

[Stephen Gibbons](#)

Gone with the wind: valuing the visual impacts of wind turbines through house prices

**Article (Accepted version)
(Refereed)**

Original citation:

Gibbons, Stephen (2015) *Gone with the wind: valuing the visual impacts of wind turbines through house prices*. [Journal of Environmental Economics and Management](#), 72. pp. 177-196. ISSN 0095-0696

DOI: [10.1016/j.jeem.2015.04.006](https://doi.org/10.1016/j.jeem.2015.04.006)

© 2015 [Elsevier Inc.](#)

This version available at: <http://eprints.lse.ac.uk/62880/>

Available in LSE Research Online: July 2015

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final accepted version of the journal article. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

Gone with the wind: valuing the visual impacts of wind turbines through house prices

Stephen Gibbons^a

March 2015

Key words: Housing prices, environment, wind farms, infrastructure, energy

JEL codes: R3,Q4, Q51

Acknowledgements: This research was produced as part of the 2011-2014 research programme of the Spatial Economics Research Centre at the London School of Economics. The programme was funded by grants from the UK Economic and Social Research Council (ES/J021342/1), the UK Government Department of Business Innovation and Skills, and the Welsh Advisory Government. The research makes use of data provided by the ESRC funded UK Data Archive, the UK Land Registry, the Office for National Statistics and the EDINA data centre at the University of Edinburgh.

^a Department of Geography and Environment, London School of Economics and Spatial Economics Research Centre, London, United Kingdom. Stephen Gibbons (email: S.Gibbons@lse.ac.uk).

Abstract:

This study provides quantitative evidence on the local benefits and costs of wind farm developments in England and Wales, focussing on their visual environmental impacts. In the tradition of studies in environmental, public and urban economics, housing sales prices are used to reveal local preferences for views of wind farm developments. Estimation is based on quasi-experimental research designs that compare price changes occurring in places where wind farms become visible, with price changes in appropriate comparison groups. These groups include places close to wind farms that became visible in the past, or where they will become operational in the future and places close to wind farms sites but where the turbines are hidden by the terrain. All these comparisons suggest that wind farm visibility reduces local house prices, and the implied visual environmental costs are substantial.

Key words: Housing prices, environment, wind farms, infrastructure, energy

JEL codes: R3,Q4, Q51

1 Introduction

Renewable energy technology clearly provides potential global environmental benefits in terms of reduced CO₂ emissions and slower depletion of natural energy resources. However, like most power generation and transmission infrastructure, the plant, access services and transmission equipment associated with renewable electricity generation may involve environmental costs. This is particularly so in the case of wind turbine developments, where the sites that are optimal in terms of energy efficiency are typically in rural, coastal and wilderness locations that offer many natural environmental amenities. These natural amenities include the aesthetic appeal of landscape, outdoor recreational opportunities and the existence values of wilderness habitats. The visual impacts of these 'wind farms' may be especially important because they are often on high ground with extensive visibility. Although views on their aesthetic appeal are mixed, there is evidently considerable dislike for their visual impact on the landscape, with 23% of respondents in a poll of 1001 residents in Scotland in 2010 agreeing or strongly agreeing that wind farms "are, or would be, ugly and a blot on the landscape" (You Gov 2010). It should be noted, however, that only 51% of respondents had actually seen a wind farm in real life. In addition to these potential impacts on landscape, residents local to operational wind turbines have reported health effects related to visual disturbance and noise (e.g. Bakker et al 2012, Farbouda et al 2013).

The UK, like other areas in Europe and parts of the US has seen a rapid expansion in the number of these wind turbine developments since the mid-1990s. Although these wind farms can offer various local community benefits, including shared ownership schemes, community payments and the rents to land owners, in the UK, and elsewhere in Europe, wind farm developments have faced significant opposition from local residents and other stakeholders with interests in environmental preservation. This opposition suggests that the environmental costs may be important. The issue is highly controversial, given that opinion polls and other surveys generally

indicate majority support of around 70% for green energy, including windfarms, (e.g. results from the Eurobarometer survey in European Commission 2006). This contradiction has led to accusations of 'nimbyism' (not in my backyard-ism), on the assumption that it is the same people opposing wind farm developments in practice as supporting them in principle. There is perhaps less of a contradiction when it is considered that the development of wind farms in rural locations potentially represents a transfer from residents in these communities and users of natural amenities (in the form of loss of amenities) to the majority of the population who are urban residents (in the form of energy). Other possible explanations for the tension between public support and private opposition to wind energy developments are discussed at length in Bell et al (2007).

This paper provides quantitative evidence on the local benefits and costs of wind farm developments in England and Wales, focussing on the effects of wind turbine visibility, and the implied cost in terms of loss of visual landscape amenities. In the tradition of 'hedonic' studies in environmental, public and urban economics, housing sales prices are used to reveal local preferences for views of wind farms. This is feasible, because wind farms in England and Wales are often close to and visible from residential areas in rural, semi-rural and even urban locations, so the context provides a large sample of housing sales that are potentially affected (at the time of writing, around 1.8% of residential postcodes are within 4 km of operational or proposed wind farm developments). The study offers a significant advance over previous studies, which have mostly been based on relatively small samples of housing transactions and cross-sectional price comparisons. Estimation in this current work is based on quasi experimental, difference-in-difference based research designs that compare price changes occurring in postcodes where wind farms become visible, with postcodes in appropriate comparison groups. These groups include: places where wind farms became visible in the past, or where they will become visible in the

future and places close to where wind farms became operational but where the turbines are hidden by the terrain. The postcode fixed effects design implies that the analysis is based on repeat sales of the same, or similar housing units within postcode groups (typically 17 houses grouped together). Kuminoff, Parmeter and Pope (2010) provide a discussion of the advantages of quasi-experimental approaches of this type in the context of hedonic methods for environmental valuation.

The overall finding is that operational wind farm developments reduce prices in locations where the turbines are visible, relative to where they are not visible, and that the effects are causal. This price reduction is around 5-6% on average for housing with a visible wind farm within 2km, falling to under 2% between 2-4km, and to near zero between 8-14km, which is at the limit of likely visibility. Evidence from comparisons with places close to wind farms, but where wind farms are less visible suggests that the price reductions are associated with turbine visibility. As might be expected, large visible wind farms have much bigger impacts that extend over a wider area.

The remainder of the paper is structured as follows. Section 2 discusses background policy issues and the existing literature on wind farm effects. Section 3 outlines the data used for the analysis. Section 4 describes the empirical strategy and Section 5 the results. Finally, Section 6 concludes.

2 Wind farm policy and the literature on their local effects

In England and Wales, many wind farms are developed, operated and owned by one of a number of major energy generation companies, such as RES, Scottish Power, EDF and E.ON, Ecotricity, Peel Energy, though some are developed as one-off enterprises. Currently, wind farms are potentially attractive businesses for developers and landowners because the electricity they generate is eligible for Renewables Obligation Certificates, which are issued by the sector regulator (Ofgem) and guarantee a price at premium above the market rate. This premium price is

subsidised by a tariff on consumer energy bills. The owners of the land on which a wind farm is constructed and operational will charge a rent to the wind farm operator. Media reports suggest that this rent could amount to about £40,000 per annum per 3 MW turbine (Vidal 2012).

The details of the procedures for on-shore wind farm developments in England and Wales have evolved over time, but the general arrangement is that applications – in common with applications for most other types of development - have to pass through local planning procedures. These procedures are administered by a Local Planning Authority, which is generally the administrative Local Authority, or a National Park Authority. Very small single wind turbines (below the scale covered by the current analysis) can sometimes be constructed at a home, farm or industrial sites within the scope of ‘permitted development’ that does not require planning permission. The planning process can take a number of years from the initial environmental scoping stage to operation, and involves several stages of planning application, environmental impact assessment, community consultation and appeals.¹ Once approved, construction is relatively quick. According to public information from the European Wind Energy Association², a 10 megawatt wind farm (3-4 turbines) can be constructed in 2 months, and a larger 50 megawatt wind farm in 6 months (the average size wind farm in this current study is around 18 Mw). Large wind farms (over 50 Mw) need approval by central government. Offshore wind farms are also subject to a different process and require approval by a central government body.

Wind farms have potential local economic benefits of various types. Interesting qualitative and descriptive quantitative evidence on the community and local economic development benefits of

¹ E.g. Peel Energy <http://www.peelenergy.co.uk/> provide indicative project planning timelines for their proposed wind farm developments

² <http://www.ewea.org/wind-energy-basics/faq/> accessed February 2014

wind farms in Wales is provided by Munday et al (2011). Potential benefits include the use of locally manufactured inputs and local labour, discounted electricity supplies, payments into community funds, sponsorship of local events, environmental enhancement projects, and tourism facilities. They argue that the local economic development effects have been relatively limited, although in many of the communities surveyed (around 21 out of 29 wind farms) payments were made to community trusts and organisations, and these contributions can be quite substantial – at around £500-£5000 per megawatt per annum. Based on these figures, a mid-range estimate of the community funds paid out to affected communities in Wales would be about £21,000 per wind farm per year. For the US, Kahn (2013) argues that wind farm counties generate benefits for their communities because the revenues to land owners spill over to the community in general, through lower property tax rates and improved public expenditures. This direct link between local taxation and school resources is more important in the US, than in the UK where schooling is financed mainly through central government grants. Using data and fairly descriptive quantitative evidence from counties in Texas, he finds some signs of increases in school resources relative to non-wind farm counties and lower property tax rates, and no evidence that wind farms have deterred higher-educated residents from moving in to the area.

There is also an extensive literature on attitudes to wind farm developments, the social and health aspects, and findings from impact assessments and planning appeals. Most existing evidence on preferences is based on surveys of residents' views, stated preference methods and contingent valuation studies and is mixed in its findings.

There have been several previous attempts to quantify impacts on house prices in the US. Hoen et al (2011) apply cross-sectional hedonic analysis, based on 24 wind farms across US states. Their study is interesting in that it makes the comparison between price effects at places where turbines

are visible compared to places where nearby turbines are non-visible (a technique which is applied later in the current paper) but finds no impacts. For the UK, Sims et al (2007, 2008) also conduct a cross-sectional hedonic analysis of around 900 property sales, which all postdate construction, near three wind farms in Cornwall. Again this study finds no effects. One study with a larger housing sample size Lang et al (2014), looks at 10 small-scale wind farms in suburban and urban locations in Rhode Island, all but one of which are single-turbine sites. The authors provide difference-in-difference estimates and repeat sales estimates, based on changes in prices over a 14 year interval. Their sample has 2670 housing transactions within 1 mile (2.25km) over this period, with 338 sales post-dating construction. They report no significant effects on housing prices from the wind farms, but these are small wind power developments in an area that is already highly developed rather than rural. The results are therefore difficult to generalise to the case of large scale wind farms like those in the UK and elsewhere in the US and Europe.³ Even so, the point estimates are in some cases large, with the repeat sales analysis suggesting falls of more than 6% within 2 miles after announcement of the wind farms, although the estimates are rarely statistically significant.

Another study from the US, Hoen et al (2013), attempts a difference-in-difference comparison for wind farms, but using cross-sectional comparisons between houses at different distances from the turbines. This study uses fairly sparse data on 61 wind farms across nine US states. The sample contains over 50,000 transactions, but very few transactions in the areas near the wind farms: only 1198 transactions reported within 1 mile of current or future turbines (p20) and only 300 post-dating construction. Their cross-sectional difference-in-difference comparison is between places beyond and within 3 miles of a wind farm site and the research design does not exploit price changes or repeat sales. The conclusions of the paper are that there is 'no statistical evidence that

³ Their regressions also control for an unspecified number of city-by-quarter fixed effects, which seem likely to absorb much of the impact of the wind farms on prices making it difficult to detect any effects even if they exist.

home values near turbines were affected' by wind turbines, which is true in a literal sense. However, as in Lang et al (2014), the point estimates indicate some quite sizeable effects; it is the fact that the point estimates are imprecise and have big standard errors that makes them statistically uninformative. A similar conclusion is reached, for similar reasons, in Vyn and McCullough (2014) who study the impact of turbines in a large windfarm in Canada on neighbouring farmland and residential sales. Their dataset includes over 5000 residential sales and over 1500 farm sales, and the authors went to considerable trouble to determine turbine visibility. Sadly though, only a very small number of sales occur after turbine construction. A total of 18 sales occur within 1km and 79 within 5km (their Table 2) after the wind farm was built. Inevitably this means the results are not very informative and are very imprecise. As in many previous studies, the standard errors are so large and the point estimates vary so much from specification to specification, that the authors can only conclude that "while the results indicate a general lack of significantly negative effects across the properties examined in this study, this does not preclude any negative effects from occurring on individual properties" and note that "a recent appraiser's report on the impacts of Melancthon's wind turbines ... found that the values of five specific properties in close proximity to turbines declined by up to 59%." (Vyn and McCullough p.388).

In contrast, the current study has nearly 38,000 quarterly, postcode-specific housing price observations over 12 years, each representing one or more housing transactions within 2km of wind farms (about 1.25 miles). Turbines are potentially visible for 36,000 of these. There is therefore a much greater chance than in previous work of detecting price effects if these are indeed present.

3 Data

Information on wind-farm location (latitude and longitude), characteristics and dates of events was provided by RenewableUK, a not for profit renewable energy trade association (formerly BWEA). This dataset records dates of operation and other events related to their planning history, number of turbines, MW capacity and height of turbines (to tip). The dates in these data relate to the current status of the wind farm development, namely application for planning, approval, withdrawal or refusal, construction and operation. Unfortunately these public data do not provide a complete record of the history for a given site, because the dates of events are updated as the planning and construction process progresses. Therefore, for operational sites, the dates of commencement of operation are known, but not the date when planning applications were submitted, approved or construction began. This limits the scope of investigation of the impact of different events in the planning and operation process, other than for cases where there is a final event recorded, and this version of the paper makes use of operational wind farms only.

A GIS digital elevation model (DEM)⁴ was combined with this wind-farm site and height data to generate 'viewsheds' on 200m grid. These viewsheds were used to differentiate residential postcodes (geographical units with approximately 17 houses) into those from which the wind farm is visible, and those from which it is less likely they are visible, using information on the underlying topography of the landscape. These viewsheds provide approximate visibility indicators, both in terms of the 200m geographical resolution of the view sheds (necessary for manageable computation times), and because they are based on wind-farm centroids, not individual turbines. This means that in the case of large wind farms, turbines may be visible from

⁴ GB SRTM Digital Elevation Model 90m, based on the NASA Shuttle Radar Digital Topography Mission and available from the EDNIA ShareGeo service <http://www.sharegeo.ac.uk/handle/10672/5>

locations which the procedure classifies as non-visible, given a large wind turbine array can extend over 1km or more. However, the median wind farm development in the data contains only 6 turbines, so the errors introduced by basing visibility on site centroids are likely to be small. Note that the error will in general result in mis-classification of sites from which the turbines are deemed non-visible, given that if the tip of a turbine at the centroid of the site is visible, it is almost certain that at least one turbine is visible. The viewsheds also take no account of intervening buildings, trees and other structures, because Digital Surface Models which take account of such features are not yet available for the whole of England and Wales. As a further refinement, to eliminate cases where visibility was highly ambiguous, I calculated the rate of change of visibility from one 200m grid cell to the next, and dropped postcodes in cells in the top decile of this visibility gradient. In general, misclassifications in terms of visibility, and measurement error in distance to wind turbines will tend to attenuate the coefficients in regression-based estimates. This implies that the results that follow may, if anything, under-estimate the effects of wind farm distance and visibility on prices.

Given the focus of this study on the visual impacts of wind farms in rural areas, a number of single-turbine wind farms in urban areas and industrial zones were excluded from the analysis (around 21 operational turbines are dropped). Land cover estimates were used first to restrict the analysis to wind farms outside zones with continuous urban land cover. Some additional turbines were eliminated on a case-by-basis where the information available in the wind farm data, and reference to web-based maps and information sources, suggested that turbines were on industrial sites within or close to major urban areas. The land cover at the wind farm centroid was obtained by overlaying the wind farm site data with 25m grid based land cover data (LandCoverMap 2000 from the Centre for Ecology and Hydrology). Land cover was estimated from the modal land

cover type in a 250m grid cell enclosing the wind farm centroid. In cases where no mode exists (due to ties), the land cover in the 25 m grid cell enclosing the centroid was used.

Housing transactions data come from the England and Wales Land Registry 'price paid' housing transactions data, from January 2000 to the first quarter of 2012. Data going back to 1995 are available at the time of writing, but was not yet available at the time the dataset for this analysis was created. The 'price paid' data include information on sales price, basic property types – detached, semi-detached, terraced or flat/maisonette – whether the property is new or second-hand, and whether it is sold on freehold or leasehold basis. The housing transactions were geocoded using the address postcode and aggregated to mean values in postcode-by-quarter cells to create an unbalanced panel of postcodes observed at quarterly intervals (with gaps in the series for a postcode when there are no transactions in a given quarter). For a small subset of the data, floor area and other attributes of property sales can be merged from the Nationwide building society transactions data. Demographic characteristics at Output Area (OA) level from the 2001 Census were merged in based on housing transaction postcodes. These additional characteristics are used in some robustness checks which appear later in the empirical results.

Postcode and wind farm visibility data were linked by first forming a panel of postcodes at running quarterly (3 month) intervals over the period January 2000-March 2012. The cumulative number of operational turbines within distance bands of 0-1km, 1-2km, 2-4km, 4-8km and 8-14km of each postcode was then imputed at quarterly intervals by GIS analysis of the information on site and postcode centroids. The 14km limit is set in part to keep the dataset at a manageable size, but also 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. Existing literature based on field work suggests that large turbines are potentially

perceptible up to 20km or more in good visibility conditions, but 10-15km is more typical for casual observer and details of individual turbines are lost by 8km (University of Newcastle 2002). In the next step, the site viewsheds were used to determine whether wind-farm sites are visible or not visible from each postcode in each quarter, again using GIS overlay techniques. Additional GIS analysis with the Digital Elevation Model provided estimates of the elevation, slope and aspect (North, East, South and West in 90 degree intervals) of the terrain at each postcode, plus visibility of coastline for use in a robustness check. These are potentially important control variables, because places with good views of wind farms may have good views generally, be more exposed to wind, or have more favourable aspects, and these factors may have direct effects on housing prices.

Finally, the housing transactions and wind farm visibility data were linked by postcode and quarter to create an end product which is an unbalanced panel of postcode-quarter cells, with information on mean housing prices and characteristics, the cumulative number of visible and non-visible operational turbines within the distance bands, plus additional variables on terrain and demographics. Note, prices in quarter t are linked to the turbine data at $t-1$, so although the price data extends to the first quarter of 2012, only wind farm developments up to the last quarter of 2011 are utilised. The next section describes the methods that are applied using these data to estimate the house price effects of wind farm developments.

4 Estimation strategy

The research design involves fixed-effects, regression-based difference-in-difference methods. In all cases, the research strategy is to compare the average change in housing prices in areas where and when wind farms become operational and visible, with the average change in housing prices in some comparison group.

4.1 Comparing the effects of new wind farms with existing and future wind farms

The simplest approach is to compare the price changes occurring around the time a wind farm becomes visible and operational, with the price changes occurring in comparable areas where wind farms are already visible and operational or where they will become so in the future. The idea is that postcodes close to existing or future wind farm locations and where these wind farms are or will be visible, provide a suitable counterfactual for places where new wind farms are becoming operational and visible in the current period. These postcodes close to and with views of new, existing and future wind farms are likely to be similar to each other in respect of: a) being physically suitable for wind farm developments; b) being viable for development in terms of the planning and construction process; and c) having topography that means that turbines are likely to be visible.

To implement this approach, I estimate the following regression specification, on the sample of postcodes which had visible turbines within a given distance radius at the beginning of the study period (2000), or will have visible turbines within these radii or bands by the end of it (2011)⁵:

$$\ln price_{it} = \sum_k \beta_k (visible, j_k < dist < k, operational)_{it-1} + x'_{it} \gamma + f(i, t) + \varepsilon_{it} \quad (1)$$

Here $price_{it}$ is the mean housing transaction price in postcode i in quarter t . The variable capturing exposure to wind-farm developments is $(visible, j_k < dist < k, operational)_{it-1}$. This is a dummy (1-0)

⁵ More precisely, a postcode is included in the sample for estimating (1) if it has a visible wind turbine development within the specified distance band before January 2000 or if turbines become visible over the course of the study period from 2000 to 2011. In this sample of postcodes the treatment indicator equals 1 for at least one quarter over the sample period. A postcode that has, for example, a visible, operational wind farm within 4km opening in the last quarter of 2004 will be included in the sample, but will have $(visible, 0 < dist < k, operational)_{it-1} = 0$ in all quarters up to t corresponding to the first quarter of 2005, and $(visible, 0 < dist < k, operational)_{it-1} = 1$ in all quarters thereafter. Postcodes with at least one visible, operational turbine from the beginning of the study period are included in the sample, but have the indicator $(visible, 0 < dist < k, operational)_{it-1} = 1$ throughout.

treatment variable, indicating that postcode i has at least one visible-operational turbine between j_k and k km distance in the previous quarter. Vector x_{it} is an optional set of control variables, including housing characteristics. The function $f(i, t)$ represents a set of general geographical and time effects which will be controlled for using postcode fixed effects plus interactions between geographical and time dummies, as described in more detail below. The coefficient of interest β_k is the average effect on housing prices of wind farm turbines visible within distance band j_k-k . The sign of β_k is ambiguous a priori, since it depends on the net effects of preferences for views of wind farms, the impact of noise or visual disturbance – at least for properties very close to the turbines – and other potential local gains or losses, such as spillovers from land owner rents, shares in profits, community grants, or employment related to turbine maintenance and services.

This wind farm visibility indicator for a given postcode $(visible, j_k < dist < k, operational)_{it-1}$ is an interaction between an indicator that turbines are potentially visible from the postcode ($visible_i$), an indicator that these turbines are within a given distance band of the postcode ($j_k < dist < k$), and a ‘post-policy’ indicator which indicates that the turbines have been built and have become operational ($operational_{it-1}$).⁶ This date of operation is taken as the date around which the wind farms impact on prices because my data contain no information on the date when the wind farm development was announced or when construction started or finished.

Two versions of the distance specifications in (1) are used in the empirical work. I start with the simplest specifications in which the regressions are estimated for different values of k (1km, 2km, 4km, 8km, 14km) and $j_k = 0$, i.e. β_k estimates the effects of visible wind farms within a radius k .

⁶ Note, it is not necessary to explicitly control for the separate components ($visible, j_k < dist < k$ and $operational$) because these are subsumed through the specification of geographical and time fixed effects $f(i, t)$ described below.

The estimation sample is restricted to postcodes with potentially visible turbines within distance k . In the second case, a series of distance band indicators is used ($0 < \text{distance} \leq 1\text{km}$, $1\text{km} < \text{distance} \leq 2\text{km}$, $2\text{km} < \text{distance} \leq 4\text{km}$, $4\text{km} < \text{distance} \leq 8\text{km}$ and $8\text{km} < \text{distance} \leq 14\text{km}$) in a single regression, and the sample is restricted to postcodes with visible turbines within the maximum 14km. The distance thresholds are chosen somewhat arbitrarily in order to give reasonably detailed delineation of the distance decay close to wind farm sites, while allowing for potential impacts up towards the limits of visibility.

Crucially, specification (1) must allow for unobserved components which vary over time and space $f(i,t)$ which are potentially correlated with the wind farm visibility indicator. This correlation with the geographical effects occurs because wind farms are not randomly assigned across space and postcodes close to wind farms and where turbines are visible may not be comparable to postcodes further away in terms of the other amenities that affect housing process. The correlation with the time effects occurs because the number of wind farms is growing over time, so there is obviously a spurious correlation between any general trends in prices over time and the indicator of wind farm visibility.

It is therefore essential to control in a quite general way for geographical fixed effects and time trends that are related to wind farm proximity and visibility. This is done firstly through the restriction to postcodes that have, or will have, visible wind farm developments close by. Secondly, postcode fixed effects are eliminated using the within-groups transformation (i.e. differences in the variables from postcode-specific means) and common time effects eliminated by including quarter-specific dummies (i.e. for the 48 quarters spanned by the data). Furthermore, in the distance-band version of the specification, separate sets of year dummies for each distance band, $j_k < \text{dist} < k$, are included control for differences in the price trends in these different distance

bands (i.e. interactions between $j_k < dist < k$ dummies and year dummies). Additional time varying geographical effects are captured by interactions between year dummies, and dummies for categories of postcode elevation (0-25m, 26-50m, 51-100, >100m), slope (0-0.5%, 0.51-1%, 1.01-1.5%, 1.51-2.5%, >2.5%), and aspect (315-45 degrees, 46-135 degrees, 136-225 degrees, 226-316 degrees). These terrain variables are potentially important, because wind farm visibility depends on the elevation, slope and direction of the land at the postcode location. Some supplementary specifications include region-by-year dummies to control for general spatial trends, where regions are defined by splitting the sample into north, south east and west geographical quadrants.

Since the specification controls for postcode fixed effects, the estimation method exploits changes in average prices between the post-operation and pre-operation periods and β_k is estimated from postcodes that have housing transaction observations before and after a wind farm becomes operational. However, postcodes that have sales only before, or only after wind farm operations, including wind farms visible at the start of the study period in 2000, form part of the control group and contribute to estimation of the time trends and other parameters that are common across postcodes. The estimates of β_k from the within-postcode fixed effects estimator should be interpreted as the average price change between the pre- and post-operation periods, given the time spanned by the housing sales data (not necessarily the step change in price occurring at the time of operation, nor the full long run price effect from the period prior to planning announcement to the post-operation period). Given the data and setting, the within-groups estimator which compares the post-operation average price with the pre-operation average price over the whole sample period, is preferable to a specification using differences between two time periods. This is because: a) there is unlikely to be a step-change in prices coincident with wind farm operation, both because price changes evolve slowly, and because there may be pre-operation price changes after announcement; and b) the panel is sparse and unbalanced, with missing

periods where there are no price transactions in a given postcode, so working with differences over specific time intervals within postcodes would result in a large reduction in sample size (e.g. a 4 quarter difference can only be observed in postcodes where there happen to be sales observed 4 quarters apart).

4.2 *Comparing the effects of visible and non-visible turbines*

It is well known that difference-in-difference based research designs suffer from the problem of pre-existing differences in trends between the ‘treatment’ and ‘control’ groups. In the method described above, this problem is mitigated by using wind-farm locations as both treatment and control groups. Postcodes with existing visible-operational turbines, and postcodes with potentially visible turbines that become visible-operational in the future, provide information on the counterfactual price changes for postcodes in which turbines have just become visible-operational. However, this method may not completely take care of more subtle short run differential trends in the affected postcodes, e.g. if wind farms are intentionally or coincidentally targeted to particular places during periods in which these places have falling or rising prices relative to places that saw wind farm developments in the past, or will see them in the future. In addition, if the aim is to interpret β_k as the visibility impact of wind farms, estimates from (1) will be biased by any price effects arising through other channels such as local benefits, or costs due to noise.

To obtain cleaner estimates of the impacts of wind farm visibility, I augment specification (1) with additional treatment indicators, for postcodes close to wind-farms, but where the turbines are likely to be hidden from view by the landscape topography. This approach provides a powerful test of the robustness of the main findings on visibility, because the postcodes with non-visible-operational turbines within a given radius of the turbines are in the same geographical areas as the

postcodes with visible turbines. These two visible and non-visible groups are thus likely to be closely comparable on unobserved dimensions, and subject to similar unobserved price trends arising through other causal channels. One concern might be that topographic features that obscure a wind farm from view from a property also reduce the noise level, meaning that comparisons between the visible and non-visible groups also capture differences in noise levels. In practice this is very unlikely. The predicted combined noise level from a wind farm with a ten turbine array, with each turbine emitting a typical 100dbA, falls to around 40dbA by 1km, which is below the background noise level in an average home.⁷ At 2km the noise level is around 34dBA. Moreover, much of the nuisance noise from wind farms is low frequency, and low frequency sound in particular is not attenuated by large topographic features due to refraction. At distances beyond 1km, comparisons between groups with visible and non-visible turbines are very unlikely to pick up noise-related effects.

The structure of the regression specifications for these visible-non-visible comparisons is identical to (1) but the sample now includes the sample of postcodes with potentially visible-operational turbines plus the sample of postcodes which are close to the same set of turbines, but where these are non-visible. Accordingly, specification (2) uses a treatment indicator that is an interaction of an indicator that there are no visible wind farms (*non-visible*) at the postcode, that the postcode is within a given radius or distance band ($j_k < dist < k$) and the indicator that the turbines are operational (*operational*):

⁷ Calculations based on the National Physical Laboratory wind turbine noise model
<http://resource.npl.co.uk/acoustics/techguides/wtnm/>

$$\begin{aligned}
\ln price_{it} = & \sum_k \beta_k (visible, j_k < dist < k, operational)_{it-1} \\
& + \sum_k \delta_k (non - visible, j_k < dist < k, operational)_{it-1} \quad (2) \\
& + x'_{it} \gamma + f(i, t) + \varepsilon_{it}
\end{aligned}$$

In this setup, the estimated parameters δ_k are estimates of the effects on house prices of proximity to operational turbines, when there is no impact from the turbines being visible in the neighbourhood. These sign of these effects is theoretically ambiguous, for reasons discussed above for visible operational turbine, because there are potential community benefits and potential costs. If there are local community benefits, then the visibility parameters β_k will be underestimates of the costs associated with wind farm visibility, because these impacts are already partly compensated by these other benefits (as in the classic wage-price-amenity trade off in the Roback model of compensating wage and land price disparities in Roback 1982). However, the difference-in-difference-in-difference estimate of $\beta_k - \delta_k$ provides a cleaner estimate of the specific impact of wind farm visibility – i.e. the increase in the gap between house prices in places where wind farm sites are visible and where they are not visible, once the turbines are built. This estimate thus provides an explicit estimate of willingness to pay through housing expenditure to avoid views of wind turbines and estimate of the monetary value of the visual dis-amenity associated with them.

In these specifications with visible and non-visible indicators, the set of geographical-by-time effects is extended to include separate quarterly trends for postcodes with visible and non-visible turbines (i.e. interactions between $0 < dist < K$, *non-visible* and quarter dummies, and interactions between $0 < dist < K$, *visible* and quarter dummies, where K is the maximum radius included in the particular specification). As before the specification also includes separate sets of year dummies for each distance band (i.e. interactions between $j_k < dist < k$ dummies and year dummies)

interactions of year dummies with elevation, slope and aspect indicators and control variables for property characteristics.

A number of other robustness checks are carried out to assess sensitivity to local price trends, changing composition of housing sales, and assumptions about the clustering of standard errors. These are described where they arise in the Results section below.

4.3 *Specifications for effects by wind farm size*

The set up described above is based around a treatment effect design with a simple 1-0 indicator of turbine visibility and operation, and thus implicitly estimates the effect of wind farms of average size. Clearly, the impacts are likely to differ by wind farm size (number of turbines) and there are likely to be interactions of size with distance, especially if visibility turns out to be an important influence on prices. I therefore estimate final specifications that look at the interactions between wind farm size and distance, using a similar set up to (1), but with separate indicators for the number of turbines visible and operational at each distance and the number of turbines.

5 Results

5.1 *Descriptive figures and statistics*

Figure 1 shows the historical development of non-urban wind turbines in England and Wales from the mid-1990s to 2011. By the end of 2011, these turbines could provide up to 3200mw of generating capacity, which, in principle, amounts to sufficient power for about 1.8 million homes (or around 7.7% of the 23.4 million households in England and Wales)⁸. Figure 2 illustrates the

⁸ This figure is estimated from DECC 2013a and DECC 2013b as follows. Total UK electricity output from onshore and offshore wind was 15.5TWh in 2011 (DECC 2013a Table 6.4) from 6500MW total capacity. Scaling down to the capacity of 3200MW in England and Wales, suggests an output of 7.6 TWh from wind farms in England and Wales. Average UK domestic household electricity consumption is 4.2×10^{-6} TWh, based on total domestic electricity consumption of 111.6TWh (DECC2013b, Table 5.1.2), and a figure of 26.4 million households in the UK (2011 Census). Therefore, wind farms in England and Wales could power approximately $7.6/4.2 \times 10^{-6} = 1.8$ million households.

evolution of the spatial distribution of these turbine sites between 2000 and 2011. These sites have, over the whole period, been predominantly in coastal and upland areas. These are often on the peripheries of areas that are valued for their natural beauty, although wind farms have not been permitted in National Parks. Examples are the Cornwall peninsula in the south west, Wales in the west, the Pennines in northern England, and around the coast Cumbria (the Lake District) in north west England. There are also concentrations around Sunderland in the north east, and Lincolnshire, Norfolk and the Wash in the east. In general, density has increased in these regions, rather than the distribution spreading across regions, although new wind farms have appeared in eastern central England in recent years. There are very few sites in the south and east of England.

Some basic summary statistics for the operational, non-urban wind farms in the dataset are shown in Table 1. There are 148 wind farms recorded in operation in England and Wales over this period (after eliminating some single-turbine urban and industrial sites). The mean operational wind farm has 11 turbines (6 median) with a capacity of 18.6 MW, but the distribution is highly skewed, with a maximum number of turbines of 103 and capacity of 150MW. These largest wind farms are off-shore. The average height to the tip of the turbine blades of just over 90m, though the tallest turbines (mainly offshore) reach to 150m. The distribution of wind farms across land cover types is given in the table notes and shows that most wind farms are in farmland locations, followed by mountain and moorland locations. Offshore sites are also included in the analysis, where these are potentially visible from residential areas on shore. Urban and most industrial locations (except where these impact on rural areas) are excluded from the analysis.

Table 2 summarises the main postcode-by-quarter aggregated panel data set, with information on property prices and characteristics, and the distribution of visible and non-visible operational turbines. The top panel with the housing summary statistics relates to the sample of postcodes

with operational turbines within 14km in 2000, or appearing within 14km at some time over the sample period up to the end of 2011. Price dataset is merged to the wind farm dataset with a one-quarter lag, so the price series runs from the first quarter of 2000 to the first quarter of 2012. Changing the lag to 6 months made essentially no difference to the regression results presented below. To show the spatial structure of the data, the second panel shows the number of postcodes in the data at different wind farm distances, categorised according to whether the wind farms are visible (based on the modelled view-shed). Note that many postcodes have both visible and non-visible turbines over the whole period. The third panel provides information on how many of the postcodes that will have visible turbines, have sales in both pre and post operation periods. This panel also shows the mean time interval between sales in the pre and post periods. There are 1125 postcodes with visible turbines within 1km, though only 468 of these have repeat sales in pre and post periods. Wind farms are visible from nearly all these postcodes. As we move further out, the number of postcodes increases to over 220000 and the proportion from which turbines become visible decreases to around 56% within 14km band. At greater distances it becomes more likely that views from the postcode neighbourhood are obscured by intervening terrain. The mean interval between sales in the pre and post operation periods is stable over all distances at around 23 quarters (5.75 years), implying that the regression estimates that follow will represent the average price change occurring over this time interval. Overall there around 7.75 repeat observations for each postcode ($=1710293/220669$ from the numbers in the table). The median number of transactions (not reported in the table) per postcode-quarter cell is 1 with a median of 1 and a 99th percentile of 5.

The methods described in 4.2 proposed comparing the price effects in postcodes with visible-operational turbines to the price effects in postcodes with non-visible operational turbines. To illustrate the basis for this approach, Figure 3 shows the viewshed for a wind farm in north east

England. This is the Haswell Moor wind farm in County Durham, which has 5 turbines, a total capacity of 10MW and the height to the tip of the turbines is 110m. This is a fairly typical wind farm development in the sample. The dark shaded areas are residential postcodes and the light grey shading indicates the land where at least the tips of the turbine blades are visible (technically, these are computed as the land surface that is visible to an observer at the tip of the turbine). Results presented in the next section compare prices changes occurring with the start of wind farm operation in these postcodes where the turbines are visible, with those occurring where they are not-visible.

5.2 *Baseline regression results on visibility and robustness tests*

Table 3 reports the coefficients from a baseline set of postcode fixed-effects regressions of prices on wind farm proximity and visibility indicators discussed in Sections 4.1 and 4.2, using separate regressions for different radii, from 1 to 14km. For each radius, the first two columns restrict the sample to postcodes which have or will have an operational wind farm within the specified distance following the approach of Section 4.1. Identification comes purely from comparing the change in mean postcode-quarter specific prices between the periods before and after the wind farm operation, with the changes occurring in postcodes that have already got visible-operational wind farms or which will do so in the future. For radii above 1km, the third column at each radius extends the sample to include postcodes which have or will have non-visible operational wind farms within the specified distance following the approach of Section 4.2 (this is infeasible at 1km since almost all postcodes have wind farms visible). The regression in the first column of each set has no control variables other than quarterly dummy variables. Other columns control for the property characteristics and the array of geographical trends described in the methods section. Standard errors are clustered at Census Output Area level (10 or so postcodes) to allow for serial

correlation in the errors over time and spatial correlation in the price changes across neighbouring postcodes.

The key finding from this table is that prices in postcodes where wind farms are close and visible are reduced quite substantially over the period in which a wind farm becomes operational. The price impact is around 6.5% within 1km, falling to 5.5-6% within 2km, 2.5-3% within 4km. Beyond 4 km the effect falls below 1% and becomes statistically insignificant, at least once control variables are included. Generally, controlling for property characteristics and the array of terrain-by year dummies makes little difference to the results, suggesting that unobserved price trends or changes in the types of housing being sold do not affect the results substantively.

Columns 5, 8, 11 and 14 include indicators of proximate non-visible wind farms, and tell us more about the specific visibility impacts of wind farms, as distinct from other costs and benefits associated with their operation. The point estimates within the 2km band are similar to those for visible-operational turbines, but statistically insignificant, given that the small share of postcodes with non-visible wind farms within 2km (5% from Table 2). In part, the coefficient on non-visible wind farms within 2km may be picking up impacts on the few sales much closer to wind farms, where turbines are not visible but noise may be an issue (the estimates later on in Table 6 present the impacts in distance bands to address this issue). Further out, a more interesting pattern emerges. Within 4km (where wind farms are hidden for 18% of postcodes) there is no effect on prices from non-visible operational turbines, while visible wind farms reduce prices by 2.4%. This comparison suggests that the negative effects from visible-operational turbines are specifically attributable to visibility. Within 8km, there are signs of some up-lift of around 1.6% for prices in postcodes where wind farms become operational, but are hidden, and the effect of visible turbines falls to zero. Given there was no detectable effect from non-visible wind farms within 4km, the up-

lift in prices is evidently within the 4-8km band (as shown in subsequent results). There are a number of possible interpretations of this price premium. Firstly there could be spurious effects due to non-random placement of wind farms although it seems unlikely that this this would show up specifically for non-visible wind farms at this radius. Secondly, there may be benefits to home owners within the 8km radius, offset by other costs at closer distances. Lastly, prices may be increased by displacement of demand from neighbouring areas where the turbines are visible. These displacement price effects are theoretically possible if buyers in these rural housing markets are relatively constrained in their choices (e.g. by family, jobs, search costs, other local amenities) and willing to pay more for housing in these localities without wind farm visibility rather than seek alternative housing in completely different non-wind farm locations. It is not possible to distinguish between these second and third hypotheses, but either way, the results for non-visible wind farms are reassuring in showing that the *negative* impacts from visible wind farms do not arise from a spurious association between price trends and the timing and location of wind farm development. Again, overall within the 14km, the regressions indicate no positive or negative effects associated with the timing of wind farm operations in the general local area.

All this evidence suggests that the estimated price reductions in postcodes where wind farms are visible are causally attributable to wind farm visibility. Later results will provide more detail on the pattern of distance decay of the wind farm price effects, and present some more formal difference-in-difference-in-difference estimates of the visibility impacts.

One concern could be that the price effects by distance and visibility status are the result of general spatial price trends, generated by other factors such as housing supply or opportunities in the labour market. Although the patterns in Table 3 are consistent with what we might expect theoretically from a causal effect of wind farms on prices, it is potentially possible that windfarms

just happen to always be opening in regions where prices are falling, relative to regions where windfarms already exist, and falling more in places close to new windfarms than in places further away. Since the sample is restricted to sales close to windfarms, and does not include any sales in the wider region beyond 14km from wind farms sites, there are limits to how flexibly the specifications can control for very general regional trends. However, Appendix Table 8 demonstrates that the patterns are similar after controlling for regional price trends, which are defined by splitting the sample (sales <14km from a windfarm) into four north, south, east and west quadrants and interacting quadrant indicators with year dummies. In this specification, comparison is being made between sales around wind farms opening in one location, and windfarms that exist or will exist in the same quadrant of England and Wales.

Table 4 and Table 5 present further assessments of the credibility of the findings by checking for spurious price trends and changes in the types of housing being sold as the wind farms become operational. The results shown are for the sample within the 4km radius, but the general picture is the same when the exercise is repeated at other distances. Table 4 presents a series of ‘balancing’ tests in which the dependent variable in the regressions of Table 3, column 8, is replaced by housing characteristics, and the housing characteristics are excluded from the set of regressors. The aim here is to see if there are within-postcode changes in the composition of the sample that coincide with the start of wind farm operations. Columns 1-6 use the few characteristics that are available in the Land Registry data set as the dependent variables. In column 7 the dependent variable in postcode quarter i,t is the cumulative sum of sales in postcode i up to period t and the regression provides a test for changes in the rate of transactions between the before and after operation periods. In the remaining columns, the dependent variables are postcode-by-year mean characteristics taken from an auxiliary dataset of transactions from the Nationwide building society and merged to the dataset. This dataset has far more information on housing

characteristics, but is only a sub-set of transactions, and hence postcodes, in the Land Registry data, therefore the sample size is much reduced. Looking across Table 4 it is evident that there are no statistically significant changes in the composition of housing transactions associated with wind farm operation, and there is no systematic pattern in the point estimates that would suggest that the price changes in Table 3 could be related to the sale of lower quality houses.

Table 5 carries out further robustness tests on the 4km sample, firstly adding in the Nationwide data set characteristics as control variables (column 2), and replacing the Land Registry prices with prices from the Nationwide (column 3). The coefficient estimates from the Nationwide sample are slightly larger than those from the Land Registry, although not by much relative to the standard errors, and changing the source of the price information does not make any difference. Column 4 adds in additional demographic characteristics from the 2001 Census (proportion not qualified, proportion tertiary qualified, proportion born in UK, proportion white ethnicity, proportion employed, proportion in social rented accommodation) interacted with linear time trend, but again this has no bearing on the results.

Column 5 shows a specification which controls for region-specific quarterly price index, based on prices in the ten standard regions of England and Wales. As noted above, it is not feasible to do this simply by including region-by-quarter dummies, because there are too few wind farms becoming operational in any region-quarter period. Instead, the region-quarter price indices are estimated a first stage postcode-fixed effects regression of log prices on region-quarter dummies in the full Land Registry price paid dataset⁹. The estimated region-quarter effects are then used as controls in the second stage estimation. Again this has no impact on the key result, even though

⁹ The sample is restricted to postcodes *beyond* the 14km wind-farm distance limit, otherwise the estimated price index would be mechanically endogenous in the price regressions based on the wind farm sample.

the region-quarter effects are strongly correlated with the prices close to the wind farms (the coefficient on the region-quarter effects is 1.059, with a standard error of 0.030).

Column 6 does something similar, but controlling for predicted pre-operational and post-operational linear price trends in the area defined by the set of postcodes that share the same nearest operational wind farm within 4km. Again it is not practical to simply include nearest-wind-farm specific trend variables, since the price changes in response to wind-farm operation are not sharp enough to successfully identify these separately from wind-farm specific price trends over the whole period. Instead, similarly to the region-quarter trends, the pre-operation and post-operation wind farm price trends are estimated in a first stage regression of prices wind farm-specific time trends using observations for the pre-operation or post-operation period only. The first stage regression predictions of the wind farms specific price trends from the pre-operation period are then extrapolated over the whole sample period and included as controls in the second stage regression. Controlling for pre and post operation price trends in this way yields a slightly bigger coefficient on visible wind farms, suggesting that the baseline estimates in Table 3 are, if anything, conservative. This is consistent with post-announcement, pre operation downward price trends, which will reduce the pre-post operation average price difference and attenuate the basic within-groups fixed effects estimates of Table 3.

Column 8 and 9 also test for robustness to other regional price drivers. Column 8 controls for differences in new housing supply across space. Highly geographically detailed data on housing supply is not available in England or Wales, but Column 8 uses the best information available and controls the number of housing construction starts (in logs) in each of the ten standard Regions in England and Wales in each year. Column 9 includes labour market variables, namely mean wages

and jobs per capita at the Local Authority, County or Region level (there are 348 Local Authorities, and 42 Counties in England and Wales).¹⁰

Many windfarms are close to the coastline of England and Wales, so there is an outside chance that the results could be influenced by coastline visibility, given that coastline visibility is presumably a desirable amenity. To check this, Column 9 controls for trends associated with coastal views. A coastal view-shed was constructed for places within 14km of the coastline, and sales categorised in quintiles of coast line visibility. The specification in Column 9 includes an interaction of top quintile coastline visibility with year dummies. Evidently, differences in coast views do not explain the estimated effects of wind farms on prices, although the unreported coefficients on the coast-view-by-year dummies indicate differential trends in coastal locations.

Overall, there is no evidence from Table 4 and Table 5 that the finding of negative impacts from wind farms on prices arises from omitted variables or unobserved price trends.

More detail on distance-decay of the wind farm price effects and the differences in the effects of visible and non-visible wind farms within the 14km limit is provided in Table 6. In this specification, estimation is from postcodes with transactions within 14km of a site, and treatment indicators for the different distance bands are included in a single regression. The coefficients indicate the effects at each distance band within this 14km radius. The estimation includes postcodes with or without wind farm visibility. The results are broadly in line with the alternative presentation in Table 3, but there are some subtle differences. These differences arise because the coefficients on the housing control variables, quarter dummies and terrain-by-year trends are estimated from the full 14km radius sample. This specification also constrains postcodes within

¹⁰ In the vast majority of cases Local Authority variables are used, but these are not always published for Local Authorities due to small sample sizes, in which case higher level geography is used.

each wind-farm distance band to be on the same general price trend in the absence of any effects due to wind farm operation and visibility (through distance-band-by-year interactions). At the same time the specification allows for differences in general price trends between postcodes with potential wind farm visibility and those without for the whole 14km radius circle (through visibility-by-quarter interactions).

Looking at Table 6, the price effect for visible turbines within 1km, and at 1-2km is around 5.5-6%. This falls quite sharply in the 2-4km distance band, to just under 2%. Beyond this there are price effects from visible turbines right out as far as 14km, although these are small at around 0.5-1%. The results in the next section show that these effects at greater distances are associated with the largest wind farms only. In contrast, the coefficients on non-visible turbines are generally positive, but small and non-significant except in the 4-8km band. Note that the coefficients on non-visible turbines look comparable in magnitude but opposite in sign to the effects of visible turbines in the 4-8km band, which might suggest some aggregate net gains in terms of total housing values. However, it should be borne in mind that only 35% of postcodes within 8km of a wind farm do not have views of the wind farms, so a much smaller share of transactions see price gains rather than price losses. The impacts of wind farms 8-14km away, where the wind farms are not visible, is, as expected, zero and insignificant.

Potential theoretical reasons for these positive effects associated with proximity to turbines where the turbines are hidden were discussed in relation to Table 3. A corollary is that the coefficients on the wind farm visibility indicators, while showing the house price changes, underestimate the value of the visual dis-amenity of wind farms. As discussed in Section 4.2 a difference-in-difference-in-difference estimate based on the difference between the coefficients on visible turbines and non-visible turbines at each distance band provides a cleaner estimate of the

willingness to pay to avoid views of wind farms. These estimates are shown in the bottom panel of Table 6. These are calculated from the coefficients and the variance-covariance matrix of the coefficients in Table 6. Given the small positive coefficients associated with non-visible wind farms, the basic price effects estimated from the visible-operational treatment dummies underestimate the marginal willingness to pay to avoid the visual dis-amenity and the difference-in-difference estimates are slightly larger in magnitude. Within 2km, the visual impact of wind farms is has an implied cost of around 8.5% of housing prices, between 2km and 8km the figure falls to around 3.5%, whilst beyond 8km there is virtually no impact (just under 0.7%).

5.3 *Further results on wind farm size.*

The results so far have looked simply at turbine development as a binary treatment effect, and have ignored the scale of the wind farm. Table 7 provides a more comprehensive analysis that investigates whether there is a greater cost associated with larger developments with more turbines, and over what distance. The setup is basically the same as in Table 6, but with interactions between dummies for wind farm size and distance. Again, the lower panel of the table reports difference-in-difference-in-difference estimates of the price differentials associated with visibility for each distance band and wind farm size group. Figure 4 illustrates the patterns in Table 7 by plotting the coefficients against the mid points of the distance bands. The results are in line with what would be expected if the price impacts are related to the dis-amenity of wind farm visibility. Bigger wind farms have a bigger impact on prices at all distances. A wind farm with 20+ turbines within 2km reduces prices by some 12% on average, and the implied effect of the visual dis-amenity is around 15%. Note though that there is a relatively small number of transactions within 2km of the centroid of a 20+ turbine wind farm (988) and given the geographical spread of the turbine array, this price effect could also relate to noise and visual flicker problems. However, even at 8-14km there is a 4.5% reduction in prices associated with large visible operational wind

farms, and the willingness to pay to avoid visibility is 6.5%. Even at 8km there is some negative impact of the large wind farms, and all of this is attributable to visibility. Medium size wind farms above average size also have strong effects throughout the distance range up to 8km, but no effect after that. The effect of smaller wind farms with less than 1-10 turbines is, as might be expected, concentrated in the first 2km where there is a 5% reduction in prices. This falls to just over 1.5% at 4km and becomes zero and insignificant beyond that, although there is an implied visibility cost in the 4-8km range due to the lift in prices of houses in the 4-8km range where turbines are not visible. All in all, the results in Table 7 and their visualisation in Figure 4 are entirely consistent with theoretical reasoning about the potential visual impacts of wind farms, and the differences across wind farm size and distance band provide reassurance that the effects are genuinely causal and not spurious.

One concern in any spatial estimation design with multiple interventions on grouped observation (wind farm developments affecting groups of neighbouring houses in this case) is the estimation of the standard errors (Moulton 1990, Conley and Taber 2011). All specifications so far allowed for serial and spatial correlation (and heteroscedasticity) in unobservable factors within neighbouring groups of postcodes defined by Census Output Areas, using clustered standard errors at this level. These standard errors may be biased by more general spatial autocorrelation in the unobservables, between Census Output Area groups. Tests on the regression residuals fail to find evidence of this spatial autocorrelation. Moran's I statistics based on the residuals have values of less than 0.001 (on a theoretical range of -1/+1), and the p-value for the test of the null of no spatial autocorrelation is 0.5 or higher.¹¹ Nevertheless, some alternative standard errors allowing for more general spatial

¹¹ Moran's I statistics are estimates of $\text{Cov}(m(x), x)/(\text{Var}(x))$ where $m(x)$ is an average of x over neighbouring observations and neighbours are defined by spatial weights. Tests were performed using inverse distance weights, and average of observations within 4km.

autocorrelation are shown for the final specification in Appendix Table A1. Standard errors using the double clustering method of Thompson (2011), allowing for serial correlation within postcodes, and cross sectional correlation within quarters, are similar to those obtained from clustering at Census OA level. Standard errors with clustering on Census Wards yields larger standard errors and lower levels of significance, although the pattern remains the same, with statistically significant coefficients for small wind farms up to 4km, and statistically significant impacts from large wind farms throughout the distance range. Standard errors clustered on nearest wind farm groups (not reported) yield similar results to the ward-based clustered standard errors.

6 Conclusions

The analysis in this paper provided estimates of the effects of wind farm visibility on housing prices in England and Wales. The fairly crowded geographical setting, with numerous wind farms developed within sight of residential property, provides a unique opportunity to examine the visual impacts of wind farms through hedonic property value methods. The analysis used a micro-aggregated postcode-by-quarter panel of housing transactions spanning 12 years, and estimated difference-in-difference effects using a quasi-experimental, postcode fixed effects methodology. Comparisons were made between house price changes occurring in postcodes where nearby wind farms become operational and visible, with the price changes occurring where nearby wind farms become operational but are hidden from view. All the results point in the same direction. Wind farms reduce house prices in postcodes where the turbines are visible, and reduce prices relative to postcodes close to wind farms where the wind farms are not visible. Averaging over wind farms of all sizes, this price reduction is around 5-6% within 2km, falling to less than 2% between 2 and 4km, and less than 1% by 14km which is at the limit of likely visibility. As might be expected, small wind farms have no impact beyond 4km, whereas the largest wind farms (20+ turbines) reduce prices by 12% within 2km, and reduce prices by small amounts right out to 14k (by around 1.5%).

There are small (~2%) increases in neighbouring prices where the wind farms are not visible, although these are only statistically significant in the 4-8km band. This price uplift may indicate some local benefits from wind farms, for example due to spillovers from rents to landowners from wind farm operation, or from community grants. However these price increases could also be explained by displacement of demand by those seeking housing in these areas towards places where the wind farms are hidden. These offsetting price effects in neighbouring places where wind farms are visible and where they are not may explain, in part, why previous studies that focus only on distance to wind farms fail to find significant effects.

These headline findings are comparable to the effects of coal power plants in the US found in Davis (2011) who finds up to 7% reduction within 2 miles (3.2 km). Of course, it takes many geographically dispersed wind farms to generate the same power as a single coal (or nuclear) plant, so the aggregate effects of wind farms and the number of households affected by their visual impact is likely to be considerably larger. The results are also in line with existing literature that suggests that other tall power infrastructure has negative impacts on prices (e.g. high voltage power lines, Sims and Dent 2005). The point estimates are comparable to the repeat sales estimates of the effects of wind farms in in Lang et al (2014) for Rhode Island, although their estimates are not statistically significant.

The paper presents a number of robustness tests, but even so the findings should be interpreted with some caution. The information on wind farm location and visibility is limited by lack of data on the precise location of individual turbines, so the classification of postcodes in terms of visibility is subject to measurement error. This is most likely to result in some attenuation of the estimated effects. Steps were taken to minimise this problem by eliminating postcodes where visibility is ambiguous. More importantly, there is no historical information on the timing of events leading up

to wind farm operation (announcement, approval, construction etc.) so the price effects reported here relate to the average difference between the post-operation and pre-operation periods for the periods spanned by the data (a gap of just under 6 years). However, the wind farm development cycle can last a number of years, and price changes evolve fairly slowly over time in response to events. Again the most likely consequence of this is that the results underestimate the full impact between the pre-announcement and post-construction phase. It should also be noted that the estimates of turbine visibility, may pick up some effects from turbine noise – especially close to large windfarms, if terrain that hides the windfarms also attenuates the noise. However, noise levels at the distances beyond 1km at which the visible/non-visible comparisons are made are likely to be very low.

Well established theories (Rosen 1974) suggest that we can interpret price differentials emerging between places where wind farms are visible and comparable places where they are not, as household marginal willingness to pay to avoid the dis-amenity associated with wind farm visibility (though Kuminoff and Pope, 2014, has recently highlighted some potential pitfalls in interpreting difference-in-difference estimates in this way). If we take the figures in the current paper seriously as estimates of the mean willingness to pay to avoid wind farms in communities exposed to their development, the implied costs are quite substantial. For example, a household would be willing to pay around £600 per year to avoid having a wind farm of small-average size visible within 2km, around £1000 to avoid a large wind farm visible at that distance and around £125 per year to avoid having a large wind farm visible in the 8-14km range.¹² The implied amounts required per wind farm to compensate households for their loss of visual amenities is

¹² These figures is based on an average house price of £145,000 (in 2010), a the visible-non-visible price differentials from Table 7 and a 5% interest rate.

therefore fairly large: about £14 million on average to compensate households within 4km.¹³ The corresponding values for large wind farms will be much higher than this, as their impact is larger and spreads out over much greater distances.

These per-household figures are somewhat higher than the highest estimates from the stated preference literature, although there are no directly comparable figures. The figures cited in Bassi, Bowen and Fankhauser (2012) are typically much less than £100 per year, though this is per individual, so household willingness to pay could be higher.

The findings of the paper are relevant on a number of policy levels. The estimates provide potential inputs into cost-benefit analyses related to the siting of wind turbines, and the net benefits of wind power relative to other forms of low carbon energy. It should be noted, however, that the price effects reflect the valuation of home buyers in locations where wind farms are visible, so may not represent the mean valuation of wind farm visibility in the general population. The estimates could also inform policy on compensation for home owners for the loss of value in their homes arising from views of new wind farms. Interestingly, the evident increase in value of for houses where local wind farms are out of site suggests some scope, at least in theory, for these 'winners' to compensate the 'losers' in places where the turbines are visible e.g. through adjusting council taxes or introducing property value taxes.

References

Bakker, R.H., E. Pedersen, G.P. van den Berg, R.E. Stewart, W. Lok, J. Bouma (2012) Effects of

¹³ Based on: around 1.8% of postcodes within 4km of a visible turbine; the number of households in England and Wales is 23.4 million; the capitalised effect of visibility within 4km is 3.5% on average; an average house price is £145000; and the number of operational turbines is 148.

- wind turbine sound on health and psychological distress, *Science of the Total Environment*, 425 42–51.
- Bell, Derek, Tim Gray & Claire Haggett (2005), The ‘Social Gap’ in Wind Farm Siting Decisions: Explanations and Policy Responses *Environmental Politics*, 14 (4) 460 – 477
- Conley, Timothy G. and Christopher R. Taber (2011) Inference with “Difference in Differences” With a Small Number of Policy Changes, *Review of Economics and Statistics*, 93(1): 113–125
- DECC (2013a) Digest of United Kingdom Energy Statistics, Chapter 6: Renewable sources of energy, Department of Energy & Climate Change, London
- DECC (2013b) Digest of United Kingdom Energy Statistics, Internet Booklet, https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/225056/DUKES_2013_internet_booklet.pdf, accessed November 2013
- European Commission (2006) Energy Technologies: Knowledge – Perception – Measures, European Union Community Research Report, Luxembourg: Office for Official Publications of the European Communities
- Farbouda, A., R. Crunkhorna and A. Trindadea (2013), ‘Wind turbine syndrome’: fact or fiction? *The Journal of Laryngology & Otology*, 127 (03) 222-226
- Hoen, Ben, Jason P. Brown, Thomas Jackson, Ryan Wiser, Mark Thayer and Peter Cappers (2013) A Spatial Hedonic Analysis of the Effects of Wind Energy Facilities on Surrounding Property Values in the United States, *Environmental Energy Technologies*

Hoen, Ben, Ryan Wisser, Peter Cappers, Mark Thayer, and Gautam Sethi (2011), Wind Energy Facilities and Residential Properties: The Effect of Proximity and View on Sales Prices, Journal of Real Estate Research

Kahn, Matthew E. (2013) Local non-market quality of life dynamics in newwind farms communities, Energy Policy 59 (2013) 800–807

Kuminoff, Nicolai V., Christopher F. Parmeter and Jaren C. Pope (2010) Which hedonic models can we trust to recover the marginal willingness to pay for environmental amenities? Journal of Environmental Economics and Management. 60 (3) 145–160

Kuminoff, Nicolai and Jaren C. Pope (2014) Do “capitalization effects” for public goods reveal the public’s willingness to pay? International Economic Review, 55 (4) 1227-1250

Lang, Corey, James J. Opaluch, and George Sfinarolakis (2014) The Windy City: Property Value Impacts of Wind Turbines in an Urban Setting, unpublished manuscript, Department of Environmental and Natural Resource Economics, University of Rhode Island

Moulton, Brett (1990) An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units, The Review of Economics and Statistics Vol. 72, No. 2, May, 1990

Munday, Max, Gill Bristow, Richard Cowell, Wind farms in rural areas: How far do community benefits from wind farms represent a local economic development

- opportunity? *Journal of Rural Studies* 27 1-12
- Roback , Jennifer (1982), Wages, Rents, and the Quality of Life, *The Journal of Political Economy*, 90 (6) 1257-1278.
- Rosen, Sherwin (1974) Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition, *The Journal of Political Economy*, 82 (1) 34-55
- Samuela Bassi, Alex Bowen and Sam Fankhauser (2012) The case for and against onshore wind energy in the UK, *Grantham Research Institute on Climate Change and Environment Policy Brief*, London
- Sims, Sally & Dent, Peter (2005) High voltage overhead power lines and property values: a residential study in the UK, *Urban Studies* 42 (4) 665 – 94
- Sims, Sally and Peter Dent (2007) Property stigma: wind farms are just the latest fashion *Journal of Property Investment & Finance* 25 (6) 626-651
- Sims, Sally, Peter Dent & G. Reza Oskrochi (2008) Modelling the impact of wind farms on house prices in the UK, *International Journal of Strategic Property Management*, 12:4, 251-269
- Thompson, Samuel (2011) Simple Formulas for Standard Errors that Cluster by Both Firm and Time, *Journal of Financial Economics*, 99 1–10
- University of Newcastle (2002) Visual Assessment of Windfarms Best Practice. Scottish Natural Heritage Commissioned Report F01AA303A.
- Vidal, John (2012) Wind turbines bring in 'risk-free' millions for rich landowners, *Guardian*

28th February online <http://www.theguardian.com/environment/2012/feb/28/windfarms-risk-free-millions-for-landowners>, accessed November 2013

Vyn, Richard J. & Ryan M. McCullough (2014) The Effects of Wind Turbines on Property Values in Ontario: Does Public Perception Match Empirical Evidence? Canadian Journal of Agricultural Economics 62 (3) 365–392

YouGov (2010) Wind Power, survey conducted on behalf of Scottish Power http://www.scottishrenewables.com/media/uploads/results_for_scottish_renewables%28wind_power%29.pdf accessed March 2014

Table 1: Operational windfarm summary data, 1992-2011 England and Wales

	Mean	s.d.	Min	Max
148 wind farms				
Turbines mean	11.2	15.4	1	103
Turbines median	6			
MW capacity	18.6	39.2	.22	300
Height to tip	90.9	29.2	42	150

Landcover of non-urban/industrial wind farms: Offshore 14; Forest 8; Farmland 82; Moorland and mountain 39; Coastal 5.

Table 2: Main estimation sample summary data, 2000-2011 England and Wales

	Mean	s.d.	Obs
Sales in postcodes operational turbine at some time 2000-2011 within 14km			
Log price	11.56	0.674	1710293
New build	0.041	0.192	1710293
Detached house	0.250	0.423	1710293
Semi-detached house	0.070	0.249	1710293
Terraced house	0.320	0.452	1710293
Flat/Maisonette	0.361	0.469	1710293
Freehold	0.849	0.351	1710293
Postcodes within 1km of wind farm, 2000-2011			1142
Where visible			1125
Postcodes within 2km of wind farm, 2000-2011			5350
Where visible			5062
Postcodes within 4km of wind farm, 2000-2011			20838
Where visible			17031
Postcodes within 8km of wind farm, 2000-2011			81820
Where visible			52980
Postcodes within 14km of wind farm, 2000-2011			220669
Where visible			123892
Time between post-pre sales in same postcode (quarters)			
Visible within 1km	23.335	5.016	468
Visible within 2km	23.379	6.189	2004
Visible within 4km	23.297	6.170	7348
Visible within 8km	23.047	6.150	24408
Visible within 14km	23.148	6.131	59852

Table 3: Postcode fixed effects estimates; samples with operational wind farm within k km, during 2000-2011

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Radius	<1km	<1km	<2km	<2km	<2km	<4km	<4km	<4km	<8km	<8km	<8km	<14km	<14km	<14km
Control vars.	No	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Visible-operational:	-0.0632*** (0.0171)	-0.0666** (0.0221)	-0.0628*** (0.0095)	-0.0554*** (0.0095)	-0.0558*** (0.0095)	-0.0300*** (0.0057)	-0.0267*** (0.0055)	-0.0244*** (0.0054)	-0.0144*** (0.0032)	-0.0046 (0.0029)	-0.0035 (0.0029)	-0.0048* (0.0019)	-0.0018 (0.0018)	-0.0027 (0.0017)
Non-visible-operational:					-0.0611 (0.0609)			-0.0018 (0.0125)			0.0165*** (0.0041)			-0.0024 (0.0020)
Obs	8,052	8,052	36,298	36,298	37,998	125,619	125,619	150,907	417,108	417,107	621,395	984,294	984,292	1,710,293
R-squared	0.8141	0.8459	0.8284	0.8580	0.8601	0.8377	0.8626	0.8642	0.8487	0.8719	0.8736	0.8461	0.8706	0.8718

Robust standard errors in parentheses, clustered at Census OA *** p<0.001, ** p<0.01, * p<0.05

Data in postcode-quarter cells, 2000-2011. Dependent variable is postcode-quarter-mean log prices.

Visible-operational is the treatment indicator (visible, 0<distance<k, operational) described in Section 4, indicating that a postcode has an operational windfarm visible within the specified radius k.

Non-visible operational is the treatment indicator (non-visible, 0<distance<k, operational) described in Section 4, indicating that a postcode has an operational windfarm within the specified radius k, but this is not likely to be visible.

Sample restricted to postcodes with visible-operational turbines within distance k at some time over the study period in columns 1,2,3,4,6,7,9,10,12,13.

Sample restricted to postcodes with visible-operational or non-visible-operational turbines within distance k at some time over the study period in columns 1,2,3,4,6,7,9,10,12,13.

Control variables in columns 1,2,3,4,6,7,9,10,12,13 are postcode slope-by-year, elevation-by-year, aspect by-year dummies, proportions of sales of detached, semi-detached, terraced, flat/maisonette.

Control variables in columns 5,8,11,14 are postcode slope-by-year, elevation-by-year, aspect by-year dummies, proportions of sales of detached, semi-detached, terraced, flat/maisonette, plus dummy groups for distance-band-by-year, and visible/non-visible-by-quarter trends.

All regressions control for quarter dummies.

Table 4: Balancing tests for various housing characteristics. 4km radius

	(1) New	(2) Detached	(3) Semi	(4) Terraced	(5) Flat	(6) Leasehold	(7) Yrly. Sales	(8) Floor area	(9) Beds	(11) Baths
Visible-operational:	-0.0036 (0.0059)	0.0011 (0.0038)	-0.0001 (0.0016)	-0.0071 (0.0044)	0.0061 (0.0038)	0.0039 (0.0021)	0.0050 (0.0074)	-0.1267 (2.0573)	-0.0636 (0.0464)	0.0509 (0.0450)
Non-visible-operational:	-0.0043 (0.0069)	-0.0128 (0.0078)	-0.0043 (0.0038)	0.0099 (0.0094)	0.0072 (0.0090)	-0.0080 (0.0052)	0.0005 (0.0157)	0.4270 (4.9602)	0.0193 (0.1198)	-0.0705 (0.1018)
Number of observations	150,907	150,907	150,907	150,907	150,907	150,907	150,907	17,931	17,931	17,931

	(11) No CH	(12) No Gar	(13) Detached	(14) Semi	(15) Terraced	(16) PB Flat	(17) Conv Fl	(18) Other	(19) Age
Visible-operational:	-0.0105 (0.0154)	-0.0178 (0.0304)	-0.0287 (0.0235)	0.0207 (0.0282)	-0.0004 (0.0242)	0.0108 (0.0150)	0.0014 (0.0092)	-0.0038 (0.0051)	-0.5586 (1.7077)
Non-visible-operational:	-0.0838 (0.0612)	0.0212 (0.0780)	0.0324 (0.0943)	-0.0575 (0.1090)	-0.0330 (0.0733)	0.0423 (0.0364)	0.0109 (0.0203)	0.0048 (0.0060)	-0.2947 (4.6216)
Number of observations	17,212	17,931	17,931	17,931	17,931	17,931	17,931	17,931	17,931

Specifications as in Table 3, column 8, but with property type control variables excluded.

Columns 8-19 based on sub-sample with transactions from Nationwide sales database.

Table reports coefficients, standard errors (clustered on OA) and sample size

Table 5: Robustness to additional control variables and trends. 4km radius

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Baseline estimate from Table 3	Sub-sample with additional Nationwide property Xs	Nationwide prices and Xs	Census output area Xs x trends	Control for regional price index from full dataset	Windfarm specific pre and post trends	Control for regional house construction starts	Control for Local Authority wages and jobs per capita	Control for coast view-by-year dummies
Visible-operational:	-0.0244*** (0.0054)	-0.0452** (0.0146)	-0.0419*** (0.0120)	-0.0260*** (0.0054)	-0.0206*** (0.0048)	-0.0326*** (0.0054)	-0.0232*** (0.0054)	-0.0194*** (0.0053)	-0.0263*** (0.0054)
Non-visible-operational:	-0.0018 (0.0125)	0.0220 (0.0608)	0.0298 (0.0356)	-0.0123 (0.0133)	0.0049 (0.0114)	-0.0016 (0.0122)	0.0105 (0.0120)	0.0170 (0.0120)	0.0119 (0.0122)
Observations	150,907	17,212	17,212	136,031	150,907	150,907	150,907	150,907	150,907

Robust standard errors in parentheses, clustered at Census OA *** p<0.001, ** p<0.01, * p<0.05

Column 2 controls for floor size, number of bedrooms, bathrooms, central heating type, garage type, and detailed property type for postcodes represented in Nationwide data. Column 3 similar, using price reported in Nationwide data. Column 3 adds linear trends interacted with census 2001 variables at output area (OA) level (OA land area, proportion with no qualifications, proportion with tertiary qualifications, proportion born UK, proportion white ethnicity, proportion employed, proportion in social rented housing).

Column 5 controls for piecewise constant quarterly price index estimated from transactions beyond 14km from any operational windfarm.

Column 6 controls for nearest operational windfarm linear time trends estimated from pre-operational and post-operational periods.

Column 7 includes control for region-by-year private housing construction starts from Department of Communities and Local Government housing statistics

Column 8 includes control for Local Authority-by-year wages and job density from Annual Survey of Hours and Earnings and Office for National Statistics job density data (from www.nomisweb.co.uk)

Column 9 includes coast-view by year dummies, where coast-view is an indicator that property is within 14km of the coastline and the coastline visibility is in the top 20% based on the number of coast outline vertices from which the property is visible.

Specifications otherwise as Table 3, column 8,

Table 6: Postcode fixed effects estimates; distance bands; sample with operational wind farm within 14km, during 2000-2011

	(1) <1km	(3) 1-2km	(3) 2-4km	(4) 4-8km	(5) 8-14km
Turbines visible	-0.0539*** (0.0164)	-0.0578*** (0.0092)	-0.0193*** (0.0052)	-0.0104*** (0.0028)	-0.0050** (0.0019)
No turbines visible	-	0.0268 (0.0498)	0.0152 (0.0105)	0.0223*** (0.0040)	0.0018 (0.0021)
Difference-in-difference-in-difference estimates relative to non-visible					
	-	-0.0847† 0.0501	-0.0345** 0.0106	-0.0327*** 0.0046	-0.0068* 0.0027

Notes as for Table 3, column 8, but with additional wind farm distance indicator

Observations 1710293, R-squared 0.8719

Robust standard errors in parentheses, clustered at Census OA *** p<0.001, ** p<0.01, * p<0.05, † p<0.10

Table 7: Effects by windfarm size and distance bands

	(1)	(2)	(3)	(4)
	<2km	2-4km	4-8km	8-14km
No turbines visible	0.0276 (0.0498)	0.0154 (0.0105)	0.0217*** (0.0040)	0.0015 (0.0021)
1-10 turbines visible	-0.0556*** (0.0084)	-0.0165** (0.0053)	-0.0032 (0.0030)	-0.0023 (0.0021)
11-20 turbines visible	-0.0512** (0.0187)	-0.0213* (0.0091)	-0.0371*** (0.0055)	-0.0013 (0.0035)
20+ turbines visible	-0.1199*** (0.0277)	-0.0530** (0.0169)	-0.0466*** (0.0059)	-0.0162*** (0.0029)
Obs. 1,710,293. R-squared 0.8719				
	(5)	(6)	(7)	(8)
	<2km	2-4km	4-8km	8-14km
Difference-in-difference-in-difference estimates relative to non-visible				
1-10 turbines visible	-0.0832† 0.0501	-0.0319** 0.0107	-0.0249*** 0.0048	0.0038 0.0029
11-20 turbines visible	-0.0789 0.0527	-0.0368** 0.0128	-0.0588*** 0.0066	0.0027 0.0039
20+ turbines visible	-0.1475** 0.0560	-0.0685** 0.0192	-0.0684*** 0.0069	-0.0177*** 0.0035

Notes as for Table 3, column 8, but with additional turbine size indicators

Robust standard errors in parentheses, clustered at Census OA *** p<0.001, ** p<0.01, * p<0.05, † p<0.10

Figure 1: Development of wind turbines in England and Wales, 1993-2011


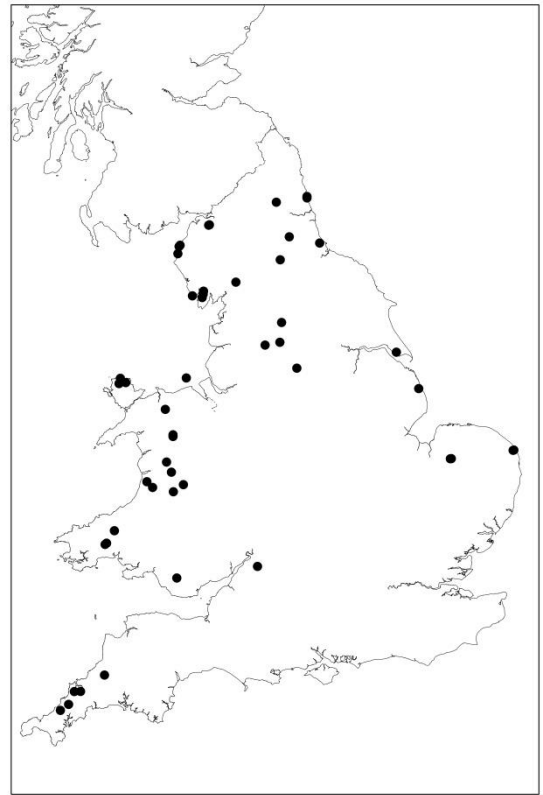
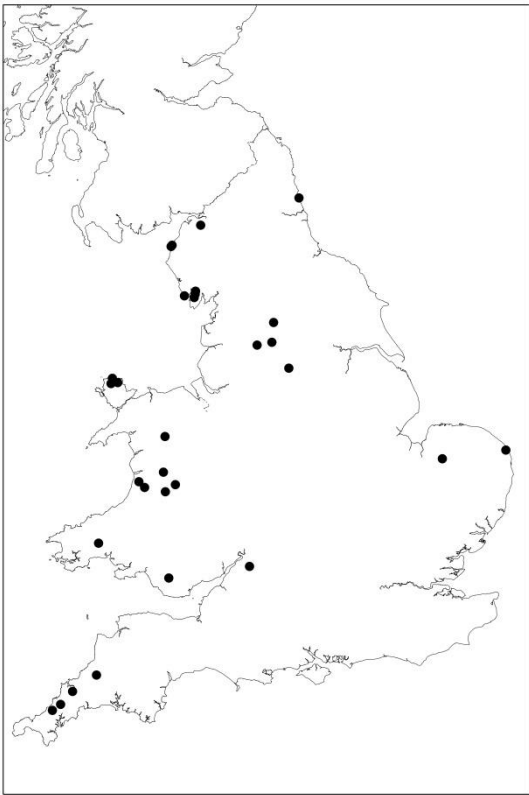


Figure includes onshore and offshore wind farms which are closer than 16km to postcodes with housing transactions

Figure 2: Development of wind turbine sites in England and Wales
2000: 30 sites
2003: +20 sites



2007: +33 sites

2011: +65 sites

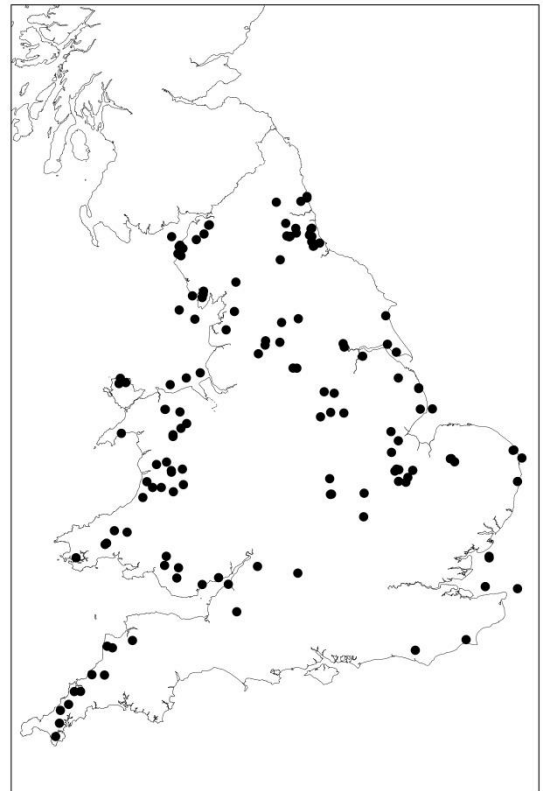
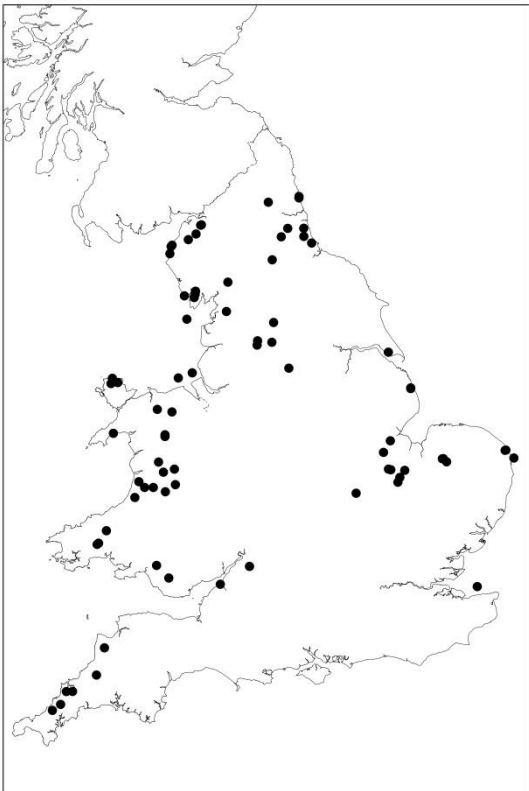



Figure 3: Example viewshed. Haswell Moor wind farm in north east England



Figure 4: Comparison by visibility: Postcode fixed effects estimates; distance bands; controls include distance-band-by-year effects and visibility-by-quarter effects.



7 Appendix

Table 8: Postcode fixed effects estimates; samples with operational wind farm within k km, during 2000-2011; additional controls for region-by-year effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Radius	<1km	<2km	<2km	<4km	<4km	<8km	<8km	<14km	<14km
Control vars.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Visible- operational:	-0.0663*** (0.0183)	-0.0438*** (0.0088)	-0.0455*** (0.0088)	-0.0306*** (0.0049)	-0.0301*** (0.0048)	-0.0105*** (0.0026)	-0.0089*** (0.0026)	-0.0117*** (0.0016)	-0.0105*** (0.0016)
Non-visible- operational:			-0.0696 (0.0595)		-0.0110 (0.0122)		0.0057 (0.0040)		-0.0074*** (0.0019)
Sample	Visible	Visible	All	Visible	All	Visible	All	Visible	All
Obs	8,052	36,298	37,998	125,619	150,907	417,107	621,395	984,292	1,710,293
R-squared	0.8505	0.8615	0.8632	0.8660	0.8666	0.8753	0.8756	0.8735	0.8742

Notes as in Table 3

Additional controls for regional trends: north, south, east and west quadrant-by year dummies

Table 9: Alternative standard errors: windfarm size and distance bands

	(1) <2km	(2) 2-4km	(3) 4-8km	(4) 8-14km
No turbines visible	0.0272 (0.0498) [0.0486] {0.0539}	0.0153 (0.0105) [0.0121] {0.0159}	0.0215* (0.0040) [0.0058] {0.0084}	0.0013 (0.0021) [0.0035] {0.0046}
1-10 turbines visible	-0.0557** (0.0084) [0.0084] {0.017}	-0.0168† (0.0053) [0.0053] {0.0101}	-0.0032 (0.0030) [0.0047] {0.0069}	-0.0022 (0.0021) [0.0033] {0.0052}
11-20 turbines visible	-0.0517† (0.0187) [0.0253] {0.0269}	-0.0217 (0.0091) [0.0102] {0.0141}	-0.0373** (0.0055) [0.0083] {0.0127}	-0.0013 (0.0035) [0.0052] {0.0073}
20+ turbines visible	-0.1207*** (0.0275) [0.0287] {0.0201}	-0.0531† (0.0169) [0.0140] {0.0290}	-0.0467*** (0.0059) [0.0073] {0.0115}	-0.0161* (0.0029) [0.0036] {0.0067}

Obs. 1,710,293. R-squared 0.8718

Notes as for Table 3, column 8, but with additional turbine size indicators

Robust standard errors in parentheses, clustered at Census OA (.), double clustering at postcode and quarter following Thompson 2011 [., ward {.]

Significance indicated for most conservative ward-clustered standard errors *** p<0.001, ** p<0.01, * p<0.05, † p>0.10