



APPENDIX 5-4

*HOUSE PRICES STUDY - CXC
SCOTLAND 2016*

Impact of wind turbines on house prices in Scotland

Dr Stephan Heblich,¹ Dr Dan Olnert,² Prof Gwilym Pryce² and Prof Chris Timmins³

With research assistance from Dr Ellie Bates⁴ and Dr Tim Birabí²

October 2016



¹ Department of Economics, University of Bristol, member of the ESRC AQMeN project.

² Sheffield Methods Institute, University of Sheffield, member of the ESRC AQMeN project.

³ Department of Economics, Duke University, USA, member of the ESRC AQMeN project.

⁴ School of Law, University of Edinburgh, member of the ESRC AQMeN project.

Summary

This report presents the main findings of a research project estimating the impact on house prices from wind farm developments. It is based on analysis of over 500,000 property sales in Scotland between 1990 and 2014.

The methodology builds on research on the impact from wind farms on house prices in England (Gibbons 2014). This study improves the way the impact is estimated by looking at the impact of both single turbines and whole wind farms.

To control for the normal fluctuations in house prices we used a 'control group' that closely resembles the characteristics of the dwellings in the study but without being exposed to a wind farm. This provides prices that can be used to interpret a wind farm's impact on the price of dwellings nearby. As such a result showing no effect means that the house price of the property with a wind farm close by has increased or decreased at the same rate as the properties in the control group.

The study looked at both natural landscape and built environment in relation to how exposed a dwelling is to the visual impact of the wind farm.

Key findings

1. **No evidence of a consistent negative effect on house prices:** Across a very wide range of analyses, including results that replicate and improve on the approach used by Gibbons (2014), we do not find a consistent negative effect of wind turbines or wind farms when averaging across the entire sample of Scottish wind turbines and their surrounding houses. Most results either show no significant effect on the change in price of properties within 2km or 3km, or find the effect to be positive.
2. **Results vary across areas:** The results vary across different regions of Scotland. Our data do not provide sufficient information to enable us to rigorously measure and test the underlying causes of these differences, which may be interconnected and complex.

Our results persist under a variety of assumptions:

- whether or not we account for the visibility of turbines;
- whether we base the analysis on individual turbines or entire wind farms;
- whether we account for building heights or use only the natural terrain when estimating turbine visibility; and
- whether we follow individual dwellings over time or use postcode averages.

The complexity of the findings may be due to:

- attitudes towards wind farms and their benefits potentially varying across regions and different social and economic groups;
- Scotland having a higher proportion of its turbines located in remote areas; and
- the fact that some wind farms provide economic or leisure benefits (e.g. community funds or increasing access to rural landscapes through providing tracks for cycling, walking or horse riding).

Additionally these factors are not mutually exclusive. It is likely that they affect house prices simultaneously, and to varying degrees in different locations.

Contents

Summary	3
Key findings.....	3
Introduction and background.....	5
Details of the house price impact analysis	7
Overview of the data and method	7
House price data.....	7
Wind turbines.....	7
Landscape and building height data	8
Analysis step 1: Which houses can likely see turbines? ‘Line of sight’ analysis.....	9
Analysis step 2: house price impact using ‘difference in differences’	12
Results.....	14
Result #1: Analysis based on Postcode Averages & Wind Farm Centre Points (‘centroids’) (Gibbons)	14
Result #2: Analysis based on Repeat Sales & Individual Turbines.....	15
Result #3: Analysis based on Repeat Sales & Individual Turbines, Taking into Account Building Heights	17
Summary and possible explanations for the results	19
Heterogeneous and changing preferences.....	19
Location of turbines.....	19
Amenity and economic benefits.....	20
Patterns of social stratification	20
Interactions between multiple causes	22
Appendix: Sensitivity analysis.....	23
Introduction.....	23
Sensitivity analysis for result #1: based on Postcode Averages & Wind Farm Centre-Points (‘centroids’) (Gibbons)....	25
Sensitivity analysis for result #2: based on Repeat Sales & Individual Turbines.....	26
Sensitivity analysis for result #3: based on Repeat Sales & Individual Turbines and Taking into Account Building Heights	27
Acknowledgements	27

Introduction and background

The Scottish Government has committed to a target for renewables to generate the equivalent of 100% of Scotland's electricity demand by 2020⁵. Onshore wind power is playing a central part in decarbonising Scotland's energy supply.

The rapid growth in onshore wind (both in Scotland and globally) has been accompanied by an interest in understanding the impacts of onshore wind development, both positive and negative. The overall economic benefits of investment and spending are relatively straightforward to measure⁶; impacts on communities less so. Survey-based approaches consistently show a majority in favour of renewable power generation in principle but paint a more mixed picture for those directly affected by nearby wind farm development⁷.

There is now a substantial body of research on the local impacts of wind farms. Some of this research has looked at measurable effects on house price in order to understand the objective effects on communities, beyond stated views. Have properties near to, or in sight of, new wind farm developments seen price changes that differ from other houses? Until recently, all extant studies had consistently found no robust evidence of any such price impact. One of the most recent studies, by RenewableUK and the Centre for Economics and Business Research, used seven wind farm case studies across England and Wales, and came to the same conclusion: either no impact or even a slight positive one⁸.

Very shortly after that study, however, Steve Gibbons looked again at English and Welsh wind farms using a larger dataset and property prices between 2000 and 2012, and found evidence for negative price impacts⁹. In Gibbons' analysis of previous house price studies¹⁰, the key problem he identifies is sample size: while some studies contain many properties, the number of observations actually used to estimate the price impact tends to be too low to be statistically reliable. Many also do not compare price changes across time. Gibbons' research design allows for comparison of much larger groups of property prices before and after wind farms became operational, allowing for more robust results.

The present study bases its price impact analysis on Gibbons' approach, including his use of a landscape analysis to determine whether properties can likely see a turbine¹¹, or whether line of sight is blocked. Line of sight analysis allows us to test whether visibility of turbines affects house prices differently to proximity alone, by separating visible and non-visible turbines into two groups. We have also explored ways of improving on Gibbons' approach, greatly increasing the resolution and precision of the data. These improvements are listed below:

1. Whilst we replicate Gibbons' approach using average house price per postcode and postcode-centre for housing location, we also repeat the analysis using individual property prices based on full address locations.
2. We use a dataset of wind turbines that includes their exact location and tip height, rather than the centre-point of wind farms. Relying on the centre-point of wind farms might be particularly problematic in a Scottish context where some wind farms are very spread out. When turbines are dispersed in this way, it is possible for a house to be a very long way from the centre of the wind farm, but very close to a peripheral turbine.
3. Our landscape analysis uses 5 metre grid squares (versus 200 metre in Gibbons). Combined with the exact property locations and turbine locations, this gives much more accurate lines of sight.

⁵ 2020 Routemap For Renewable Energy In Scotland – Update, 2015, <http://www.gov.scot/Resource/0048/00485407.pdf>

⁶ RenewableUK, 'Onshore Wind: Direct and Wider Economic Benefits', 2015, <http://www.renewableuk.com/en/publications/index.cfm/BiGGAR>.

⁷ See e.g. Christopher R. Jones and J. Richard Eiser, 'Understanding "Local" Opposition to Wind Development in the UK: How Big Is a Backyard?', *Energy Policy* 38, no. 6 (2010): 3106–17.

⁸ RenewableUK, 'The Effect of Wind Farms on House Prices', 2014, <http://ruk.pixl8-hosting.co.uk/en/publications/index.cfm/RenewableUK-Cebr-Study-The-effect-of-wind-farms-on-house-prices>.

⁹ Stephen Gibbons, 'Gone with the Wind: Valuing the Visual Impacts of Wind Turbines through House Prices', *Journal of Environmental Economics and Management* 72 (July 2015): 177–96, doi:10.1016/j.jeem.2015.04.006.

¹⁰ Ibid. p.179

¹¹ Why 'likely'? - The real landscape may differ in ways the model has not captured - for example, vegetation may be blocking a view.

4. Taking advantage of this higher resolution, we have also added building height data (where available) to test whether buildings may block a property's view.

The following section describes the data used in more detail, and then explains the two key steps in producing the analysis: the line of sight analysis and the econometric house price analysis. The full results are then presented, before concluding with some possible explanations for the findings.

Details of the house price impact analysis

Overview of the data and method

In this section, we outline the data sources for the project and explain how they were used to produce the house price impact analysis. The following four sub-sections describe the **four sources of data** used:

1. House price data for Scotland from January 1990 to March 2014.
2. Wind turbines that became operational between November 1995 and December 2014.
3. Digital Elevation Model (DEM) data for the Scottish landscape, giving height above sea-level for 5-metre grid squares covering the whole of Scotland.
4. Building height data, added to the DEM data.

We shall then detail the two steps of data preparation and analysis. The first step was to carry out a line of sight analysis identifying which houses could most likely see at least one turbine. This provided full details for each house of the number of visible turbines and their distance. The second step was to use this information, along with property price change over time (and a number of other control variables; see below), to produce the final house price impact analysis.

House price data

Data for property prices in Scotland comes from two previously unlinked versions of price data from Registers of Scotland (RoS). By linking these, the house price record covers just over 23 years (1990 to March 2014). While RoS record every Scottish sale, the analysis here drops any sales that, for a number of reasons, were not suitable. For example, not all properties could be exactly geocoded because the RoS record contained insufficient address information to obtain a location match and had to be excluded.

Only repeat sales (properties that sold more than once within the time period of the data) were used in the house price analysis. Following properties over time in this way helps us to compare like with like when estimating the house price impact of turbines being constructed. One limitation of this repeat sales approach is that we do not know whether there have been major changes to the dwelling over time. However, provided changes to dwellings are fairly randomly distributed across all dwellings in the data, this should not have a big effect on the results. In total, the RoS data provided 637,000 repeat-sale properties, accounting for just over 1.7 million sales.

Following Gibbons, we restricted the properties used in the analysis to those within 15km of at least one turbine (i.e. within the green circles in Figure 2). This is done, as Gibbons says, because "as the distance to the wind farm increases, the number of other potential coincident and confounding factors increases, making any attempt to identify wind farm impacts less credible"¹². This reduces the total number of properties in the analysis to 509,275.

Wind turbines

Three sources have been combined to produce the wind turbine dataset:

1. Precise wind turbine locations were acquired from Ordnance Survey's "Points of interest" (POI) data, freely available through an academic license¹³. Its latest incarnation (as of late 2015) is much more comprehensive than previous versions. This data is collated for Ordnance Survey by PointX (www.pointx.co.uk). The POI turbine data itself is mainly supplied to Ordnance Survey by RenewableUK.
2. Dates that wind farms became operational were 'scraped' from RenewableUK's website (www.renewableuk.com) and then matched to turbines.

¹² Gibbons, 'Gone with the Wind'. p.180

¹³ Code and guidance for extracting specific types of POI data are accessible at the Sheffield Methods Institute github page: github.com/SheffieldMethodsInstitute/windfarmsHousePrices

- Turbine tip height information was collated through direct research of planning applications and other publicly available sources¹⁴.

Figure 1 shows the cumulative rise in the number of turbines becoming operational in Scotland from 1995 onwards; the total reaches just over 2,500 turbines by the end of 2014.




Figure 1: Number of operational wind turbines in Scotland, cumulative from 1995 to 2014

Landscape and building height data

To determine whether a turbine is likely to be viewable from a particular property, we need to know if any landscape features intervene to block the view. This requires using a 3D 'Digital Elevation Model' (DEM) of the Scottish terrain, onto which houses and turbines can be added. We use Ordnance Survey's "OS Terrain 5" DEM, which provides height above sea level for every 5-by-5 metre grid point.

The OS Terrain 5 data can be used to identify which houses have their lines of sight blocked by the physical landscape, but this does not account for the effect of other buildings. To correct for this, we also use building height data for the majority of properties in Scotland, combining Ordnance Survey's Mastermap with LIDAR data from the Centre for Environmental Data Analysis (CEDA). The OS Terrain 5 DEM data's 5 metre resolution is fine enough to allow addition of building footprints and heights derived from the Mastermap and CEDA data.

On the map of Scotland in Figure 2, areas for which we used building data are shown with the yellow (Mastermap) and red (CEDA) grid areas. Where both sources covered the same area, we used the slightly better quality Mastermap data. These two sources do not cover all buildings in Scotland, but because data exists for all the larger conurbations, 84% percent of properties have a line of sight that crosses building height data and so could potentially have that view blocked. Calculations are run both with and without building heights for comparison, with the latter using the 84% subset of houses that may have had a line of sight blocked by a building.

¹⁴ The majority of the work tracking down tip heights was done by Dr Ellie Bates, University of Edinburgh.




Figure 2: Scotland - housing data location (dark blue), turbine 15km radii and building height data location

Analysis step 1: Which houses can likely see turbines? 'Line of sight' analysis

The econometric analysis requires the following information for each repeat-sale property:

- Which turbines, if any, are within 15km?
- How close is each of them to the property?
- Of those turbines within this 15km range, which are visible to this property and which likely cannot be seen?

We used Pythagoras' Theorem to compute distances between each dwelling and turbine. To estimate turbine visibility, we used 'line of sight' analysis (also known as "intervisibility" analysis)¹⁵. Figure 3 and Figure 4 illustrate how this process is carried out using the example of a particular property in Glasgow that has its line of sight blocked by another building. 136 batches of housing, turbine and landscape data are processed - these figures use a batch covering the Cathkin Braes wind turbine, installed in 2013¹⁶. (Other batches process larger groups of turbines together, e.g. the Whitelee wind farm to south of Glasgow in Figure 3 is processed in one batch.)

The dotted line on the map of Glasgow in Figure 3 marks an 8.7km line of sight between this example property and the Cathkin Braes turbine. Figure 4 gives the landscape cross-section for this same line (with horizontal distance at 1/8th scale, relative to height), showing how the DEM landscape data - both with and without building heights - is used. The line starts two metres aboveground level on the site of the house¹⁷ and 'looks' towards the turbine blade tip height. If the highest point of the tip is visible above landscape and buildings, the line of sight is clear. In this example, for landscape alone, the house (left-hand side of graph) has a clear line of sight. If building heights are used, however (green in Figure 4), line of sight is blocked.

This process was repeated for all properties. The addition of building height data blocked a great many more from view of a turbine. Without building heights, 80% of properties within 15km of a turbine are identified as having a line of sight to at least one. This drops to 32% when building heights are used - an unsurprising result given how many properties are located in conurbations. Note that this binary visibility result says nothing about a turbine's actual visual impact which will depend on proximity. For example, a visible turbine will presumably have a much bigger visual impact when viewed from nearby properties compared with the view from houses 15km away. As Gibbons says:

"Existing literature based on fieldwork suggests that large turbines are potentially perceptible up to 20km or more in good visibility conditions, but 10 to 15km is more typical for a casual observer and details of individual turbines are lost by 8km."¹⁸

¹⁵ Code and guidance for this is available at the Sheffield Methods Institute github page: github.com/SheffieldMethodsInstitute/windfarmsHousePrices

¹⁶ See e.g. "£5m city turbine will be visible around world (From Evening Times)." 2013. www.eveningtimes.co.uk/news/13256714.5m_city_turbine_will_be_visible_around_world

¹⁷ The building data for the house is discounted: for the building height check, line of sight is only checked once the line has got past the building's edge.

¹⁸ Gibbons p.180




Figure 3: Digital Elevation Model for Glasgow area. Repeat-sales properties in green. Wind turbines are yellow triangles. Dotted line is an example line of sight (matches figure below) for a sample Glasgow property to Cathkin Braes turbine tip.




Figure 4: example line of sight blocked by buildings that would not be blocked by landscape alone. Matches dotted line in above figure. Property on left, Cathkin Braes turbine tip on right. Note horizontal distance is 1/8 of actual scale, relative to height.

Analysis step 2: house price impact using 'difference in differences'

The aim of the econometric analysis described in this section is to assess the house price impact as distance increases, both for visible and non-visible turbines and wind farms.

We use a "difference in differences" approach to identify the causal effect of wind turbine proximity and visibility. This approach seeks to estimate how rates of change in house prices differ between properties "exposed" to wind turbines (through proximity and/or visibility) compared with those that are not exposed. We use only 'repeat sale' properties, as described above. We label properties exposed to wind turbines - those we want to identify any price impact for - as the "treatment group".

To measure the causal effect of wind turbine exposure, we would ideally like to know how the same dwelling's change in price over time is affected by the presence or absence of a wind farm. Clearly, observing both states at the same time is not possible. Instead, we construct a "control group" that closely resembles the characteristics of the treatment group but has not been exposed to a wind farm. The control group thus provides us with a counterfactual dwelling price, which we interpret as what the price would have been if the treatment group had not been in proximity to, or in sight of, wind turbines. This setup allows us to compare the average change in 'exposed' dwellings' house price to the average change in 'unexposed' dwellings' house price before and after turbines become operational - a so called *difference-in-differences* framework.

The first difference is how much the treatment and control groups change price between the chosen time periods. The second difference is how these two changes compare. This second difference is labelled the "treatment effect", i.e. the causal impact of wind farm developments on house price growth. If we were to produce the same findings as Gibbons, with the treatment group's price increasing **less** than the control group, then the impact of wind turbines on house price growth would be negative. For example, if we find a house price impact of -10%, this means that prices in the treatment group went up by 10% less than they did in the control group. On the other hand, if we find a positive effect, say 10%, this means that prices in the treatment group went up by 10% more than in the control group.

Note that a key assumption in the difference-in-differences framework is that the treatment and control groups show the same trends in house price growth in the pre-treatment period (the 'common trends assumption'), which means that they are subject to the same influences on price before the turbine is installed.

For all results, we repeated our difference-in-differences analysis using a large variety of additional controls that control for possible unobserved factors. This is the same as the "fixed effects" approach used by Gibbons (2014). The essential principle of a fixed effects approach is to allow fixed (i.e. constant over time) differences in subsets of the data to be accounted for. Including fixed effects allows the analysis to control for factors that we cannot easily measure (such as cultural differences or unknown economic, political or physical factors) but are likely to be fairly constant over time and may cause different price trends. The most intuitive fixed effects are regional. For example, there might be different house price trends across NUTS2 regions because of differences in the fixed characteristics across regions, such as their physical geography. These differences can be controlled for using fixed effects even if we do not have detailed data on the different underlying characteristics. This may be important if wind farms are sited taking these features into account.

All of the results presented in this report include basic fixed effects that control for variations in overall house price trends and differences in property characteristics. We use annual and quarterly fixed-effect controls to flexibly account for house price trends. Since we are looking at repeat sales, our estimations further include a set of house fixed effects - allowing each property its own trend line - that absorb any time-invariant house characteristics such as its footprint size or number of bedrooms. These are the "basic" controls used in all the results reported here.

We then add a number of additional controls to the models in order to test sensitivity. First, a number of geographic controls are added, allowing different house price effects over time by including fixed effects for slope (for each individual property), elevation (height above sea level for each property) and aspect (which compass direction the property's slope is facing, indicating which direction their predominant view is likely to be). Second, we add controls for different price effects across distance rings. These controls are in line with the ones used by Gibbons (2014). In addition, we allow house prices to differ between Scotland's four NUTS2 regions and include a set of region-by-year interactions. These additional fixed effects results are provided in the appendices.

Results

We present three sets of results. We start with the Gibbons (2014) approach, which is based on postcode averages for house prices and computes proximity and visibility using the centre point of entire wind farms (rather than individual turbines). We then compare these baseline results with outputs based on more fine-grained analysis that follows individual dwellings over time and calculates turbine proximity and visibility based on individual wind turbines. This is done both for visibility based just on terrain, and also visibility that also accounts for any buildings that may block the view.

Result #1: Analysis based on Postcode Averages & Wind Farm Centre Points ('centroids') (Gibbons)

Figure 5 shows the percentage impact on house price growth of a dwelling close to a wind farm being able to see the wind farm (blue line) compared with not being able to see the wind farm (red line). The approach used to derive this first set of results is similar to Gibbons (2014). They are based on:

- the change in average house prices in a given postcode before and after a wind farm became operational (rather than individual dwellings); and
- the effect of entire wind farms (rather than individual turbines).

Compared to the individual-property-level repeat sales analysis, one may think of this as a repeat sales estimation at the postcode level. However, instead of looking at the same house selling multiple times, we now look at multiple transactions in the same postcode. The implicit assumption is that houses within the same postcode unit are very similar and could be used interchangeably.




Figure 5: Result #1: Percent difference in the change of house price

(Postcode/wind farm centroids¹⁹, **whole wind farm visible/not visible.**)

The horizontal axis shows the distance between the postcode of dwellings and the centre of the wind farm. These distances are grouped into 6 bands: (i) 0-2km, (ii) 2-3km, (iii) 3-4km, (iv) 4-5km, (v) 5-8km, and (vi) 8-14km. The vertical bars show the confidence intervals for each estimate. If the confidence interval is narrow, depicted by a short vertical bar, it means the estimate is precise. The longer the bar, the wider the confidence interval,²⁰ and the less precise the estimate is. If this vertical bar is entirely above zero, it means the result suggests a significant²¹ positive effect on house price change caused by the construction of the wind farm. If the vertical bar lies entirely below zero, it means that the effect is significantly negative. If the vertical bar extends above and below zero, as is the case for most of the estimates in Figure 5, it means that there is no significant effect, either positive or negative. In other words, we cannot rule out a zero effect at the 95% confidence level.

A zero effect does not mean that house price growth has flat-lined. Rather it means that the treatment group (those properties that are in close proximity to a wind turbine) have a similar house price growth trajectory as the control group (those properties that are not in close proximity to a wind turbine).

The results in Figure 5 suggest that visible turbines have a positive effect on house prices (the blue line is above zero for the first four distance bands). However, the majority of confidence intervals extend above and below zero. This suggests that there is no significant house price effect in the first three distance bands, but a possible slight positive effect for visible turbines in the 4-5km band, dropping to a negative effect in the 8-14km band.

As discussed above, we repeated our analysis using a large variety of different specifications that control for a variety of possible unobserved factors using the same “fixed effects” approach used by Gibbons (2014). The results of the key variations from this exercise are presented in Figure A1 in the appendix, where Figure 5 is replicated in Figure A1(A) for comparison. We can see that the results are broadly consistent with Figure 5 in that none of the graphs show significant negative impacts of wind turbines on house price growth in the first three distance bands. Some graphs do, however, suggest a significant **positive** impact on house price growth, particularly in the second distance band (2-3 km), and particularly for visible turbines (see graphs (B), (C), (D), (F), and (H) of Figure A1). A more detailed description of the results in Figure A1 is presented in the Appendix.

Result #2: Analysis based on Repeat Sales & Individual Turbines

Figure 6 shows results based on the repeat sales of individual properties and the impact on house price growth after individual turbines become operational²². Here we see a significant positive impact on house price growth in the first distance band (1-2km) for properties that cannot see any turbines, but this effect is much smaller and statistically insignificant for properties in the same distance band that can see turbines. Note that the positive effect on properties, for which turbines are visible, becomes statistically significant in the second, third and fourth distance bands. The two furthest distance bands, however, do indicate negative price impacts. Though these results are mixed, as confidence intervals for visible/not visible turbines cross or touch the zero line.

Results of the sensitivity analysis—comparison with a variety of different fixed effects—are presented in Figure A2 in the appendix. Again, these different versions of the results tell a similar story with the positive impact on house price growth

¹⁹ Centroid means centre point of an aerial unit (e.g. postcode) or multiple points.

²⁰ Based on the 95% level of confidence, which is the standard threshold used in statistical studies.

²¹ Statistical "significance", in this context, means that there is less than a 5% chance that an estimated negative or positive house price impact is purely due to random variation in the data.

²² Again, this is replicated in the appendix, figure A2(A).

tending to diminish with distance for properties that cannot see turbines, but rising then falling with distance for properties that can see turbines.

Crucially, there are no consistent signs of negative impacts on house price growth in the first three distance bands. In these results, the negative signal in the furthest two bands is again mixed, with no completely consistent pattern either side of zero.

Note that at shorter distances, confidence intervals tend to be larger. This is unsurprising, as sample sizes at shorter distances are smaller (there are not many houses very close to turbines) and so there will necessarily be more uncertainty in our estimates at close distances.




Figure 6: Result #2: Percent difference in the change of house price (All repeat sales, turbine visible / not visible)

Results for individual repeat sales properties (Figure A2, appendix) show much the same pattern, but with larger percentage effects. The larger non-visible turbine effects at very close distance do, again, have large confidence intervals - but these do not cross the zero line. For both the centroid and repeat-sales results, any impact on house price growth tends to drop off as distances increase, though there is a great deal of variability in this response.

Repeat-sales results take advantage of having individual turbine data to distinguish between responses to turbines over and under 100 metres to tip height (appendix, figures A2(E) and A2(F); A3(E) and A3(F)). Sub-100 metre turbines are associated with consistent negative house price impacts, if they can be seen - but, again, confidence intervals cross the zero line. This is not the case for those out of sight, however.

Turbines over 100 metres in height are very similar to the main results - with perhaps a more clear decay of positive effect over distance for non-visible turbines. It is worth noting that: (a) Aberdeenshire has a large proportion of the sub-100 metre turbines and (b) most of the above 100 metre turbines were built after 2006, so this difference in response could be rooted in these different times and places.

Result #3: Analysis based on Repeat Sales & Individual Turbines, Taking into Account Building Heights

One disadvantage with both Result #1 (the Gibbons approach) and Result #2 (the individual houses/turbines approach) is that the visibility estimates do not take into account the possibility of buildings (as opposed to natural features) blocking the line of sight to turbines and wind farms.

Figure 7 shows the results of an analysis based on the repeat sales of individual properties and the impact on house price change after individual turbines become operational taking into account the height of buildings that might block the view of turbines. (Again, the appendix shows the results of the sensitivity analysis for these results in Figure A3). While the main findings remain similar to Results #1 and #2 in that there are no consistent signs of negative house price effects in the first three distance bands, it is clear that the estimates of impacts of visible and non-visible turbines on house price changes appear to be much closer in Result #3. Looking across all the results in Appendix figure A3, for both visible and non-visible turbines, the impact on house price growth seems to be more positive in the second distance band (2-3km) than in the closest distance band (0-2km), but then declines in distance bands three and four. As with the previous result, there appear to be negative price impacts in the last two distance bands, particularly for visible turbines, but these results are less consistent in the sensitivity analysis.




Figure 7: Result #3 Percent difference in the change of house price

(All repeat sales, turbine visible / not visible, using building height data for line-of-sight)

While results using building height data in Figure 7 are broadly similar to those relying on terrain-based line of sight, for some of these regressions there are quite different results even for properties that cannot "see" a turbine. This is because it uses a different sample of houses - only those that have lines of sight that cross areas that have available building height data. If this is not done, it is impossible to know whether a property has a clear line of sight due to no

buildings blocking it, or just that no building height data was available. As mentioned above, this still accounts for 84% of properties - but these are all in the larger conurbations. The properties that "can see" and "cannot see" are, of course, also different. The building height results, then, say more about the impact of wind turbines in urban areas than the non-building height sample.

The main difference in the building height result is in the nearest distance band where the effects on house price growth for properties whose line of sight is blocked by a building are noticeably smaller in comparison to those with line of sight blocked by terrain. With terrain only, visible and non-visible appeared to show a quite different response (Figure 6), but when the building height data are included (Figure 7), the impact of visible and non-visible turbines both have the same direction of change as distance is increased (though again, the wide confidence intervals mean there is considerable uncertainty surrounding the estimates).

The pattern of difference between sub-100-metre turbines (Figure A3(E)) and those over (Figure A3(F)) is similar to the terrain-based results once the uncertainty surrounding estimates is taken into account. For turbines less than 100 metres that can be seen despite building height, there appear to be large impacts on the price growth of properties in close proximity, and these impacts diminish at further distances, but the confidence intervals are so wide, we cannot be sure that the effects are different to zero for any of the distance bands, visible or non-visible. Much more precise results are available for turbines over 100m with statistically significant positive effects for the second distance band (2-3km) in Figure A3(F).

Summary and possible explanations for the results

In summary, we have not found any consistent evidence of a negative impact of wind turbines on house price growth. Generally speaking the effect is either positive at particular distance bands (2-3km) or not distinguishable from zero.

Note again that a zero effect does not mean that house price growth has flat-lined. Rather, it means that there is no significant difference between the house price growth of the treatment group (properties close to turbines) and that of the control group (properties far away from turbines).

A positive effect means that the treatment group has a higher rate of house price growth than the control group. The repeat sales analysis, for example, finds a positive effect of 2% for houses in the 2-3km distance band that can see a turbine (Figure 6). This means that the value of those houses went up by 2% more than the increase in value of dwellings in the control group.

We also find some evidence that the impact of wind turbines on house price growth appears to vary across different regions of Scotland. This finding has not, as far as we are aware, been systematically tested in previous UK studies using the rigorous methods applied here.

There is some evidence from the results that property prices respond differently to wind turbines in different parts of Scotland. It must be emphasised, this finding is somewhat tentative. Using the current method, sample sizes are too small to be fully reliable. However, it does suggest that while some areas see the positive impacts described above, others may see negative impacts.

Results for Angus/Dundee and Clackmannanshire/Fife regions, all clustered north of the Firth of Forth, appear to see some negative impacts for visible turbines, though most of these have confidence intervals crossing or just touching zero. In contrast, North and South Lanarkshire show the most positive price impacts at close distances. Other regions either produce no geographical results due to data limitations, or are very mixed.

Our data do not provide sufficient information to enable us to rigorously measure and test the underlying causes of these differences which may be interconnected and complex. Differential impacts may arise, for example, from interactions between variations in physical terrain, urban social structures, local approaches to turbine development policy and community engagement.

We now conclude the report by offering a number of possible explanations for our findings.

Heterogeneous and changing preferences

The reason our results are consistently different to those reported by Gibbons (2014) might be because attitudes towards wind farms may be different in Scotland than in other parts of the UK, and may also vary significantly within Scotland, and between individuals. Attitudes may also have varied over time – e.g. in response to public debates about energy futures or rural economic development. So our complex findings may reflect genuine complexity and fluidity in the preferences and attitudes of homeowners across Scotland over the time period considered.

Location of turbines

In Scotland, a much higher proportion of turbines are likely to be located on moors and mountains, and in much more remote areas than in England and Wales. These differences in terrain might be another important reason for the discrepancies between our results and those of Gibbons (2014), as might the potential alternative uses of the land on which turbines are constructed. For example, in remote mountain locations, there may be fewer alternative commercially viable uses for the land and so the opportunity cost in terms of foregone alternative revenue streams from the land may be smaller. In contrast, high quality farmland locations in England and Wales may well have more valuable

alternative uses that have to be foregone, both now and in the future, if turbines are constructed. This may itself affect the attitudes of, and financial impact on, local residents and businesses.

Amenity and economic benefits

The positive house price impacts presented above may also reflect the fact that some wind farms provide economic and leisure benefits to the surrounding areas.

- E.g.1: The Whitelee wind farm had 25,000 visitors in the first two years of opening²³ and provides 130kms of tracks for walkers, cyclists, horse riders and dog walkers. These benefits may be substantial and may offset any negative aesthetic or noise effects. The positive effect of such amenities might be particularly strong if the previous land use was essentially barren and of little aesthetic merit. The effects, positive and negative, are likely to vary geographically but not necessarily in the same way.
- E.g.2: Some renewable energy companies provide community and development funds to fund a range of projects that benefit the locality and potentially generate employment. The SSE Clyde wind farm fund²⁴, for example, is expected to provide a total of £17.5 million for local projects that boost local investment and employment, offer training, prevent poverty, or benefit the local or social environment in some way. Such initiatives may improve the quality of life of local residents and increase house prices accordingly.

Patterns of social stratification

Attitudes towards wind turbines and the economic benefits may vary across different social and economic groups. If the location of these groups relative to the location of wind farms varies (e.g. because affluent households are more concentrated in the outskirts in some cities than in others) then the house price responses might vary depending on location.

For example, Kavanagh, Lee and Pryce (2016)²⁵ find that poverty is much more concentrated in the inner city in Dundee than it is in Edinburgh. The maps in Figure 11 below make the same point using the Scottish Index of Multiple Deprivation. Note also that Kavanagh, Lee and Pryce (2016) identify significant change in the geographic patterns of poverty between 2001 and 2010. Since wind turbines tend to be located in rural areas, households living near the edge of the city are most likely to be affected, either positively or negatively, and variations in the pattern of wealth over time and between cities might affect the pattern of house price impact.

²³ <http://www.pfr.co.uk/doich/15/Wind-Power/23/Tourism/>

²⁴ See for example

http://www.southlanarkshire.gov.uk/info/200168/getting_involved_in_your_community/571/sse_clyde_wind_farm_fund

²⁵ Kavanagh, L., Lee, D. and Pryce, G. (forthcoming) Is Poverty Decentralising? Quantifying Uncertainty in the Decentralisation of Urban Poverty. *Annals of the American Association of Geographers*, freely available here: <http://bit.ly/2dAihAX>




Figure 11: 2011 Scottish Index of Multiple Deprivation in Edinburgh and Dundee. Lower values (darker blue) are more deprived areas, higher values are less deprived.

Overall, those who are likely to be able to see a wind turbine typically live in lower value houses (and presumably have lower incomes) than those who cannot (Figure 12). It may be that those on lower incomes are less averse to wind turbines, perhaps because the marginal benefit of any community fund or other positive spillover from wind farm projects is larger relative to their disposable income.




Figure 12: Average annual house prices (plotted on log scale) for houses that will have a turbine in sight at some point within the timeframe of the study vs. those that do not

Interactions between multiple causes

These explanations are not mutually exclusive. It is likely that they affect house prices simultaneously, and to varying degrees in different locations.

These forces may also reinforce or negate each other. They may each wax and wane over time and have different effects at different spatial scales leading to a complex and fluid set of potential outcomes at each point in time.

Further research would be needed to identify which of these effects is most prevalent and persistent. However, it should be noted that the data we collated for this project are unlikely to be sufficient to disentangle these effects in a robust way.

Appendix: Sensitivity analysis

Introduction

We noted above that we use a “fixed effect” methodology to control for a wide range of factors that we cannot observe or measure directly. Provided these factors remain fairly constant over time, we can control for their impact on price trends by introducing additional categorical variables into the analysis. All of the results presented in this report include basic fixed effects that control for differences in dwelling attributes, such as number of bedrooms, which we assume remain constant over time.

We also experimented with a wide number of additional controls. This allows us to test whether our results are robust to changes in how the analysis is set up. For example, we included fixed effects that allow different house price effects to occur over time for: the land gradient (for each individual property); elevation (height above sea level for each property); and aspect (which compass direction the property's slope is facing, indicating which direction their predominant view is likely to be). We also included controls for different price effects across distance rings and we allowed house prices to differ between Scotland's four NUTS2 regions and include a set of region-by-year interactions.

The impacts of these different specifications are presented in the graphs below for each of main categories of results presented under the labels A1, A2, and A3 which relate to the headings used in the main body of the report:

- Figure A1 reports sensitivity analysis for Result #1: Analysis based on Postcode Averages & Wind Farm Centre-Points ('centroids') (Gibbons),
- Figure A2 reports sensitivity analysis for Result #2: Analysis based on Repeat Sales & Individual Turbines
- Figure A3 reports sensitivity analysis for Result #3: Analysis based on Repeat Sales & Individual Turbines, Taking into Account Building Heights

You will see that each of the three figures contains eight sub-graphs, labelled (A) to (H) which give results for each type of fixed effects analysis. The labels for each are explained below:

The first sub-figure, labelled (A), is the "basic" fixed effects used in all analyses:

- **(A) “properties”**: includes fixed effects for time and properties. Note that these results are the same as the results used in the main sections above: they include the same time fixed effects and the property-level fixed effects as those used in Figures 5, 6 and 7 and follow the method described in the "Analysis Step 2" section above. We reproduce them below for ease of comparison with the additional results.

Sub-figures (B) to (D) in Figures A1, A2 and A3 below each add an extra fixed effect on top of the last. In order, these are:

- **(B) "geography"**: fixed effects for slope, elevation and aspect;
- **(C) "rings"**: fixed effects for properties in each distance ring from turbines (or wind farms for figure A1);
- **(D) "NUTS2"**: fixed effects for Scotland's four NUTS2 regions.

Each sensitivity analysis includes a further four sub-figures. These run separate analyses on a particular subset of the data, with each of them using the full set of fixed effects. All three break down properties by their distance from the Scottish coast:

- **(G) “Coast < 2km”**: contains only coastal properties – i.e. those within 2km of the coast;
- **(H) “Coast > 2km”**: contains only inland properties – i.e. those located 2km or more beyond any coastal point.

Sub-figures (E) and (F) vary depending on whether the analysis is based on postcodes/wind farm centre-points or individual dwellings/turbines:

In **Figure A1** the analysis is based on postcode and wind farm centre-points and the results are broken down by wind farm size:

- **A1(E) “Single turbines”**: looks just at single turbine sites;
- **A1(F) “More than one turbine”**: looks at sites with more than one turbine.

In **Figures A2 and A3**, the analysis is based on individual turbines (rather than entire wind farms), and so we can estimate the impact of turbine height:

- **A2(E) and A3(E) “Turbines < 100m”**: plots the impact of turbines that are less than 100m tall;
- **A2(F) and A3(F) “Turbines > 100m”**: plots the impact of turbines over 100m tall.

Note that all graphs in the appendix have the same scale for the vertical axis, which is limited to the plus/minus 15% price change interval. This was done to make each sub-figure directly comparable. Any confidence intervals (i.e. the vertical bars plotted for each estimate) beyond this range are cut off at the 15% limit.

Sensitivity analysis for result #1: based on Postcode Averages & Wind Farm Centre-Points ('centroids') (Gibbons)

The results in the graphs (E) and (F) of Figure A1 allow us to compare the effects of "wind farms" consisting of single turbines (graph A1(E)) and those with two turbines or more (graph A1(F)). Single-turbine effects have wider confidence intervals making the estimates less precise and not statistically different from zero. The estimates are also noticeably less precise for coastal locations (A1(G)) than for inland properties (A1(H)). Controlling for "geography" using fixed effects for slope, elevation and aspect (A1(B)), distance rings (A1(C)) and NUTS2 region (A1(D)) yields relatively precise positive house price effects particularly for the 2-3km distance band.

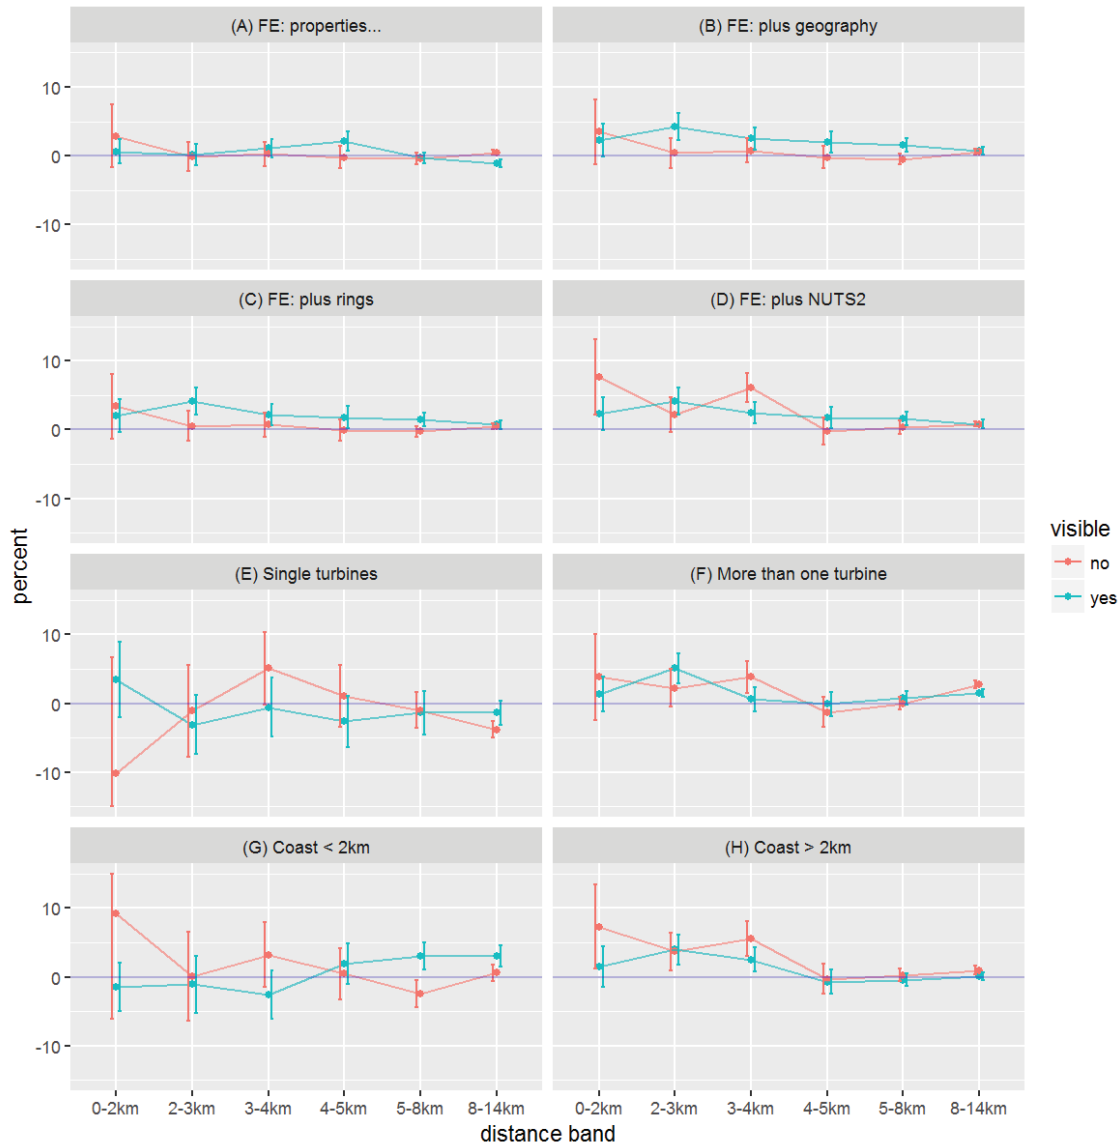


Figure A1: Percent difference in the change of house price
(Postcode/wind farm centroids, whole wind farm visible / not visible)

Sensitivity analysis for result #2: based on Repeat Sales & Individual Turbines




Figure A2: Percent difference in the change of house price
(All repeat sales, turbine visible / not visible)

Sensitivity analysis for result #3: based on Repeat Sales & Individual Turbines and Taking into Account Building Heights




Figure A3: Percent difference in the change of house price
(All repeat sales, turbine visible / not visible accounting for building heights)

Acknowledgements

This work was partly funded by the Economic and Social Research Council (ESRC) through the Applied Quantitative Methods Network: Phase II, Grant Number ES/K006460/1