversified portfolio of wind farms. Data on the mean and variance of outputs are much easier to obtain, but we find they give a poor estimate of financial performance. We found a correlation of just 0.149 between a portfolio's efficiency with respect to profits and with respect to output. Seven portfolios were on the efficient frontier with respect to revenues, and six with respect to output, with none being efficient on both measures. Furthermore, the optimal portfolios with respect to output had an efficiency of just 0.440 with respect to profits. In other words, investment decisions should take account of the interactions between patterns of output and market prices across a portfolio.

The next section of the paper describes the background to this study and some relevant previous work. Section 3 sets out our model, while section 4 describes the data we have used for demand, wind generation and the costs of conventional plant. Results are given in section 5, and conclusions in section 6.

2 Background

The UK is one of a number of European countries that has installed large amounts of wind generation over the last decade, and is expected to continue into the next. The problems that the intermittency of wind generation can cause, and the need for back-up plant, are well-known, as are the potential benefits from evening out this intermittency by dispersing the stations over a large area, reducing the impact of any one weather pattern. Nørgaard and Holttinen (2005) and Sinden (2007) model the potential output from wind farms dispersed around Scandinavia and Great Britain respectively, showing that this can significantly reduce the variability of output. However, Oswald *et al* (2008) point out that some of the coldest winter weather coincides with high-pressure systems that produce very little wind across large areas of North-West Europe.

The impact of wind output on power prices has also become well-known. The so-called “merit order effect” (Sensfuß *et al*, 2008) means that prices are lower when wind output is high. Twomey and Neuhoff (2010) point out that this means that wind stations will tend to receive less than the time-weighted price for their output, except to the extent that average wind speeds are positively correlated with average prices. Green and Vasilakos (2010) simulated this effect when they used predictions for wind output in a market model to simulate the price distributions that might be expected if Great Britain built 30 GW of wind stations by 2020, showing that it could have a noticeable impact on wind generators’ revenues, particularly if conventional generators were able to exploit market power.

The other branch of research that we draw on is that of portfolio theory, grouping assets together to achieve the desired trade-off between risk and return. Markowitz (1952) showed that the combination of two assets with returns that were not perfectly correlated could achieve a lower variance than either asset in isolation, and this insight has been applied in many other fields. Awerbuch (2000) was the first to apply it to energy economics, showing that adding renewables to a portfolio of conventional power stations with uncertain fuel prices could allow a given level of risk to be achieved for a lower expected generation cost, even if the renewable sources were more expensive on average than the fossil-fuelled stations. Several other applications are contained in Bazilian and Roques (2008), while Madlener (2012) provides a more recent survey. Roques *et al* (2006) make the distinction between costs and profits, pointing out that the latter can be affected by the correlation between fuel and electricity prices. For renewable stations, where this correlation is low, the socially beneficial reduction in the variance of generation costs may lead to an increase in the variation of generators’ profits that the latter would seek to avoid. Green (2008) used a supply function model to estimate electricity prices and hence profits for different types of thermal plants under varying fuel and carbon prices, showing that nuclear stations would face greater risks in the presence of carbon prices that were correlated with those of fossil fuels. Delarue *et al* (2011) show that it is important to consider the expected operating pattern of each kind of plant when building a portfolio (and that this will depend on the capacity mix chosen), for the optimal portfolios constructed while taking this into account can differ significantly from those built around assumed load factors. Lynch *et al* (2013) use a sophisticated dispatch model to derive the prices and operating patterns behind their portfolio analysis, taking account of start constraints and no-load costs.

Doherty *et al* (2006) model the role of wind in a future Irish power system and find that it can help to reduce both the average level of generating costs and their volatility. Their stations are dispersed around the system to reduce intermittency, but the paper does not suggest that this was done via a formal optimisation process. In contrast, Roques *et al* (2010) apply portfolio theory to consider the optimal siting of wind farms across five European countries, treating the average load factor as the equivalent of the return to a portfolio, and the hour-to-hour change in output as its volatility. They constructed optimal portfolios for the year as a whole, and for peak hours.

Rombauts *et al* (2011) consider the impact of transmission constraints on the efficient frontiers that can be created from seven sites across three countries, illustrating their approach with a relatively short sample of wind data from the Netherlands. The absence of transmission constraints allows each country to choose a somewhat less even distribution of wind power across its sites than would be optimal if there was no cross-border transmission capacity. Grothe and Schnieders (2011) concentrate on maximising the lower quantiles of wind output rather than minimising its volatility, finding that a different distribution of wind farms across Germany could lead to a 150% increase in the proportion of wind capacity that was online with a given probability.

All of these papers focus on output risk rather than matters directly of concern to investors. Dunlop (2004) addresses profitability and proposes that the correlation between the volatility of a wind farm's quarterly output and that of all farms (i.e. its beta) can be used as a measure of undiversifiable risk to guide investors. He points out that for a highly indebted wind farm, small variations in output can lead to large variations in free cash flow. Gómez-Quiles and Gil (2012) build optimal portfolios of Illinois wind farms, considering annual variations in output but taking prices as fixed. This captures the problem facing investors in a system using Feed-in Tariffs (FiTs), but the European Commission (2014) requires that from 2016 onwards, renewable generators must sell their output in the wholesale market and can only receive support in the form of a premium to the market price.

Obersteiner (2010) shows that the market value of wind power decreases as its market share rises, and that portfolios of weakly correlated farms will have a higher market value than those with a high correlation. Schmidt *et al* (2013) use an econometric model to capture the impact of Austrian wind output on electricity prices and hence to calculate the revenue-maximising allocation of wind stations, which are paid market prices plus a premium. They find that dispersing stations raises their average revenue by weakening the impact of the merit order effect on revenues and the optimal portfolio has a lower variability of wind output and residual load than that constructed under a fixed-price FiT scheme.

To the best of our knowledge, the present paper is the first to use mean-variance portfolio techniques to assess wind farm profits when these depend upon the interaction of output and market prices. We show below that output is poorly correlated to profit, and that explicit modelling of station dispatch and electricity prices is needed to assess the attractiveness of a portfolio to investors.

3 Modelling

Our modelling approach is to simulate hourly electricity prices from 18 years of electricity demand and wind data, and to calculate the profits that each of a sample of wind farms would have earned in each year. We then combine those farms into portfolios and calculate the mean and standard deviation of the portfolio's annual profits, following the standard approach of portfolio theory.

3.1 Electricity Market Model

To calculate the annual revenues for each wind generator, we use a short-run electricity market model which simulates the mix of power stations that will be available in 2020 and how they will operate from hour to hour (Green and Staffell, 2013).

Within the electricity market model, thermal power stations are dispatched to meet the demand for electricity, net of the output of the wind stations. This demand is price-sensitive, with an assumed constant slope of minus 10 MW per £/MWh to model demand shifting by large industrial consumers. This gives an elasticity of around –0.1 for high levels of demand, which rises as demand falls.

Our model is built around a merit order stack in which the cheapest available stations are assumed to be physically capable of meeting demand; in other words, we ignore dynamic operating constraints such as start-up times and ramp rates. In each hour, the price of electricity is set equal to its marginal cost, which equals the fuel, carbon, and other variable operating costs of the most expensive station needed. When one group of power stations is running at full available capacity, the price rises to the level at which demand is equal to that capacity. In some cases, the price will then be high enough for the next group of stations to start running at their own marginal cost.

When the net demand is particularly low, generally because wind output is high at times of relatively low gross demand, it may fall below the minimum stable generation of nuclear power stations. We assume this to be 90% of their available capacity, based on historic UK operation. At these times, it is necessary to constrain off some wind stations. These stations would lose output-based subsidies, and would therefore require a payment equal to the subsidy before they are willing to spill output. We assume that the market price is therefore equal to –£50 per MWh at times when onshore wind is spilled, and –£100/MWh for offshore.

Our assumption for the levels of capacity installed in 2020 are taken from the National Grid (2013a) 'Gone Green' scenario. We do not attempt to calculate an investment equilibrium for either wind turbines or conventional plants, although we find that most types of power station make approximately normal profits, implying that the capacity mix suggested in Gone Green is close to market equilibrium, given the assumed fuel prices.

We differentiate between winter and summer in terms of the availability of conventional power stations and the price of gas. In winter, when some gas has to be taken out of storage, its price is 6% above the base level, whereas in summer, the price is 6% below – based on the average differential observed in the UK in the 2000s. There is little planned maintenance in the winter, and so we assume that 90% of the capacity of all types of thermal power station is available, while in the summer, availability falls to 80% as scheduled maintenance takes place. We do not adjust the output of wind stations for maintenance, implicitly assuming that this happens during low-wind periods.

3.2 Wind Farm Output Model

The half-hourly output of individual wind farms is simulated using the *Virtual Wind Farm* model, which is described and demonstrated in Staffell and Green (2014b). As summarised in Figure 1, this takes hourly wind speeds data from NASA, interpolates them to the location and height of the wind farm's turbines, and then converts to power outputs using information on the model of turbine installed.

Figure 1: Overview of the Virtual Wind Farm model used to produce hourly power output time series.

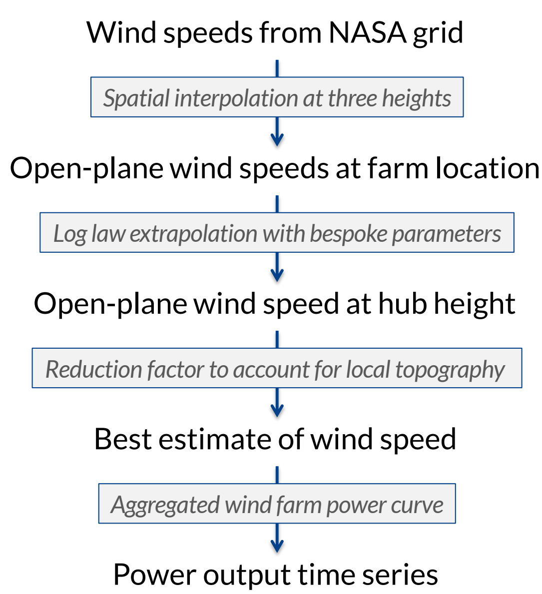

The NASA wind speeds are provided on a regular grid of ½° latitude and ⅔° longitude (approx. 55 by 44 km) at three heights above ground. By having wind speeds at multiple heights, we are able to use the logarithmic wind profile law with site-specific parameters to extrapolate wind speeds to the turbine's hub height, which improves accuracy over using generic formulae (Staffell and Green, 2014b). These wind speeds are representative for smooth terrain, and do not account for the local environment surrounding the farm. Nearby obstructions such as trees, buildings and neighbouring wind turbines will reduce the wind speeds actually experienced. A scale factor is determined for existing wind farms so that modelled average load factors match those actually achieved by existing wind farms; based on historic output data from Ofgem's ROC Register.[2] For onshore farms, the average scale factor was 0.69, while for offshore farms it was 0.85 (reflecting the lower number of obstructions at sea). For farms not yet built, the scale factor was randomly drawn from the historic distributions. The resulting wind speeds are then combined with the power curve for the model of turbine used at each farm (or that is proposed for farms which have not yet been built) to calculate the expected energy output from the farm.

While this process can only give an estimate of the wind speed at any given site, we found it validates well against metered output data reported to Elexon. At present, 46 of Britain's largest wind farms report their half-hourly output to the Balancing Mechanism Reporting System (BMRS).[3] When simulating the aggregate output of these farms at hourly resolution, as in

Figure 2, the R^2 between simulation and reported output was over 0.95, and the root mean square (RMS) error was 233 MW. In other words, any instantaneous estimate of power output has an uncertainty of ± 233 MW, or 4.5% of the capacity which reports to Elexon.

Figure 2: Comparison of simulated hourly wind output to metered data from the UK, showing a three month sample from 2012 (left) and the correlation over 2012–13 (right).

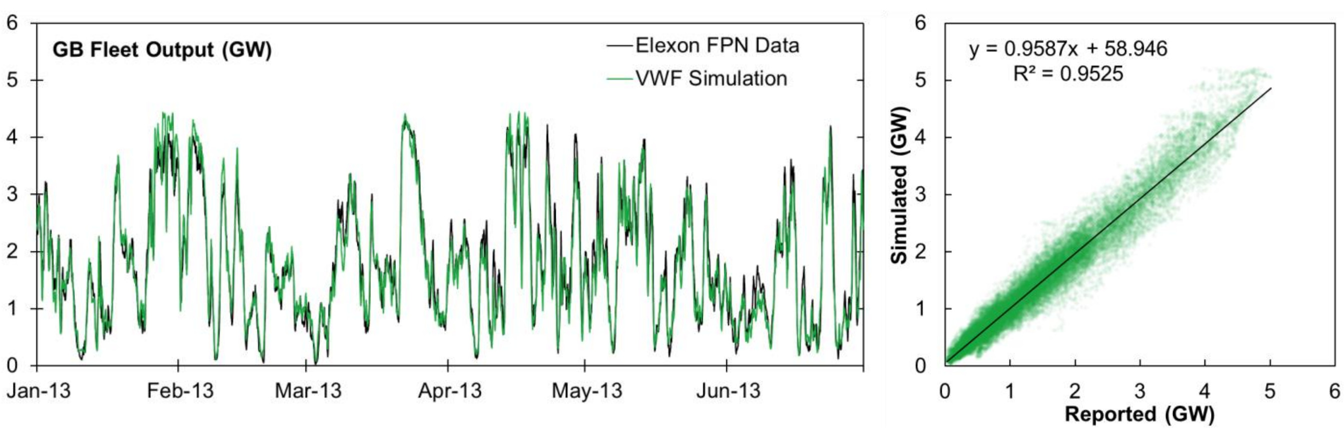

[2] https://www.renewablesandchp.ofgem.gov.uk/

[3] http://www.bmreports.com/

4 Data

The electricity market model relies on four sets of data: the available capacities and costs of thermal generation technologies, and hourly time-series of net national electricity demand and load factors for individual wind farms across the country.

4.1 Installed Capacity

We consider ten types of thermal generation, which are listed in Table 1. Capacities for each technology are taken from the 'Gone Green' scenario presented by National Grid (2013a), which is designed to meet the UK's environmental target of 15% renewable energy (across all sectors) by 2020.

Nuclear capacity is slightly below current levels due to planned retirements. There is no new coal, as environmental policy in Britain prevents unabated coal from being built and no utility-scale CCS plants are expected to be online by 2020. Most new capacity is therefore CCGT and wind, the latter expanding to 26 GW. The system also contains 2.75 GW of hydro which is allocated in multiple tranches using a peak-shaving algorithm (Staffell and Green, 2014a), and 3 GW of interconnection which is inserted into the stack at a position which reflects its historic utilisation.

Table 1: Assumed parameters for Britain's power stations in 2020.

	Capacity (GW)	Thermal Efficiency (%)	Variable O&M (£/MWh)	Marginal Cost (£/MWh)
Wind (offshore)	12.2	–	0.00	–100.00
Wind (onshore)	13.8	–	0.00	–50.00
Nuclear	9.0	35%	2.90	5.00
Coal (best)	4.0	42%	2.13	55.08
Coal (mid)	5.2	40%	2.13	57.72
Coal (worst)	6.4	38%	2.13	60.65
CCGT (best)	6.2	59%	1.39	48.99
CCGT (mid)	5.1	57%	1.39	50.66
	5.1	55%	1.39	52.45
CCGT (worst)	7.1	53%	1.39	54.38
	7.1	51%	1.39	56.46
OCGT	1.2	35%	1.84	152.58

4.2 Station Costs

Capital and operating costs were taken from Green and Staffell (2013), which in turn is based on reports commissioned by the UK's Department of Energy & Climate Change (DECC, 2013), and published by the International Energy Agency (IEA, 2010).

Fuel costs are based on DECC's central scenario for 2020, which sees coal rise by 30% to £11.10 per MWh and gas rise 16% to £25.18 per MWh. Carbon emissions are priced at £32.67 per tonne of CO_2, which is the floor price established for 2020 under the government's carbon price support scheme (HM Treasury, 2011). Variable O&M costs include the costs associated with nuclear decommissioning and waste disposal (approximately £2/MWh). The marginal costs of generation that result from combining these fuel and carbon costs with station efficiencies are given in Table 1, and agree with those presented in the 'Gone Green' scenario. The resulting marginal cost curve is shown schematically in Figure 3.

Figure 3: Marginal cost curve for generation used in the model.

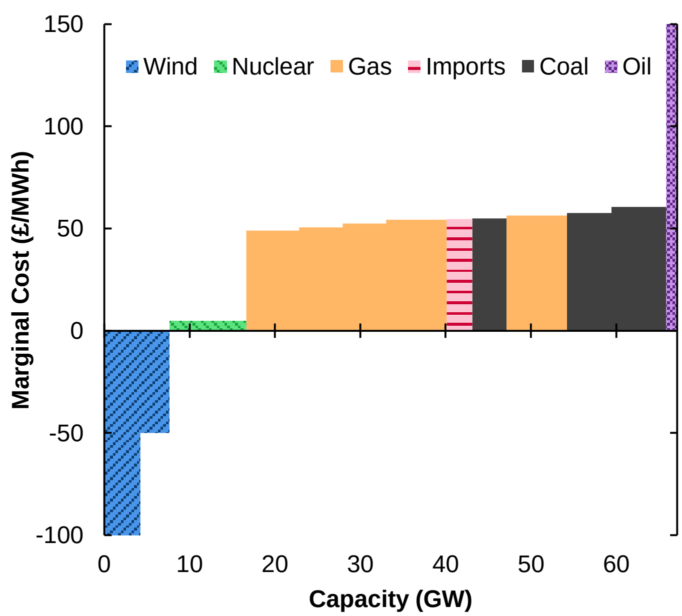

Fixed costs for onshore and offshore wind farms were set to £208 and £371 per kW-year based on DECC (2013). These were not differentiated by location (e.g. to account for differences in terrain or ease of access) as our focus is on the inter-annual variation and correlation between farms rather than the precise level of profits.

Transmission costs vary across the country, and between onshore and offshore wind farms. For onshore farms in Scotland we have taken the Transmission Use of System (TUoS) charges from National

Grid (2013b), which range from £30.25 to £11.07 per kW-year. Onshore farms in England and Wales are connected to the distribution system and do not pay TUoS charges. Instead, distribution-connected generators receive a small payment from the local distribution network operator for each unit of energy generated, reflecting the saving in network costs (under current conditions) from being able to import less power from the transmission system. We used the rates for intermittent generators connected at High Voltage.

Offshore generators have to pay the appropriate transmission charge for the zone in which they bring their power ashore, together with individual charges for substations and offshore cables. The first few stations (which are close to shore but small) pay between £31 and £61 per kW per year to cover these costs (National Grid, 2013b). A study for the Crown Estate (Renewable UK, 2012) has predicted that a 500 MW wind farm (large enough to gain some economies of scale) would have to pay a sum equivalent to £54 per kW per year if it was 50 km from shore, rising to £124/kW if it were 150 km from land. One-fifth of this cost is for the cable, and we scale up this component for Dogger Bank, which will be 220 km from shore, giving it a farm-specific charge of £148 per year. Near-shore stations pay £40 per kW per year, based on the revenue requirement for the Thames Array (Ofgem, 2013).

4.3 National Electricity Demand Profiles

We take the hourly pattern of electricity demand from eighteen years of data published by the National Grid for the period 1994–2011. We follow National Grid (2013a) and DECC (2012) projections for demand levels in 2020, assuming that the level of demand will continue falling by 0.5% per year to 318 TWh in 2020. This total includes approximately 10% of gross demand which will be met by distribution-connected wind, domestic solar photovoltaics and other on-site generators. National Grid's figures are based on transmission-connected generation, and therefore exclude this smaller-scale generation and the demand it meets, whereas we include both.

The annual weather-corrected demands are scaled to a common level, which means that our hourly observations preserve any correlations between the weather and electricity demand. We scale the hourly demand values for each year using a linear scale factor for each year, and so do not reflect changing patterns in the underlying demand due to de-industrialisation or lifestyle changes. We also leave out the potential electrification of heating and transport demands, as these are not expected to be significant factors in the near term. The system peak and minimum demands therefore also scale linearly to 64.6 and 22.7 GW respectively.

4.4 National and Individual Wind Farm Generation Profiles

Two sets of wind profiles were required for this study: the total output of the national fleet that will be installed by 2020, and the individual outputs of the 30 farms which we assemble into portfolios.

For the national profile, we follow National Grid's 'Gone Green' scenario in assuming there is 26 GW of wind capacity installed by 2020, 13 GW onshore and 13 GW offshore. To model this output, we consider farms that are either operating, under construction or have obtained planning permission, giving the 730 farms which are mapped in Figure 4. The location, hub height and turbine model of these farms was taken from The Wind Power and UKWED databases,[4] so that output from each farm could be simulated using the Virtual Wind Farm model.

From these 730 farms, thirty were simulated individually for assembly into portfolios. One farm was chosen in each of the National Grid charging zones, which are defined in (National Grid, 2014) and shown in Figure 4, so that each farm pays a different transmission charge to reflect its location. The 19 chosen onshore sites are large existing farms; the 11 offshore sites include one existing farm, seven that have been allocated in the Crown Estate's Round 3 and one from the Scottish Territorial Waters.

Figure 4: Map of Great Britain divided into the thirty regions used in this study, showing the location of all current and planned wind farms, and those that we assemble into portfolios.

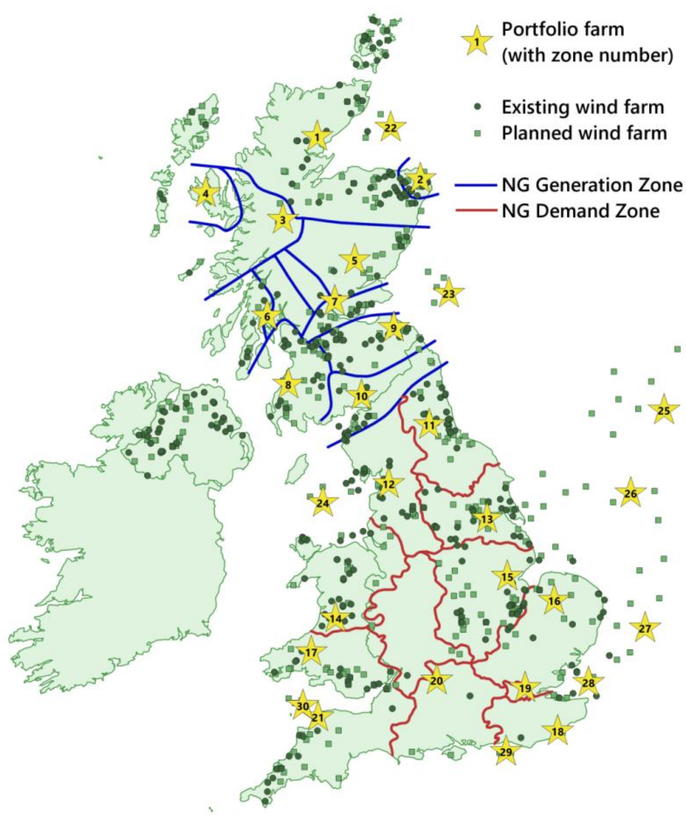

[4] http://www.thewindpower.net/ and http://tinyurl.com/reuk-ukwed

For each farm (the individual 30 and other 700), we simulated the output that would have produced if it had been operating from 1994 until 2011, giving 18 years of load factor data that are coincident with our demand data. The average simulated capacity factor was 24.4% for onshore farms (in line with historic averages) and 34.8% for offshore (which is two percentage points higher than historic experience, due to the anticipated move to larger turbines further from shore).

The distribution of load factors across the 30 portfolio farms is shown in Figure 5. The within-year variation across the country is represented by the spread of each individual bar, and the inter-year variation over the last thirty years by the difference between bars. Two individual farms are highlighted in Figure 5, showing that the variation within a given farm tends to follow that of the national average: a good farm tends to always out-perform an average one. This selection of farms was unbiased, having an average capacity factor of 24.8% onshore and 35.4% offshore.

Figure 5: Spread of simulated load factors for the 30 wind farms in the UK over the last 18 years of wind conditions, with two individual farms singled out.

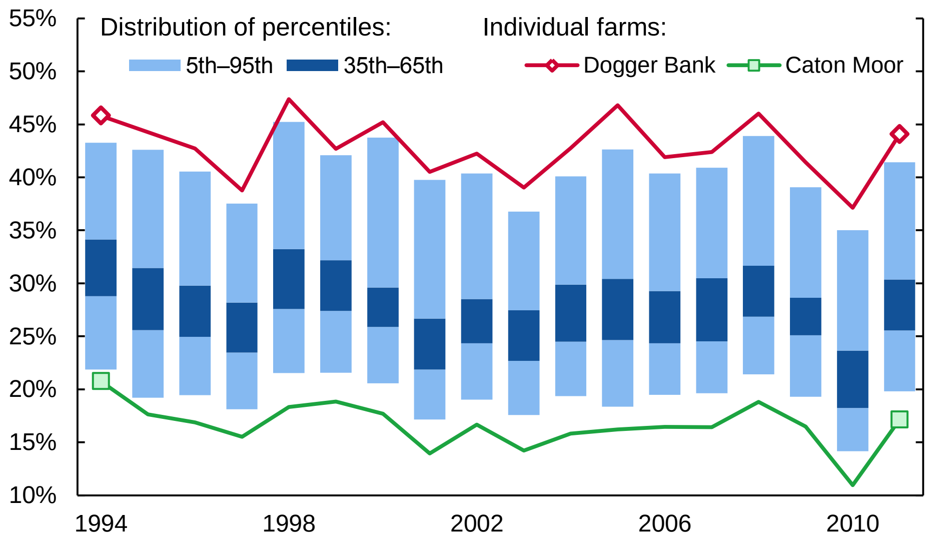

5 Results

The revenue for each wind farm is equal to the value of its output at market prices plus the Renewables Obligation Certificates it earns (valued at an average of £42 per MWh for onshore and double for offshore, and scaled inversely to total national wind output). The annual revenue for our 30 portfolio farms averaged £244/kW, giving an IRR of 12 ± 2% per year. The annual fixed costs were £163–188 per kW for onshore, and £285–380 per kW for offshore farms, depending on the location (and thus TNuoS charges). Given our assumptions, average annual (super-normal) profits were £38/kW-year across all farms, with an average standard deviation (measured across the years for each station individually) of £19/kW-year.[5]

As in previous works (Nørgaard and Holttinen, 2005), we find that the cross-correlation between hourly output from pairs of wind farms falls exponentially with the distance between them, halving with each 440 km of separation. As seen in Figure 6, the correlation between annual outputs falls more slowly, halving every 1,100 km, as this depends on the broader wind resource rather than local phenomena. Profits depend on the broad wind resource and the short-term (and local) effects of prices, and we find that the correlation between annual average profits halves every 650 km. The two measures are themselves almost perfectly correlated at the farm-pair level, with output correlation (Q) being amplified to give profit correlation (π) via $\pi = 1.46\ Q - 0.46$. Investment decisions depend both upon expected profits and their variability, and we find that these are not so closely correlated with the same measures for output.

Figure 6: Correlation between annual profits and annual outputs against distance.

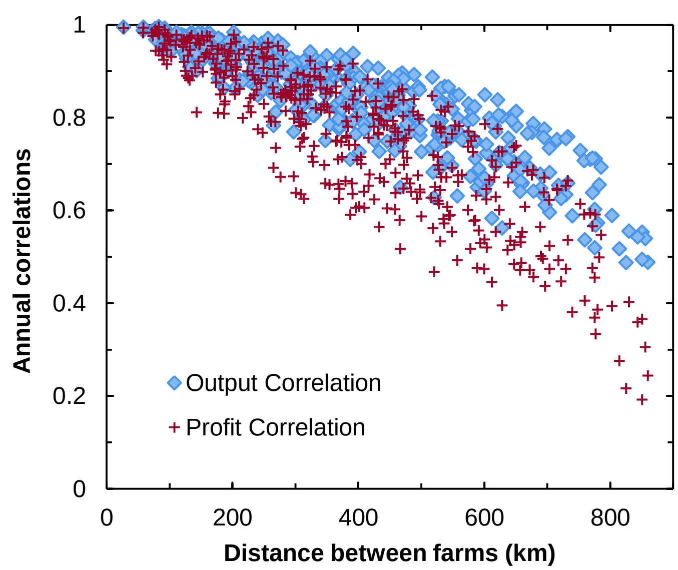

[5] It is possible that a well-informed landowner would be able to capture some of the profits from a windier site in the form of higher rental payments.

With 30 farms there are just over one billion possible portfolios if all sizes are permissible. We limited our study to portfolios of up to nine farms (giving 23 million possible portfolios), finding that larger portfolios revert towards the population's average performance and will not be on the efficient frontier.

Figure 7: Mean and standard deviation of annual profits for individual farms and a selection of portfolios. The efficient frontier is shown as a line, with the means of calculating the efficiency of a given portfolio.

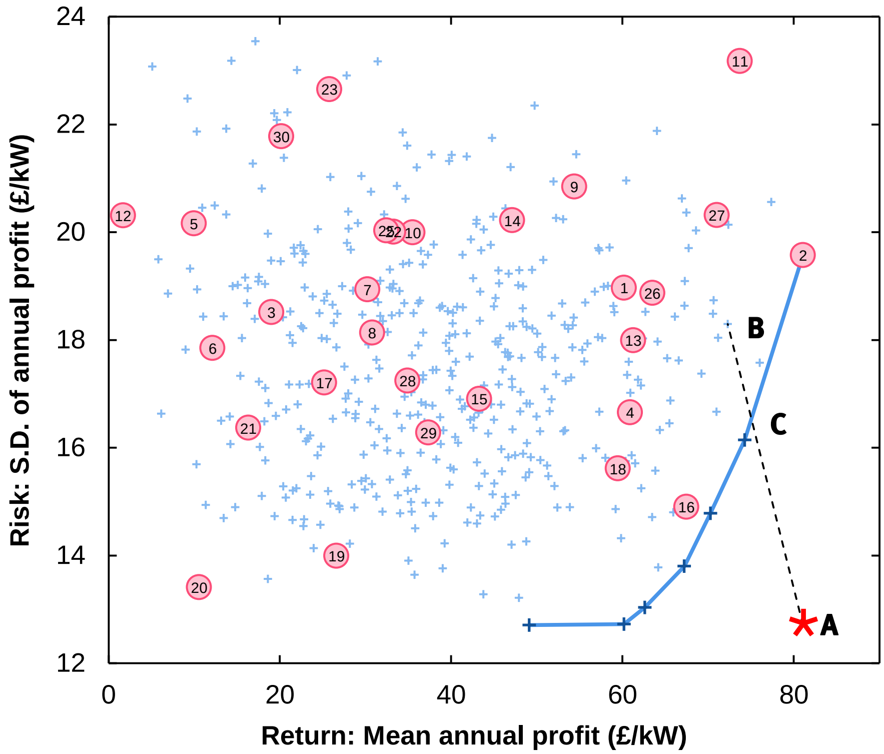

Figure 7 shows the mean and standard deviation of the annual profits for each of our portfolios. The best portfolios are those towards the bottom right of the Figure, showing high expected profits with little variability. Seven points form the efficient frontier, where it is impossible to increase the expected profits of a portfolio without increasing their standard deviation. Most of the portfolios lie above the left-hand half of the efficient frontier, with only a small number offering expected profits of more than £50/kW-year. The portfolios on the frontier are given in Table 2:

Table 2: Optimal portfolios of onshore wind farms

Return: Mean Profit (£/kW-year)	**Risk: Standard deviation (£/kW-year)**	**(2) Aberdeenshire**	**(4) Hebrides**	**(16) Kent**	**(18) East Anglia**	**(29) Rampion**
81	20	×				
74	16	×		×		
70	15	×			×	
67	14	×	×	×	×	
63	13		×	×	×	
60	13		×		×	
49	13		×			×

The most profitable portfolio is that of a single region, Aberdeenshire (Zone 2) in the east of Scotland, which has the highest annual outputs for onshore farms. The Hebrides, the islands to the north-west of Scotland, also have high outputs. The portfolios with less risk include stations from distant parts of Great Britain – Aberdeen and the Hebrides in the north-east and north-west of Scotland, Kent and East Anglia in the south-east and east of England, and the Rampion offshore farm off the southern coast.

We define the efficiency of a portfolio in terms of its closeness to the efficient frontier. We take our underlying concept from the measurement of productive efficiency. The classic presentation of Farrell efficiency graphs firms in terms of their capital-output and labour-output ratios and draws a ray from the origin to the firm's data point. The firm's efficiency is given by the distance along this ray to the efficient frontier, divided by the distance to the firm. In our context, the origin is not a desirable place, and we instead define the point shown by the star (labelled "A") in Figure 7. This shows the combination of the highest observed profits and the lowest standard deviation from any of our portfolios – since these came from two different portfolios, this point is not attainable in practice. For any other point, such as B, its efficiency is assessed along the ray to point A, relative to the point where this ray intersects with the convex hull of the efficient frontier, point C. The efficiency score is then the ratio of AC to AB. It is equal to 1 for a point on the efficient frontier, and falls as the distance between the point being assessed and the frontier increases. The average efficiency score of our portfolios is 0.339 and the minimum score is 0.150.

A developer might not wish to build a full market model to assess the profitability of its proposed farms; furthermore, the exact profits that we predict are sensitive to the details of the model. A question arises: would it be sufficient to assess the prospects for a station on the basis of its expected annual output and its variance? We could draw an efficiency frontier analogous to the one in Figure 7, but using the data for each portfolio's annual outputs. In this case, there are nine portfolios on the efficient frontier, an average efficiency score of 0.571, and a minimum score of 0.354.

The average efficiency scores are higher for output than for profit as there is a smaller range in mean output levels between stations. No portfolios are efficient when measured against both outputs and profits, although two of the profit-optimal portfolios come close, with efficiencies of 0.966 and 0.975 with respect to output. The efficient portfolios based on output are somewhat larger than those based on profits, mostly containing three and four stations – still a small proportion of the thirty being considered.

The correlation between the two efficiency scores is only 0.149. The optimal portfolios with respect to output have an average efficiency with respect to profits of just 0.440, and can be as little as 0.277. Selecting a portfolio that offers a good combination between the average level of annual output and its variability is far from a guarantee of a similar relationship for annual profits. Figure 8 shows the pattern of efficiency scores across all of portfolios of up to four farms.

Figure 8: Portfolio efficiency, measured by profits and by output.

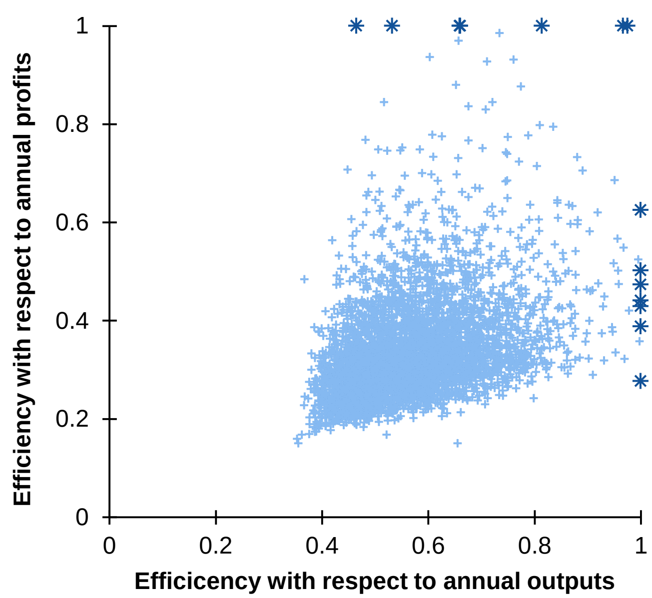

6 Conclusions

We have modelled the expected level and annual variation in the profits from portfolios of onshore wind farms distributed around Great Britain in the 2020s. For investors, this is a more relevant timescale than the hourly variation which has been the subject of previous work in this area, and is of course critical to system operators.

We find that the optimal portfolios consist of stations sited in no more than four regions at a time out of the thirty we study, suggesting that the benefits of diversification can be achieved quite easily on this time scale. The choice of portfolio is important, however, because the average efficiency across all the 23 million portfolios that we assessed was just 0.339, relative to the optimal frontier. This means that the portfolios must be carefully chosen if they are to reduce the variation in annual profits without sacrificing too much expected profitability.

We also calculated the efficiency of every portfolio in terms of the expected level and variability of annual outputs. The mean efficiency was somewhat higher, at 0.571, but we found a very weak relationship between the two measures, with a correlation coefficient of 0.149. Developers should not rely solely on measures of output when designing an optimal portfolio of wind farms, but should consider the interaction between wind output and electricity prices. A small number of high-priced hours (typically on winter days with relatively low wind) are responsible for a significant proportion of each generator's profits, and it will be the correlations between the stations' outputs in these hours, rather than over the year as a whole, that govern the behaviour of their profits.

Our results will be affected by the market design in force. We have modelled the current system of renewable energy support in Great Britain, which gives generators one revenue stream from market prices and a second from selling Renewable Obligation Certificates that depend on the level of output, but not its timing. This kind of premium on the market price is consistent with the new EU guidelines for renewable support described above; a pure Feed-in Tariff, which fixes the price received by the generator and means that profit risk is simply a levered version of output risk, will not be allowed. Another dimension to electricity market design concerns the degree of geographical differentiation in prices – the impact of this on incentives for diversification is a subject for further research. Persistent price differences would make stations in low-priced regions less attractive to generators, but variation around the mean may reinforce the incentive to own a diversified portfolio.

Generators can often get a better trade-off between the expected level and the annual variability of their outputs by building a portfolio of geographically dispersed stations. The interactions between output levels and prices are such that a portfolio designed to optimise the mean and variance of output levels may well prove sub-optimal from the point of view of revenues. Careful analysis of the market would be required for generators that wish to maximise their expected revenues without excessive volatility.

7 References

Awerbuch, S. (2000) "Getting It Right: The real cost impacts of a renewables portfolio standard" *Public Utilities Fortnightly* Feb 15: 44-52

Bazilian, M. and F. Roques (2008) *Analytical Methods for Energy Diversity and Security: A tribute to Shimon Awerbuch*, Oxford: Elsevier

Delarue, E., C. De Jonghe, R. Belmans and W. D'haeseleer (2011) "Applying portfolio theory to the electricity sector: energy versus power" *Energy Economics*, 33(1): 12–23

DECC (2012) "Updated energy and emissions projections – Annex E: Total electricity generation by source."

DECC (2013) "Electricity Generation Costs".

Doherty, R., H. Outhred and M. O'Malley (2006) "Establishing the Role That Wind Generation May Have in Future Generation Portfolios" *IEEE Transactions on Power Systems*, 21(3), pp. 1415-22

Dunlop, J. (2004) "Modern Portfolio Theory Meets Wind Farms" *The Journal of Private Equity*, 7(2): 83-95

European Commission (2014) "Guidelines on State aid for environmental protection and energy 2014-2020 (2014/C 200/01)". Brussels: Commission of the European Communities

Gómez-Quiles, C. and H.A. Gil (2012) "The Value of Wind Resource Geographic Diversity for Wind Farm Profitability" *IEEE Transactions on Power Systems*, 27(4): 2074 - 2083

Green, R.J. (2008) "Carbon Tax or Carbon Permits: The Impact on Generators' Risks," *Energy Journal*, 29(3): 67-89

Green, R.J. and N. Vasilakos (2010) "Market Behaviour with Large Amounts of Intermittent Generation" *Energy Policy*, 38(7): 3211-3220

Green, R.J., I. Staffell and N. Vasilakos (2014) "Divide and conquer? k-Means Clustering of Demand Data Allows Rapid and Accurate Simulations of the British Electricity System." *IEEE Transactions on Engineering Management,* 61(2): 251–260.

Green, R.J. and I. Staffell (2013) "The Impact of Government Interventions on Investment in the GB Electricity Market". Annex 3 of "SA.34947 Electricity Market Reform – Investment Contract (early Contract for Difference) for the Hinkley Point C New Nuclear Power Station". European Commission DG Competition.

Grothe, O. and J. Schnieders (2011) "Spatial dependence in wind and optimal wind power allocation: A copula-based analysis," *Energy Policy*, 39(9): 4742–4754.

HM Treasury (2011) "Carbon price floor consultation: the Government response".

International Energy Agency (2010) "Projected Costs of Generating Electricity". Paris, IEA, NEA, OECD

Kubik, M., D. Brayshaw and P. Coker (2012) "Reanalysis: an improved data set for simulating wind generation?" Presented at *WREF 2012*: Denver, CO.

Lynch, M. Á., A. Shortt, R.S.J. Tol and M.J. O'Malley (2013) "Risk–return incentives in liberalised electricity markets" *Energy Economics*, 40(11): 598–608.

Madlener, R. (2012) "Portfolio Optimization of Power Generation Assets" in (eds) Q. P. Zheng, S. Rebennack, P. M. Pardalos, M.V.F. Pereira and N.A. Iliadis, *Handbook of CO_2 in Power Systems* pp 275-296, Berlin: Springer-Verlag

Markowitz, H.M. (1952) "Portfolio Selection". *The Journal of Finance*, 7(1): 77–91

National Grid (2013a) “UK Future Energy Scenarios”.

National Grid (2013b) “Transmission Charging Statement”.

National Grid (2014) “The Statement of Use of System Charges”.

Nørgaard, P. and H. Holttinen (2004) “A Multi-Turbine Power Curve Approach”. Presented at *Nordic Wind Power Conference*: Gothenburg, Sweden.

Obersteiner, C. (2010) “The Market Value of Wind Power from a Generator’s Perspective – An Analysis of Influencing Parameters”, PhD Thesis, Technischen Universität Wien

Ofgem (2013) “Encouraging Investment in Offshore Wind”.

Oswald, J., M. Raine and H. Ashraf-Ball (2008) “Will British weather provide reliable electricity?” *Energy Policy*, 36(8): 3212-3225

Renewable UK (2012) “Potential for Offshore Transmission Cost Reductions: A report to the Crown Estate”, Renewable UK: London.

Rienecker M.M., M.J. Suarez, R. Gelaro, R. Todling, J. Bacmeister, E. Liu, *et al.* (2011) “MERRA: NASA’s Modern-Era Retrospective Analysis for Research and Applications.” *Journal of Climate*, 24(14): 3624-3648.

Rombauts, Y., E. Delarue and W. D’haeseleer (2011) “Optimal portfolio-theory-based allocation of wind power: Taking into account cross-border transmission-capacity constraints” *Renewable Energy*, 36(9): 2374–2387

Roques, F.A., W. J. Nuttall, D.M. Newbery, R. de Neufville and S. Connors (2006) “Nuclear Power: a Hedge against Uncertain Gas and Carbon Prices?” *The Energy Journal*, 27(4): 1-23

Roques, F., C. Hiroux, and M. Saguan (2010) "Optimal wind power deployment in Europe--A portfolio approach," *Energy Policy*, 38(7): 3245-325

Schmidt, J., G. Lehecka, V. Gass and E. Schmid (2013) “Where the wind blows: Assessing the effect of fixed and premium based feed-in tariffs on the spatial diversification of wind turbines” *Energy Economics*, 40(11): 269-273

Sensfuß, F., M. Ragwitz and M. Genoese (2008) “The merit-order effect: A detailed analysis of the price effect of renewable electricity generation on spot market prices in Germany.” *Energy Policy*, 36(8): 3086-3094.

Sinden, G. (2007) “Characteristics of the UK wind resource: Long-term patterns and relationship to electricity demand”, *Energy Policy*, 35(1): 112-127

Staffell, I. and R.J. Green (2014a) “Is There Still Merit In The Merit Order Stack? The Impact of Dynamic Constraints on Optimal Plant Mix”. *IEEE Transactions on Power Systems (under review).*

Staffell, I. and R.J. Green (2014b) “How does wind farm performance decline with age?” *Renewable Energy*, 66: 775–786.

Twomey, P. and K. Neuhoff (2010) “Wind power and market power in competitive markets” *Energy Policy*, 38(7): 3198-3210

UK Meteorological Office (2006) “MIDAS Land Surface Stations data (1853-current)”. NCAS British Atmospheric Data Centre.