University of Strathclyde
Department of Economics

Renewable Energy and the Housing Market

Presented in fulfilment of the requirements for the degree of Doctor of Philosophy

2021

Orion Meyer

June 30, 2021

Declaration of Authenticity and Author's Rights

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

Signed:

Date: June 30, 2021

Contents

List of Figures

List of Tables

Acknowledgments

I would like to extend my sincere thanks to all those who have helped me through the process of completing my PhD thesis.

First and foremost, I am grateful to my supervisors Grant Allan and Stuart McIntyre. Over the past five years Grant and Stuart have helped me to revise and refine research ideas, edit draft chapters and given valuable advice. They have always gone the extra mile to help me develop professionally and personally. Without their continuous support, encouragement and patience, completing this thesis would not have been possible.

Special thanks are due to my PhD colleagues and friends Ayse, Ben, Chirsty, Daniel, Ross, and Sharada for their company and conversation. I would also like to thank my close friends and coffee morning companions Edward and Kayla for their support and companionship. From pizzas at Paesano to playing games over zoom during lock-down, the last five years have been all the better with you all in them. I am grateful to the Strathclyde Economics Department for their help and encouragement throughout the PhD.

I would like to thank the University of Strathclyde for providing funding for my PhD. I am also greatly indebted to Ordnance Survey, Digimaps, RenewableUK and Simon Mallet, who have granted me free access to the data required to complete the empirical work within this thesis.

Finally, special thanks to my parents Pamela and Daniel, and my brothers Isaac and Noah, who have kept me laughing and supported me throughout the PhD.

Synopsis

Renewable energy sources now provide 27% of global electricity generation, an increase of nearly 38% since 2010. This trend is expected to not only continue, but to accelerate over the coming decades as nations set out domestic policies and international agreements to curb emissions (IEA 2020*b*). Achieving the ambitious emissions commitments within the Paris Agreement (BEIS 2020*b*) as well even more ambitious regional targets (ASP 15 2019, WGCCC 2021, PMO 2021) require the deployment of both large and small-scale renewable generation technologies. While some of the infrastructure and capacity necessary to achieve these goals will be located far from populous areas, developments such as onshore wind are increasingly being sited near more densely populated residential areas. At the same time, there is a rapidly growing deployment of residential solar generation installed directly on homes.

This thesis contributes to the literature on the externalities associated with renewable energy developments through an assessment of house price trends arising from both large-scale commercial windfarms and residential solar photovoltaic (PV) systems. The research area is England and Wales, places experiencing a rapid deployment of both commercial and domestic renewable energy technologies. Commercial developments are those which export electricity directly to the grid and domestic are primarily used to generate electricity to be used on-site in residential homes. The aim of this thesis is to apply econometric methods to generate new insights into how these developments impact upon house prices. It also sheds light on how the assumptions made and analytical approaches influence the results of such analyses.

There is a substantial literature examining the effects of proximity to and visibility of wind energy developments on house prices (Hoen et al. 2011, Brown et al. 2012, Heintzelman & Tuttle 2012, Jensen et al. 2014, Lang et al. 2014, McCarthy & Balli 2014, Vyn & McCullough 2014, Gibbons 2015, Hoen et al. 2015, Dröes & Koster 2016, Heblich et al. 2016, Hoen & Atkinson-Palombo 2016, Sunak & Madlener 2016, 2017, Jensen et al. 2018). This thesis presents original research which extends this literature in a number of ways. In Chapter 3, I extend this literature by replicating a leading paper in the literature (Gibbons 2015). Here I apply an average sales analysis, using the average postcode transaction price to estimate effects, and test the robustness of modifications to the assumptions underpinning the analysis. In Chapter 4, I take the same dataset from Chapter 3 to analyze the effects of windfarm proximity and visibility through the use of a repeat sales analysis, where changes in transaction prices of properties which sell multiple times over the study period are used to create an alternative estimation of price effects from windfarm siting. Chapters 3 and 4 together provide a comparison between the results under average price and repeat sales respectively.

Chapter 5 contributes to the nascent literature examining the capitalization of residential PV systems into house prices (Dastrup et al. 2012, Hoen et al. 2013, Wee 2016, Ma et al. 2016, Qiu et al. 2017, Lan et al. 2020). This is the first analysis using data from the housing markets of England and Wales to estimate whether residential solar panels are capitalized into property values. Additionally, this chapter makes use of the largest dataset to date within the wider literature and is the first to use exact matching between solar and non-solar properties through the application of propensity score matching.

Chapter 1

The first chapter of the thesis provides an introduction to the topic area. I outline the key policy issues and debates around renewable energy - specifically onshore wind and solar PV - in the UK. These two technologies are particularly relevant within the context of the UK emissions reductions goals and policies set out to achieve them. I also present a review of the academic literature on environmental amenities, disamenities, and the valuation of non-market goods through the application of hedonic pricing methods; which decompose property transaction prices to place values on individual characteristics. Given the reliance of this thesis on these approaches this chapter contextualizes much of the work in Chapters 3, 4, and 5.

Chapter 2

The second chapter reviews the literature on the house price effects of windfarm proximity and visibility in detail, to provide a clear context for the empirical work and the contributions made in Chapters 3 and 4. I examine the methodological approaches and the empirical findings of the papers included here. The literature is well-developed, but remains very much in disagreement as to what - if any - house price effects arise from windfarm proximity and visibility. Broadly, some research has found either a negative price effect (Heintzelman & Tuttle

2012, Jensen et al. 2014, Gibbons 2015, Sunak & Madlener 2016, 2017, Jensen et al. 2018) while other papers find no statistically significant effect (Hoen et al. 2011, Brown et al. 2012, Lang et al. 2014, McCarthy & Balli 2014, Vyn & McCullough 2014, Hoen et al. 2015, Hoen & Atkinson-Palombo 2016) from windfarm proximity and visibility. A single study by Heblich et al. (2016) found mixed results including some evidence of positive price effects.

Chapter 3

As set out in Chapter 2, there are some key limitations of the current literature. First, there has been an substantial increase in the number of wind turbines and their distribution throughout England and Wales leading to a very different landscape since the most recent empirical analysis was performed (Gibbons 2015). Second, there are a variety of methods applied within the literature to incorporate windfarm visibility into the analyses, but no consensus on the best practices for achieving this - or even if it is necessary. Third, there is very limited discussion around the importance of the study period to driving results or the persistence of price effects over time. The chapter begins with a replication of Gibbons (2015) by applying a staggered Difference-in-Difference spatial fixed effects model, before extending it in a number of ways.

While traditional electricity generation technologies such as coal may affect property values through lowered air quality, visibility is assumed to be the key driver of any observed price effect from proximity to wind energy developments. Testing this assumption through detailed sensitivity analysis with alternative visibility estimation techniques provides important insights to the results found by Gibbons (2015), who defines a windfarm as visible if there is an unobstructed line of sight between the center of a windfarm and the center of a postcode. Gibbons (2015) restricts his analysis to include only windfarms located in rural areas, citing the possibility that there may be difficulties accurately estimating their visibility and analyzes transactions covering the period 2000-2012. I apply the same visibility definition but test it for sensitivity to the relaxation of restriction criteria and the use of a more refined digital elevation model. This sensitivity testing generates useful insights into the variety of results within the literature which may stem from the varied methods applied within it.

Through this analysis, I confirm that the inclusion of urban windfarms does significantly skew the estimated price effect, something that has not been tested in detail within the literature. I also find evidence that visibility estimation is a key driver of results by comparing the use of alternative digital elevation models when calculating visibility estimates. Additionally, I test visibility estimation accuracy by applying a variety of visibility threshold tests. Gibbons also tested for treatment intensity by comparing effects by windfarm size, however this chapter provides an alternative intensity measure based on the presence of multiple windfarms, producing evidence that larger visual impacts from wind turbines lead to a larger observed price effect.

I am able to extend the study period beyond that considered by Gibbons, to include a total of 23 years of property transactions placing it just behind Heblich et al. (2016) with 24 years and Dröes & Koster (2016) with 26 years of housing transactions covered by the analysis. However, it includes the most recent set of transactions within the literature, which is an important addition to the analysis, as the latest analysis in the literature (Heblich et al. 2016) includes transactions up to 2014. The increasingly rapid deployment of larger wind turbines and larger windfarms has led to a substantial change in the population's exposure to windfarms, requiring an updated analysis, for England and Wales in particular. The expansion exploits the additional data to show: 1) over the full period, there is a reverse of the negative impacts found under the replicated analysis; and 2) results are sensitive to the duration of the analysis: when subdivided into smaller time periods, the observed negative price effects are largest at the earliest period and become either small or positive by the most recent period. This is the first analysis which finds that the negative price effects from windfarm visibility and proximity dissipate over time, and ultimately become positive under some assumptions.

To summarize: in this chapter I extend the literature by exploring how alternative visibility measures influence estimated price effects, highlighting a divergence between urban and rural areas, and finding evidence that price effects are not stable over time. However, there are some limits to the average sales analysis of this chapter, such as the fact that the average transaction price in a postcode may not reflect the average properties within that postcode; postcode boundaries change over time, the fact that the sole property characteristic controlled for is the construction type; and the use of postcode and windfarm centroids to estimate visibility and proximity. I use the analysis of this chapter as a baseline and in the subsequent chapter I apply a repeat sales analysis, which addresses these limitations. I then use the exact locations of individual properties to further refine the modeling of visibility, found to be a key driver of results.

Chapter 4

Chapter 4 takes an alternative, but complementary, approach to understanding the effects of windfarm proximity and visibility on house prices. Specifically, I use a repeat sales analysis while maintaining the key assumptions to those applied within Chapter 3 to further test them under a repeat sales model. These assumptions are 1) visibility is the main driver of an observed price effect; 2) the magnitude of visibility is largest for properties closest to wind turbines. The repeat sales approach estimates the change in price for individual properties by comparing transaction prices of the same property before and after a nearby windfarm becomes operational. This serves to test the methodological framework of the analysis for consistency in estimated effects, which has previously been left unexplored within England and Wales. Here, properties which sell multiple times over the study period are used to generate an estimate of the price effect from windfarm proximity and visibility.

This chapter marks the first repeat sales analysis exploring windfarm visibility and proximity in England and Wales, and the repeat sales and average sales analyses together provide several key contributions. First, Chapter 3 assumes that any transaction(s) within a postcode accurately reflects the average property, which may not be the case due to differences in property characteristics and the limited number of simultaneous transactions within a postcode. This is a potential source of bias in the estimated effects. The analysis of this chapter focuses on changes in price of the same properties, comparing effects before and after windfarm operationality. This allows for a cleaner estimation of price effects as it eliminates the issues of comparing price effects on homes with differing characteristics, and - by its nature - location is automatically controlled for. Second, by defining visibility at the property level, the repeat sales approach allows for two properties in the same postcode to have differing visibility estimates. I find some evidence for slight differences in home buyers preferences towards properties with views of windfarms versus properties situated in postcodes where windfarms may be visible.

As with the average sales analysis, I find that the two factors which have the greatest influence on the estimated price effects are the method and granularity of visibility estimation and the timing of the analysis. When the most detailed and refined visibility estimates are used, there is no negative price effect from windfarm proximity and visibility, but there is a positive price impact. As in Chapter 3, I find further evidence that attitudes toward windfarms have changed over time, despite the increasing size of wind turbine height and blade diameter. This implies that over the study period, windfarm visibility may have shifted from a disamenity to an amenity or at the very least is a preference-neutral environmental feature.

In summary, Chapters 3 and 4 apply a similar set of underlying data and assumptions but use complementary analyses to estimate the impact of windfarm visibility and proximity on house prices. The two analytical approaches applied serve both to display the robustness of the findings, due to their alternative strengths and weaknesses. The two chapters both show evidence that alternative visibility calculations heavily influence the observed price effect of windfarm siting on house prices - as does the study period of the analysis. This is the first academic research to heavily test the importance of alternative specifications of the models used to estimate house price impacts from windfarm siting within the same study. This is a significant contribution to the literature, and may explain the disagreement in the direction and magnitude of reported effects therein.

Chapter 5

Whereas the previous empirical chapters dealt with large-scale renewable energy generation, in this chapter I analyze house price effects from residential-scale installations. As shown in Chapter 2, the installation of small-scale installations has become a popular policy to reduce emissions within England and Wales. I make several contributions to the literature on residential photovoltaic (PV) capitalization into house prices. I generate the first estimates of a solar property price-premium and the capitalization of PV systems into English and Welsh house prices. The empirical analyses of the chapter include the application of a hedonic regression model, a repeat sales analysis, and propensity score matching.

This chapter contributes to a limited literature currently consisting of only six peer-reviewed papers in this topic area. All of which have found evidence that a solar property price-premium does exist, and that residential PV systems are capitalized into property transactions. The majority of the papers in the literature (Dastrup et al. 2012, Ma et al. 2016, Wee 2016, Lan et al. 2020) have applied a simple hedonic regression approach as their baseline estimation for comparison against an alternative approach - although Hoen et al. (2013) use this approach to generate their headline results. The most common alternative approach is a repeat sales analysis, as applied in (Dastrup et al. 2012, Ma et al. 2016, Wee 2016), but Qiu et al. (2017) and Lan et al. (2020) make use of fuzzy matching techniques to create a set of control properties to estimate premiums and capitalization. This chapter makes use of all three approaches and the results of each model tell the same story - that there is evidence of a solar property premium in England and Wales.

To summarize, this chapter is the first analysis in the literature to estimate a solar property premium and residential PV capitalization in England and Wales. I show that solar PV systems are capitalized into English and Welsh house prices. Further, I find that the solar property premium is large enough to recover and profit from the cost of installing the typical English and Welsh residential PV system. In this chapter, I extend the literature in two key ways. First, the larger sample of data gives the analysis considerably more statistical power than any of the other papers in the literature. Second, its size is exploited to perform an exact match between treatment and control properties through the use of exact matching criteria based on property characteristics. This is an improvement upon the fuzzy matching methods employed by Lan et al. (2020) and Qiu et al. (2017), as it ensures that treatment and control properties are as similar as possible, with the exception being the installation of a residential PV system. The analysis contributes to a very limited existing literature which has to date been focused on the United States and Australia. I find that compared to the existing literature, the price premium is smaller, possibly explained by the smaller solar energy endowment of England and Wales relative to other study areas.

Chapter 6

This chapter concludes the thesis, contextualizes its contributions to the literature and provides suggestions for future research into the topic area.

Chapter 1

This Thesis in Context

1.1 Introduction

This thesis contains a series of empirical analyses which evaluate externalities arising from renewable energy developments, particularly how these impact on the housing market within England and Wales. The technologies analyzed are onshore wind, offshore wind, and residential solar photovoltaic generation. Within this chapter, I provide a contextual background for these analyses. First I discuss the legally binding emissions reductions commitments made by the UK through international agreements, second I discuss the domestic policies implemented to achieve these targets. Third, I explain how the resulting energy developments may impact the housing market of England and Wales. Lastly I overview the broad set of methods used both within this thesis and the relevant to estimate the value of these impacts.

1.2 International Climate Change Policy

Over the past 50 years, there has been a growing consensus among climate scientists that climate change is due to anthropogenic emissions of greenhouse gases (GHGs) (Oreskes 2004). Powell (2017) find that among climate scientists, there is no peer-reviewed research taking a dissenting posture from this consensus. As the United Nations Fifth Assessment Report on Climate Change states:

> “Continued emission of greenhouse gases will cause further warming and long-lasting changes in all components of the climate system, increasing the likelihood of severe, pervasive and irreversible impacts for people and ecosystems. Limiting climate change would require substantial and sustained reductions in greenhouse gas emissions which, together with adaptation, can limit climate change risks[1]."

More recently, social and economic costs of climate change have been the subject of increasing interest among researchers, and concern among policymakers. Overall, these costs are expected to be negative, though they are not evenly distributed globally or even nationally (IPCC 2014). Within the United States, Hsiang et al. (2017) find that across the sectors analyzed, the estimated cost of climate change will be approximately 1.2% of GDP per 1°C increase in the global average temperature. They do find positive economic impacts for some regions within the US, but the negative impacts in other regions, particularly the South, outweigh these. The estimated impacts for the least developed countries are even more dire and range from 2-20% of country income (Hsiang et al. 2017). Within the African Continent, the economic costs are estimated to reach 1.5-3% of GDP per year by 2030 (Watkiss P 2010).

The economic costs associated with climate change have motivated the global community to attempt to reduce global emissions of GHGs. This has led to a series of international agreements aimed at reducing GHG

[1]United Nations Fifth Assessment Report on Climate Change, 2014, pp 8.

emissions globally including the *Kyoto Protocol* (1997) and *Paris Agreement* (2015). Both the Kyoto Protocol and Paris Agreement build on the *United Nations Framework Convention on Climate Change* (1992) (UNFCC). This framework convention has been ratified by 197 countries and is aimed at preventing dangerous human interference with the climate system. Created prior to the scientific consensus on the anthropogenic nature of climate change, the UNFCC binds states to act in the interests of human safety even in the face of scientific uncertainty. The Kyoto Protocol was adopted in 1997, but did not come into force until 2005. It committed developed countries to both limit and reduce their GHG emissions by setting individual targets for each country. These limits were binding for developed nations, but not for developing nations as it recognizes that developed and developing nations will have different capacities to reduce their emissions while maintaining economic growth. However, there have been a series of criticisms about the Protocol and claims that it has failed to achieve its objectives of addressing climate change caused by GHG emissions (Rosen 2015).

The 2015 Paris Agreement addresses many of the limitations of the Kyoto Protocol. This is because it is legally binding on all ratifying governments, and its goal is to limit global warming to below 2°C, but preferably to below 1.5°C relative to pre-industrial average global temperatures. To achieve this, countries set out their nationally determined contributions (NDCs) to GHG reductions which are aimed at reaching a climate neutral world by 2050. These NDCs set out the actions that each country has committed to in order to reach its targeted emission reductions, and these are published every 5 years. The United Kingdom has been working to achieve its own emissions reductions targets within this framework, and has produced some of the most ambitious reductions targets in the world (PMO 2021). I will now discuss the domestic policies that have been implemented as part of these goals, and how these relate to the externalities of renewable energy developments on house prices.

1.3 The United Kingdom and the Low-carbon Transition

The United Kingdom has committed to substantial reductions in GHG emissions, and has worked towards these targets with increasing ambition. In its previous NDC, the government had legally committed to reducing its emissions by 68% compared to the levels in 1990 by 2030 with the objective of reaching net zero emissions by 2050. As the UK continues to progress towards this goal, it has now committed to a reduction of 78% by 2035, and this will now include the UK's share of international and aviation shipping emissions (PMO 2021). Though these legally mandated emissions targets are based on the commitments made in the Paris Agreement, the ability to meet and exceed them are greatly influenced by past and present policies promoting the transition to a low-carbon economy. I will focus my discussion first on policies supporting large-scale renewable energy developments, and then turn to policies supporting small-scale or micro-generation developments.

1.3.1 Large-scale and Commercial Renewable Energy

The first key policy which has led to the UK's success in the the transition to a low-carbon economy is the Non-Fossil Fuel Obligation (NFFO). This policy went into effect in 1990 and it was created to provide financial incentives to energy providers to invest in non-fossil fuel electricity generation such as Nuclear plants or renewables. This was partly achieved through mandating suppliers make non-fossil fuel electricity purchases and by setting the price that these purchases would be made at. The NFFO has since been replaced, but these contracts will continue until the NFFO fully expires in 2019 (OFGEM 2020*d*).

The Renewables Obligation (RO) replaced the NFFO in 2002, and provided additional support for the transition towards low-carbon electricity generation. This policy forced electricity suppliers in the UK to source an increasing share of their generated electricity from renewable energy sources. The mandated proportion of supply was initially 3% in 2002, but increases to 49.2% of supply in the 2020/2021 period (OFGEM 2020*e*). Suppliers receive Renewables Obligation Certificates (ROCs) which are equivalent to the amount of renewable energy they produced. Suppliers which exceeded the mandated proportion of their total generation from renewables could then sell their excess ROCs to those that had not met the minimum threshold. Those suppliers who do not meet their targets are required to purchase the number of ROCs equivalent to to the difference between earned ROCs and the number which would have been generated had they met the mandated proportion of renewable generation. Just as the mandatory proportion of renewable electricity increased each year, so did the buy-out price of ROCs (OFGEM 2020*e*).

These two policies, but particularly the RO were great successes in regards to their encouragement of diversifying the UK electricity generation away from fossil fuels and towards renewable generation at the commercial-scale. These policies lead to substantial growth in wind energy developments across the UK, which is home to the largest wind energy endowment in Europe (Dalla Longa et al. 2018). I focus the discussion in this Chapter as well a the Thesis on England and Wales as this is the study area of the empirical analyses contained herein.

The growth in electricity generation from wind in the ten years following the RO going into force is presented in Appendix A1. The growth in the English and Welsh wind turbine stock between 1992 and 2017 is highlighted in Figure 3.2 of Chapter 3. This trend towards increasing generation from wind turbines is also exemplified in Figure 3.3 which displays the growth in the total generation capacity of wind energy developments in England and Wales. There has also been substantial government support for micro-generation technologies in England and Wales. I will discuss these policies, and how they have influenced the deployment of small-scale renewable generation in the following section.

1.3.2 Small-scale and Residential Renewable Energy

The UK has implemented two policies - the first of which has been key to supporting the growth in small-scale renewable generation developments and the second will be key to continuing this support. These are the Feed-in Tariffs (FIT) Scheme and the Smart Export Guarantee (SEG). The aim of the FIT was to support the development and adoption of renewable and low-carbon electricity generation at small-scales. Generation capacity was required to be below 5MW, and eligible technologies include: solar photovoltaic (PV), wind, micro combined heat and power (CHP), hydro, and anaerobic digestion (AD). I will focus the discussion on solar PV (OFGEM 2020*c*).

Participants in the FIT scheme are paid for all electricity they generate as well as additional payments for any excess generation which is exported back into the grid. These payments are guaranteed and based on the year of entry into the FIT, though the rate decreased with each year which incentivized early adoption. The payments continue for 25 years after entry into the scheme. The FIT opened in 2010, but closed to new entrants in 2019, and participants could be homeowners but any system smaller than 5MW in capacity was eligible (OFGEM 2020*a*). Figure 5.3 shows the cumulative installed residential solar PV systems in England and Wales, with a substantial increase from the opening of the FIT scheme in 2010.

Replacing the FIT is the Smart Export Guarantee, which came into force January 2020. While the FIT participants received a payment for all electricity generated, the SEG guarantees participants will be paid only for electricity exported to the grid. The same technologies which were eligible for the FIT are eligible for the SEG. Though the SEG only ensures payment for exported electricity, the rates paid are guaranteed to be above £0, even if the wholesale price of electricity is at or below £0. Although the SEG is too recent to be of consequence to the empirical analyses presented in Chapter 5, this policy is the main source of continued UK government support for small-scale renewables.

1.3.3 Summary of Domestic Policies

The policies that have been implemented in the UK to support the transition to low-carbon electricity generation and to reduce GHG emissions have been highly successful. As a result, the UK is on track to achieve among the most ambitions emissions reductions targets in the world and ahead of schedule on meeting its Paris Agreement commitments (PMO 2021). The NFFO and RO policies were responsible for the substantial commercial-scale wind energy developments throughout the UK by mandating an ever increasing share of electricity be provided by renewable energy. At the same time, the levelized cost of generation for renewables has fallen sharply and are now produced at lower costs than competing fossil-fuel sourced electricity IEA (2020*a*).

Evidence of the success of these policies can be seen in the share of UK electricity generation by source in 2002 and 2020. Coal accounted for 35.61% of electricity generation in 2002, but dropped to 3.11% in 2020. Over the same period, wind and solar generation accounted for only 0.25% of the UK's electricity in 2002, but by 2020 it comprised 27.98% OFGEM (2020*b*). The FIT has supported a similarly rapid deployment of small-scale and domestic low-carbon generation, with residential solar PV seeing substantial adoption rates through the life of the scheme. These micro-generation technologies will continue to be supported through the SEG which came into effect in 2020 OFGEM (2019).

1.4 Externalities, Amenities, and Renewable Energy

The rapid deployment of both commercial-scale wind generation and residential solar have led to an intersection of the transition to a low-carbon economy and the housing market. This is a result of the trade off between the emissions reduction from adopting low-carbon energy developments and their potential impacts on house prices. Researchers have a choice between applying stated or revealed preference methods to estimate values of the house price impacts of wind turbine proximity and visibility - or residential PV systems. Within the context of the housing market, both methods seek to value individual's willingness to pay for (WTP) or willingness to accept (WTA) the presence of some feature of a property.

Stated preference methods, such as contingent-choice or contingent-valuation, use surveys to estimate the value of an environmental feature or housing characteristic. These approaches are relatively straightforward as they simply ask individuals how they value a given feature of a home (Denant-Boemont & Hammiche 2019). For example, a survey may ask individuals how much they would be willing to pay to prevent the construction of a wind turbine located near their home, or how much compensation for its construction would make them willing to accept it. These are essentially hypothetical questions, and individuals may not have an accurate valuation. As such, he main issue with stated preference methods is the well-documented discrepancy between stated WTA and WTP and actual behaviors (Cohen 1990, Plott & Zeiler 2005, Song et al. 2012).

Revealed preference methods are the preferred approach to valuing the characteristics of residential properties. This is because preferences are revealed through actual choices made by individuals Richter (1966). Estimating consumer preferences towards these developments, and valuing their impact on house prices has been extensively applied within the literature, reviewed in Chapter 2. The underlying assumption of any hedonic model is that the price of a given product is a function of its characteristics (Rosen 1974). Therefore the ultimate goal of a hedonic analysis is to decompose price into the characteristics of the product Rosen (1974). This can then be applied to housing transactions, where the residential property is the product whose price is a function of its characteristics.

1.4.1 Rosen and Hedonic Theory

The analyses of this thesis rely on Hedonic Valuation and in this section I will briefly discuss the origin and theory which underpins it. The seminal paper within the Hedonic Valuation literature is Sherwin Rosen's 'Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition' Rosen (1974). Rosen outlines how markets match buyers and sellers of multidimensional goods and this framework has been widely applied within economics as nearly all goods contain multiple characteristics. The paper sets out a method for estimating the willingness to pay for goods which do not have explicit markets.

A multidimensional good can be described by a vector of its characteristics. For homes the characteristics could be intrinsic (number of bedrooms, size of garden) or extrinsic (local school quality, or local air quality). Therefore the market price for the ith home is written as follows:

$$P_i = P(c_{i1}, c_{i1}, ..., c_{in}) \tag{1.1}$$

Where P_i is the price of house i and is comprised of the value of the characteristics of the house. The partial derivative of P with respect to the jth characteristic ($\partial P/\partial c_j$) is the marginal implicit price. This is the marginal price of the jth characteristic, when all other characteristics are held constant and is implicit to the overall transaction price of the house Greenstone (2017).

The intersection between house prices and a given characteristic is the hedonic price schedule (HPS) which is generated by the equilibrium interaction of buyers and sellers. If we assume that housing markets are competitive and buyers purchase at the market price, buyer's utility is maximized when individuals choose maximum levels of characteristics which satisfy their budget constraint. Consumer utility depends on consumption of the numeraire X and the vector of characteristics:

$$u = u(X, C) \tag{1.2}$$

The budget constraint is expressed $I - P - X = 0$, where I is income. Maximization of 1.2 with respect to the budget constraint reveals that individuals choose levels of each characteristic whch satisfy $(\partial U/\partial c_j)/(\partial U/\partial x)$ = $\partial P/\partial c_j$) Therefore, the marginal willingness to pay for a characteristic (school quality) will equal the cost of an extra unit of that characteristic within the market Greenstone (2017).Substituting the budget constraint into 1.2 gives $u = u(I - P, c_1, c_2, ..., c_n)$ and inverting this while holding all characteristics but j constant gives the expression for willingness to pay for c_j:

$$B_j = B_j(I - P, c_j; C*_j, u*) \tag{1.3}$$

Here, $u*$ is the highest level of utility attainable within the budget constraint and $C*_j$ is the vector of optimal quantities of other characteristics. Within the literature, this is referred to as a bid curve (indifference curve) because it reveals the maximum amount that a buyer would pay for different values of c_j holding utility constant Greenstone (2017). Heterogeneity in buyer's indifference curves due to differences in preferences or income levels leads to differences in the chosen quantities of a given characteristic. Buyers will choose homes where their marginal willingness to pay for c_j is equal to the market-determined marginal implicit price Greenstone (2017).

On the supply side of the housing market, it is assumed that suppliers are heterogeneous. This arises due to differences in the cost functions across suppliers. To determine the curve of the offer price for suppliers, we invert their profit function and can then determine the offer curve for characteristic c_j:

$$O_j = O_j(c_j; C*_j, \pi*) \tag{1.4}$$

Where π* is the maximum available profit given the supplier's cost function and the given Hedonic Price Schedule. The HPS is formed by the tangencies between a consumer's bid and the suppliers offer curves. At all points along the HPS, the marginal price of a given housing characteristics is equal to both the buyer's marginal willingness to pay for that characteristic and the suppliers marginal cost of providing it. The HPS reveals the marginal willingness to pay for the set of consumers that have sorted themselves by the prices and quantities of the characteristics of interest. We can then infer welfare gains or losses associated with marginal changes to the quantities of given characteristics within the relevant population of buyers Greenstone (2017).

However, the HPS on its own cannot be used to determine welfare impacts from non-marginal changes in characteristics of interest Rosen (1974) this is because the points observed along the HPS are only for other buyers who may have different preferences or income levels. Rosen's solution is to introduce a two-step econometric procedure that delivers both the HPS and bid functions. Using neighborhood vista quality as a an example, the first step is to regress house prices against all characteristics including vista quality. This allows their effects to be non-linear Greenstone (2017):

$$p = \alpha + f(c_1, c_2, ..., c_n) + \varepsilon \quad (1.5)$$

The marginal implicit price of vista quality is the derivative of house prices with respect to vista quality and this quantity is then used in the second stage equation:

$$\partial p(C)/\partial c_{vistaquality} = \alpha + g(c_{vistaquality}) + \varepsilon' \quad (1.6)$$

Here ε includes all demand influences such as preferences or incomes, which are observable and can be included as covariates. We then evaluate $g(c_{vistaquality})$ at different values of $c_{vistaquality}$ to determine the bid function. If Rosen's approach is successfully applied it reveals a buyer's preferences, and measures welfare impacts arising from non-marginal changes in characteristics of homes.

1.4.2 Challenges of Hedonic Estimation

The foundation of any welfare calculation rests upon the consistent estimation of Equation 1.1 because welfare impacts from marginal changes in a given characteristic are given directly by the HPS. Another risk is that inconsistent estimation of the HPS will result in an inconsistent marginal willingness to pay function which essentially invalidates any welfare analysis of non-marginal changes. This is particularly challenging to overcome in a cross-sectional and panel data settings Greenstone (2017). This presents itself as severe sensitivity to the inclusion of additional covariates and has been illustrated within the application of hedonics under a variety of settings - ie the housing market Chay & Greenstone (2005), as well as health economics Black & Kniesner (2003).

These will largely arise due to omitted variable bias. To address this, more recent works applying a hedonic valuation approach make use of quasi-experimental variation within the covariates of interest. These include, for example, variations caused by nature, politics, accidents, or actions/influences beyond the control of the researcher because these are assumed to be exogenous. Within quasi-experimental contexts, researches will understand the source of the variation.

A further challenge arises when researchers attempt to infer the bid function or offer function directly from the HPS. It has been demonstrated that hedonic methods are highly reliant on strong functional form assumptions Brown & Rosen (1982). Particularly crucial are assumptions on the house price function and whether it is assumed to be linear or non-linear. Additionally, there is the risk that inferring the bid functions from the HPS may be undermined by taste-based sorting of home buyers.

Despite these challenges, hedonic valuation with a quasi-experimental research design remains a workhorse within many research areas. Although consistent estimation of the HPS will not recover the underlying bid and offer functions it can estimate welfare impacts of non-marginal changes which arise from past changes in amenities. This is the context in which the approach is applied within the subsequent chapters of this thesis. I make use of a quasi-experimental design, and compare early to later periods which allows for measurement of non-marginal changes to house prices. In the subsequent section I will describe the application of the reduced form hedonic estimation which underpins the analyses of Chapters 3,4, 5.

1.4.3 Reduced Form Hedonic Estimation

In the context of hedonic modeling - or hedonic pricing - applied to house prices, we assume that the sales price of a given home is based on the sum of its total characteristics and their amenity value. Amenity value is simply the currency value that buyers ascribe to each characteristic of a home. These could be positive or negative values, meaning that the characteristic may increase or decrease the total price of a given property (Monson 2009). In Equations 1.8 and 1.9 I have split the characteristics into intrinsic and extrinsic groups. Intrinsic characteristics are those contained in the property itself, such as the number of bedrooms or lot size. Extrinsic characteristics refer to those which may contribute to the value of the property, but are not features of the property itself such as proximity to the beach. For the purpose of this thesis, as well as the hedonic literature, when a feature increases the value of a property, it is considered an amenity. Any feature which decreases the value of a property is a disamenity.

$$House\,Price = f(House\,Characteristics) \tag{1.7}$$

$$Price = f((Intrinsic\,Characteristics)\ +\ (Extrinsic\,Characteristics)) \tag{1.8}$$

$$\begin{gathered} Price = f((Bedrooms + Bathrooms + Lot\ Size + ...))+ \\ ((School\ Quality + Proximity\ to\ Beach + Transport\ Links + ...)) \end{gathered} \tag{1.9}$$

Hedonic Pricing methods have been applied to evaluate a wide variety of Environmental features, for example: Proximity to Rivers (Anderson & West 2006, Gibbons et al. 2014, Tapsuwan et al. 2015) , Urban Greenspace (Chen & JIM 2010, Brander & Koetse 2011, Schläpfer et al. 2015), and Wildfires (Stetler et al. 2010). The approach has also been applied to estimate the value of intrinsic property characteristics in-

cluding the number of bedrooms, bathrooms, fireplaces, garage size, square footage, pools, yard size, etc. (Coley 2005, Cebula 2009, Liao & Wang 2012, Pride et al. 2018). I will discuss the classification of intrinsic and extrinsic features as amenities or disamenities using examples relevant to the empirical analyses of this Thesis.
Installed residential PV systems are an intrinsic characteristic - much like the number of bedrooms, etc. That is to say the presence of a PV system us a feature specific to the property and as such, is expected to influence the price. There is a growing literature on the influence that residential PV systems have on property values, which is discussed within Chapter 5. Though the peer-reviewed literature applying hedonic pricing methods to estimate the amenity value of such systems is limited, all published papers have found that it positively influences property prices. Therefore residential PV systems are considered an intrinsic amenity.

Wind turbines and windfarms fall into the extrinsic set of property characteristics, and have therefore been included within the Environmental Amenity literature (Brinkley & Leach 2019). Schaeffer & Dissart (2018) defines environmental amenities to be the set of local attributes enhancing the residential quality of life. I then take this to define environmental disamenities as the set of local attributes which reduce the residential quality of life. There may be improvements to residential quality of life through the reduction of emissions that wind energy facilitates, however there may also be potential costs. Similarly, there may be amenity value to the installation of residential PV systems. The value of these benefits or costs can be estimated through changes in house prices.

There has long been concern that wind-turbines may have negative impacts on nearby house prices, and the first academic paper on the potential influences of wind turbine proximity and visibility was published in 2007 by Sims & Dent. This paper did not contain a dataset which allowed for a highly detailed analyses, but since its publication there has been a substantial academic interest in estimating house price effects arising from wind energy facilities. This research has focused on the potential impacts of windfarm or wind turbine proximity and visibility as environmental amenities or disamenities by analyzing their impacts on housing markets. Though the academic research on this subject is well-developed, there is not yet a consensus on whether wind turbines should be classified as environmental amenity or disamenity. A detailed discussion of this literature is provided in Chapter 2.

1.5 Hedonic Pricing and Renewable Energy

The previous section explained how renewable energy developments may be defined as amenities if they increase a property value or disamenities if they decrease a property value. In this section, I will discuss the application of Average Price and Repeat Sales Hedonic Pricing Approaches. As described in the previous section, the value is ascribed by analyzing property characteristics as shown in Equation 1.9. The process is the same for both intrinsic and extrinsic characteristics, as both types of characteristics contribute to a property's value. Here I will outline the average and repeat sales approaches to value property characteristics, but with a focus on examples such as wind turbines and residential PV.

In Equation 1.10 I again show the basic setup of a hedonic analysis, but describe its application under an Average Price analysis. On the left-hand side is the average transaction price, averaged over a specific geographic region at a specific point in time - usually at quarterly or monthly intervals. Some examples of geographic regions used within the literature discussed in Chapter 2 are postcodes (Gibbons 2015) and cities Hoen & Atkinson-Palombo (2016). In a Repeat Sales analysis, the left-hand side of the equation compares the prices of the properties which sell multiple times over the study period. For both approaches, the right-hand side of the equations will include variables representing intrinsic and extrinsic characteristics of the analyzed properties.

When the goal is to estimate the impacts of a nearby wind turbine, researchers will include a term to capture the influence of this characteristic on the average price at the geographic or property level. Time is also included in the analysis as there may be time trends which influence house prices. The key advantages and disadvantages of the Average Price approach are thoroughly discussed in Chapter 3, but the most common justification for their use is that it does not require information on the exact location of the properties in the dataset. This usually allows for an analysis with a larger group of properties to be included. The key advantages of the Repeat Sales Approach discussed in Chapter 4 but the main justifications for applying this approach is that it is less likely to suffer from omitted variable bias as the properties are compared to themselves.

$$Price = ((Intrinsic) + (Extrinsic) + (Time)) \tag{1.10}$$

In their most basic setup, both repeat and average price transactions will capture the impact of renewable energy developments through a basic indicator dummy variable on the right-hand side of the equation. This simply captures the difference in price between properties affected by the development and the properties unaffected by it. More robust analyses will incorporate the timing that a renewable development, allowing for a comparison of price impacts both before and after a property is impacted. I describe the hedonic frameworks of the literature estimating wind turbine impacts on house prices in great detail in Chapter 2. I do the same for house price impacts from residential PV systems in Chapter 5.

1.6 Conclusions

This chapter provides a contextual overview of the Thesis topic area. First I present the policy environment relevant to the empirical analyses of Chapters 3, 4, and 5. This included a description of the UK's binding commitments to reduce its emissions in line with the targets set in the Paris Agreement, as well as an outline of the policies implemented to achieve this at a record-setting pace. The NFFO and RO both supported the development of large-scale renewables - particularly on and offshore wind - which are the focus of the analyses within Chapters 3 and 4. The FIT and SEG schemes support small-scale renewables and the FIT has been particularly successful in the rapid uptake of residential PV systems which the subject of the analyses of Chapter 5.

I then describe how these policies which have led to the rapid transition to low-carbon electricity generation may create externalities within the housing market, through the amenity value they bring to homes in England and Wales. In this section, I discuss the definition of amenities and disamenities. I then overview the approaches used to estimate the preferences towards these developments through the application of Hedonic Pricing approaches such as Average Price and Repeat Sales. The contextual background of this chapter sets the stage for the literature which is reviewed in the subsequent chapter, as well as the empirical analyses contained within this thesis.

Chapter 2

Literature Review: Wind turbine proximity, visibility and house prices

2.1 Introduction

Although the literature examining house price impacts from wind turbine proximity and visibility is well-developed, there is not a consensus within the reported findings. These papers apply differing statistical methods as well as differing sets of assumptions regarding the data utilized within their analyses. The variety in the analytical approaches and decisions applied occurs not only across papers which find different impacts (no effect, decrease, increase) but it is also present even across papers which find a similar effect. In this chapter, I will review this literature in detail, discussing the findings, analytical approaches, and methodological assumptions applied for each analysis. These papers are published in academic journals at the time of writing, with the exception of Heblich et al. (2016), which is a report from a research project exploring this issue in Scotland. Each paper reviewed within this chapter attempts to value the impacts on house prices arising from nearby windfarm siting, some account for visual impacts while others rely solely on proximity to the developments to estimate the impact. I focus the discussion around the fifteen relevant papers which comprised the literature at the time of writing.

Six of the papers have found statistically significant, negative price impacts arising either through proximity to or visibility of nearby wind turbines (Heintzelman & Tuttle 2012, Jensen et al. 2014, Gibbons 2015, Sunak & Madlener 2016, 2017, Jensen et al. 2018). The papers reporting negative impacts range from a price decrease of 1.4% for properties within a 2km distance from a wind turbine (Dröes & Koster 2016) to a decrease of up to 14% for properties where visible turbines strongly affected the vistas surrounding the properties (Sunak & Madlener 2016, 2017). These papers apply differing analytical approaches, for example, Heintzelman & Tuttle (2012) and Jensen et al. (2018) do not model visibility of wind turbines and therefore report effects of proximity alone. Dröes & Koster (2016) do model wind turbine visibility from the properties in their analysis, but they find a statistically insignificant positive effect, while finding a statistically significant decrease arising from proximity alone.

Eight papers find no statistically significant house price effect from proximity to, or visibility, of wind energy facilities (Sims & Dent 2007, Sims et al. 2008, Hoen et al. 2011, Lang et al. 2014, McCarthy & Balli 2014, Vyn & McCullough 2014, Hoen et al. 2015, Hoen & Atkinson-Palombo 2016). Of this group, all papers with the exception of Sims & Dent (2007), Hoen et al. (2015), and Hoen & Atkinson-Palombo (2016) include a measure of the visibility of wind turbines within their analysis. Lang et al. (2014) and Hoen & Atkinson-Palombo (2016) report negative, but not significant effects. McCarthy & Balli (2014) and Vyn & McCullough (2014) find positive, but insignificant impacts from windfarm visibility. The rest of the papers find a mix of positive or negative impacts depending on the specification of their models - though all are statistically insignificant. Lastly, one paper (Heblich et al. 2016) also finds both statistically significant positive and negative impacts from visibility, though these depend on the distance to the visible wind turbine.

The geographic distributions of study areas within the literature spans across North America, Europe, and Oceania. The largest concentration is within the United States (Hoen et al. 2011, Heintzelman & Tuttle 2012, Lang et al. 2014, Hoen et al. 2015, Hoen & Atkinson-Palombo 2016); followed by the United Kingdom (Sims & Dent 2007, Sims et al. 2008, Gibbons 2015, Heblich et al. 2016); Denmark (Jensen et al. 2014, 2018); Germany (Sunak & Madlener 2016, 2017); the Netherlands (Dröes & Koster 2016); Canada (Vyn & McCullough 2014); and New Zealand (McCarthy & Balli 2014). Within the United Kingdom, the findings reflect the disagreement within the broader literature with two papers finding no statistically significant effect Sims & Dent (2007), Sims et al. (2008); one finding statistically significant negative effects Gibbons (2015); and one finding both positive and negative effects Heblich et al. (2016).

Within this chapter of the thesis, I will detail the underlying data and assumptions applied within the literature to estimate house price effects from windfarm siting. Each analysis within the literature has made decisions regarding both the statistical approaches and assumptions towards the data used to answer these questions, and these analytical decisions are potentially responsible for some of the disagreement in their reported findings. I have structured the discussion of the literature around these decisions and assumptions.

2.2 Inclusion Criteria

To perform a hedonic analysis, housing transaction data are required. Therefore the very first research decision for all papers within the literature is to determine which transactions will be included in any analysis. This choice involves firstly determining which windfarms or wind turbines the analysis will estimate siting impacts for, and then determining which properties are considered 'nearby.' Once the selection of windfarms and property transactions has been made, many studies further restrict the transactions to remove transaction prices above or below a certain price threshold. All papers within the literature use Geographic Information System (GIS) techniques to calculate distance from either windfarms or wind turbines to properties. This distance measure is then used to determine which housing transactions will be used within the study, but the maximum distance varies across studies noteNot all studies make use of the GIS calculated distance within their empirical models.

2.2.1 Windfarm Selection

A key decision of all papers within the literature is the selection of which windfarms will be included within the study. For all papers, windfarms are those which were operational within the study period, even when pre-operation impacts are estimated[1]. Many of the papers make a very

[1]To be considered operational, a windfarm will have been fully assembled and connected to the grid. Pre-operation is any time prior to the connection to the grid, though some papers subdivide the pre-operation status as follows: pre-announcement - the pe-

simple restriction to include all wind turbines within an administrative region including: Rhode Island (Lang et al. 2014); Massachusetts (Hoen & Atkinson-Palombo 2016); Scotland (Heblich et al. 2016); and North Rhine-Westphalia (Sunak & Madlener 2016, 2017). These papers include all wind turbines which were operational within those boundaries. Gibbons (2015) deviates slightly and includes all rural-sited windfarms in England and Wales, but drops those sited in urban areas[2]. Dröes & Koster (2016) and Jensen et al. (2014) are the only papers to perform a national analysis and include all wind turbines within the Netherlands and Denmark respectively. Other papers select either a single windfarm (Sims & Dent 2007, Sims et al. 2008, Vyn & McCullough 2014) or a group of windfarms within a specific region (Heintzelman & Tuttle 2012, McCarthy & Balli 2014) or regions across a country (Jensen et al. 2018, Hoen et al. 2011, 2015).

2.2.2 Transaction Selection

Once the selection of windfarms or wind turbines has been made, the next key decision is determining which transactions will be included within the analysis. All papers within the literature attempt to estimate the impacts on residential properties, but there are three exceptions to this rule. Firstly, Vyn & McCullough (2014) include both residential properties and agricultural land together in their analysis. Heintzelman & Tuttle (2012) also include both residential and agricultural property[3],

riod before it is publicly announced that a windfarm has been approved for construction at a given location; pre-construction - the period between announcement and the beginning of the windfarm's construction; post-construction/pre-operation - the period where construction has commenced, but the windfarm is not yet operational.

[2]Here, urban is defined using the 200m grid land classification codes for England and Wales. Gibbons also removes some single-turbine windfarms on a case-by-case basis when it seems likely they are located on an industrial estate.

[3]Residential properties are those where the primary use of the land is residential. Agricultural properties are those where the primary use of the land is agricultural - such as the growing and harvesting of crops or livestock.

but estimate their price effects separately. Lastly, Jensen et al. (2018) include exclusively residential properties but split these into primary residences and vacation homes to test for a difference in price impacts between these two groups. Sunak & Madlener (2016) attempt to estimate price effects for residential properties, but use only the value of the parcel of land surrounding the property due to data restrictions[4]. It is worth noting here that Sunak & Madlener (2016) also use 'asking price' and 'sales price' interchangeably so there is some confusion regarding exactly which price effect they are modeling[5].

Each paper within the literature makes a decision on which transactions to include based on their proximity to the windfarms within their analysis. McCarthy & Balli (2014), simply include all transactions occurring within the two towns nearby the windfarm developments of their study. For the rest of the literature, this decision involves selecting properties located within a specific distance from each wind turbine - or the center of a windfarm. For most, this involves making assumptions around where they expect any siting effects to dissipate to zero. Sims & Dent (2007) and Sims et al. (2008) are the most restrictive and include only properties within 1 mile of the wind turbines of their studies. All but two studies limit their analysis to transactions of properties up to 16km from an operational wind turbine. Vyn & McCullough (2014) include properties within 50km of a wind turbine and Heintzelman & Tuttle (2012) include properties located nearly 150 miles from the nearest turbine. Heblich et al. (2016) are unique in that they include properties located within 15km of an operational wind turbine, but also properties

[4]The property price data provided by the expert advisory boards contains arm's length transactions of properties in terms of parcels of land, which were assigned for residential utilization according to the regional land-use plan of the local administration. The price of the structure (i.e. a certain type of building) that has been built on a given land plot (if any) is not included.

[5]The asking price is the price that a property is advertised on the market, the sales price is the value of the transaction.

outside of this radius with similar characteristics which are used as a control group.

In addition to selecting the properties which are within a specified distance of wind energy developments, the analyses within the literature will generally apply either an average price (Sims et al. 2008, Hoen et al. 2011, Jensen et al. 2014, McCarthy & Balli 2014, Gibbons 2015, Hoen et al. 2015, Hoen & Atkinson-Palombo 2016, Sunak & Madlener 2016, 2017, Jensen et al. 2018) or a repeat sales analysis (Heintzelman & Tuttle 2012, Heblich et al. 2016). Average price models examine impacts to the average transaction price within a geographic region while Repeat Sales models examine price impacts between two transactions of the same property. There is a selection of papers from the literature which include both average and repeat sales analyses to check for consistency across these two approaches (Lang et al. 2014, Vyn & McCullough 2014, Dröes & Koster 2016). For the papers applying an average price analysis, the transactions are aggregated at some spatial level such as the postcode or city. This allows for a larger pool of property transactions because properties must be sold at least once to be included. Repeat sales analyses restrict the transactions to come from properties selling multiple times over the study period, which generally leads to a smaller sample of transactions within the analysis.

Once the windfarms and property transactions of an analysis have been chosen, each paper then defines the treatment. Treatment within the context of the literature is the feature of windfarm siting that is expected to impact the prices of the properties within the sample. The two main effects within the literature here are proximity and visibility. There have been some attempts to account for the effects of noise from a wind turbine on house prices, though this is assumed to be captured using a distance term and effects are expected only at the very closest proximities. Dröes & Koster (2016) state that within 400-500m the noise level produced by a wind turbine is similar to that of a typical refrigerator. In the subsequent sections I will detail the approaches to defining and estimating visibility and proximity from within the literature.

2.3 Wind Turbines as a(n) (Dis)Amenity

There are three broad mechanisms through which wind turbines may be considered environmental amenities or disamenities, and ultimately may impact upon house prices. If there is a positive price impact, wind turbines are considered an environmental amenity (they add value to properties). If there is a negative price impact, wind turbines are considered an environmental disamenity (they subtract value from properties). Hoen et al. (2011) defines the three mechanisms through which a wind turbine may be classified as a disamenity, though I have expanded these to include the possibility of a positive price impact. These include: The 'Scenic Vista Stigma', the 'Area Stigma', and the 'Nuisance Stigma'.

The Scenic Vista Stigma refers to the perception that views from properties may add or subtract considerable value to that property, based on the quality of or desire for that view. When the vista feeds through into the value of a given property, it is an amenity if the vista adds to the value of a property and a disamenity if it subtracts from the value of a property. Within the context of wind turbines, the Scenic Vista Stigma would occur if the installation of wind turbines are considered a detriment to the pre-turbine vista from a property and decrease the value a buyer places on the view leading to a reduction in the value of that property. Conversely, the addition of wind turbines to a vista could increase the value of that view, which ultimately increases the property value. The effects observed due to the Scenic Vista Stigma may be influenced by the quality and type of existing vista.

The Area Stigma is related to the Scenic Vista Stigma, but refers to the perception that the general area surrounding a windfarm will appear more developed and therefore influence the property values of the local community - regardless if whether any individual homes have views of the wind turbines. Again, it is possible for the Area Stigma to lead to positive or negative price effects, as this will reflect the preferences of home buyers towards windfarms. There may also be influences here arising from community development funds, recreational opportunities etc. (Gibbons 2015).

The Nuisance Stigma refers to the potential for disrupting factors arising from wind turbines post-construction. These are expected to occur only very close to turbines. Nuisances include factors such as noise and shadow flicker from wind turbine blades. Noise is expected to be an issue for properties within a proximity of 4-5 times the hub height of a wind turbine, assuming there are no intervening features to deafen or disrupt this noise (Dröes & Koster 2016). So, a turbine with a hub height of 100m could be expected to be heard 400-500m away. The flickering of a shadow from the turbine could project further though there is no estimate from the literature regarding the extent of a shadow. Dröes & Koster (2016) do find evidence of a negative price effect arising from the flickering shadow of turbines, though it is not statistically significant.

The price effects found withing the literature are explained as arising through these mechanisms. When creating an econometric model to evaluate the impacts then it is very important for each analysis to estimate the proximity to and visibility of a wind turbine from a nearby property. In the subsequent section, I will detail the ways that proximity and visibility have been defined, or proxied within the literature. As stated previously, any effects arising from the Nuisance stigma have thus far been proxied by distance rather than directly measuring the sound or presence of a flicker.

Many papers within the literature include both turbine proximity - to account for the Area Stigma or the Nuisance Stigma (noise) - as well as visibility to account for the Scenic Vista Stigma and Nuisance Stigma (flicker). However in the next two sections I will focus on how proximity and visibility are included separately within the literature and then discuss the application of both together subsequently.

2.3.1 Proximity

Proximity of properties to wind turbines is a key variable in the analysis of every paper within the literature with the single exception of McCarthy & Balli (2014) where proximity, while known, is not included in the model. For the rest of the literature, GIS software is used to measure proximity of homes to wind turbines. Gibbons (2015) use the center of a windfarm and the center of a postcode to determine distance. This is a method replicated by Heblich et al. (2016) under one specification of their analysis. The remainder calculate proximity from a property to the nearest wind turbine, or in some cases to all wind turbines within the maximum distance allowed by their selection criteria. Proximity to wind turbines or windfarms is primarily used as a means to estimate price effects arising from the Area and Nuisance stigmas though some papers also use it as a proxy for visibility[6].

Within the literature, there are a few different methods of defining the distance term for the analyses they perform. Firstly, some papers use a continuous measure of distance within their analysis (Sims & Dent 2007, Sims et al. 2008, Heintzelman & Tuttle 2012, Jensen et al. 2014, 2018). Of the papers taking this approach, Jensen et al. (2014) and Jensen et al. (2018) measure the exact distance from each property to the turbines which were operational during their study period, and are able to show that there is a decreasing marginal price effect from each additional proximate wind turbine. The other papers measure the distance to the chronologically first operational wind turbine. The key drawback of this approach is that when the continuous measure of distance is included within a pricing model, the coefficient is reporting the average treatment effect for the average distance[7]. The issue here is that when utilizing a continuous distance measure, the coefficient of the distance term is the effect at the average distance from a wind turbine. This is particularly problematic as these papers extrapolate their findings to represent the impacts observed at the closest distances to the turbines, where both local and property characteristics may be quite different than those at the greatest distances (Hoen et al. 2015).

An alternative approach is to classify the continuous distance measures into discrete distance bands - like rings in a target[8] - to more accurately estimate price effects within specific distances to nearby wind turbines. This is considered to impose the least non-existent structure

[6]The earliest papers within the literature use distance as the proxy for all effects (noise, visibility, etc) though more recent papers create term for each of the stigmas within their model, with the exception of noise which is generally assumed to be wrapped into the nearest distance band when these are applied.

[7]For example, Heintzelman & Tuttle (2012) include properties located 148 miles from the nearest turbine, and the mean distance from a turbine was 10 miles, approaching a point beyond the distance where any price impacts from wind turbines are anticipated.

[8]I provide an illustration of distance bands in Appendix A4

on the data, as well as being the most transparent means of estimating true proximity effects[9] (Hoen et al. 2015, Hoen & Atkinson-Palombo 2016, Sunak & Madlener 2016). The majority of papers in the literature take this approach (Hoen et al. 2011, Lang et al. 2014, Vyn & McCullough 2014, Hoen et al. 2015, Dröes & Koster 2016, Heblich et al. 2016, Hoen & Atkinson-Palombo 2016, Sunak & Madlener 2017) with Gibbons (2015) applying this to the distances between postcode and windfarm centroids rather than between properties and turbines. Sunak & Madlener (2016, 2017) also use discrete distance bands, but Sunak & Madlener (2016) do not include a stand alone proximity measure in their model. Instead they combine visibility and distance to all nearby turbines to estimate a categorical visibility impact indicator. This brings us to the next key component of the analyses within the literature: Effects from visibility.

2.3.2 Visibility

Within the literature, visibility of a wind turbine - or lack-thereof - is assumed to be a key driver of any observed house price effects, and this feeds through into all three of the stigmas described in the previous section. The Scenic Vista Stigma assumes that a price effect occurs due to changes in the views from nearby properties (for better or worse), while the Area Stigma assumes that direct views of windfarms from properties are not necessary for there to be an impact but that ultimately living in an area with views is enough for there to be price effects. Lastly, the Nuisance Stigma includes visual impacts arising from the flicker of a shadow of moving turbine blades. Despite the assumed importance of visibility to any house price effects, there is considerable variation

[9] A continuous functional form imposes structure because the researcher must decide how price is related to the underlying variables (distance in this case) through the selection of a specific functional relationship between the two.

within the literature regarding how, or if, visibility is estimated.

Broadly, the literature can be split into three categories by the approach taken to estimating windfarm/turbine visibility. The first group are the papers which do not directly measure visibility, but rather assume that it is accounted for within the proximity measure (Sims & Dent 2007, Heintzelman & Tuttle 2012, Hoen et al. 2015, Hoen & Atkinson-Palombo 2016, Jensen et al. 2018). The justification for this approach is that these papers are measuring the effects arising from the Area Stigma (Hoen & Atkinson-Palombo 2016), although some claim that properties very close to turbines will have a direct view and therefore they measure all effects from windfarms (Heintzelman & Tuttle 2012), or that because they focus on the properties very close to wind turbines, it is likely that they have a direct view (Hoen & Atkinson-Palombo 2016).

Within the literature, there are five papers which used in-person assessments of windfarm visibility by a researcher (Sims et al. 2008, Hoen et al. 2011, Lang et al. 2014, McCarthy & Balli 2014, Vyn & McCullough 2014). The number of sites visited ranges from 147 (Sims et al. 2008), to more than 1,000 (Hoen et al. 2011, Lang et al. 2014). Although there is potential for the assessments of visibility to be subjective, each study used one individual assessor to ensure consistency throughout. Among these studies, all create a visibility impact scale of some kind - ie they do not use a simple dummy variable to indicate visibility, but rather categorize the visibility into a scale from not visible to extreme visual impact. The visual impact scales may incorporate the number of turbines which are visible, the visibility of turbines interacted with proximity[10], or even whether the vistas the turbines interrupt were high or low quality.

[10]In this case, a higher visual impact score is assigned when turbines are located closer to the properties than if turbines are farther away.

Sims et al. (2008) define their visual impact categories by both the number of visible turbines and whether properties have partial or full views of the turbines. They also use the full property (ie, front, back, sides) to determine whether there are visible turbines. Hoen et al. (2011) take a similar approach and their visibility scale includes the number of turbines visible, as well as the viewing angle. They also created a ranking scale for the vista excluding the wind turbines to compare effects from turbines which interrupted 'good' views to those which interrupted 'poor' views. Lang et al. (2014) also use a visibility scale, though rather than creating a direct a ranking system, they estimate the percent of the property from which turbines are visible and use this to classify properties into four categories of visibility. They generated these by what could be seen from the street in front of the house, as well as some walking in both directions away from it.

Vyn & McCullough (2014) also take the site visit approach to generate a visibility scale, though they only classify the visibility of the nearest wind turbine, and visibility only accounts for visibility of the blades rather than the tower. Lastly, McCarthy & Balli (2014) are unique to the literature in that they use both site visits to assign a visibility impact scale as well as applying an assessment using GIS software. Though their analyses using the GIS and site-visit visibility measures show no differences, they do report that there was a roughly 5% misclassification of visibility in the GIS model, largely due to intervening flora and buildings.

The final category of visibility estimation includes those which apply GIS techniques to model visibility. This is the most popular means of accounting for windfarm visibility within the literature, as for most studies it was not feasible to physically visit each individual property or location within the datasets used. The most basic visibility estimation technique within the literature involves generating viewsheds or line-of-sight estimation from properties to wind turbines using a Digital Elevation Model (DEM). This is the approach taken by Gibbons (2015) and McCarthy & Balli (2014), both of which use a 200 meter grid resolution. DEMs represent surface elevation on a grid, where each tile contains the average elevation within each grid-tile. DEMs are essentially topographical maps which can be used to determine if a property has a view of a wind turbine, or if the view is obstructed by geological features such has hills or mountains. Therefore higher resolution (smaller tiles) DEMs will generate more accurate visibility estimations[11]. Jensen et al. (2014) utilize a 1.6 meter resolution DEM, which is the most granular analysis taking this approach within the literature.

However, DEMs only account for topographic features and on their own, do not account for features such as buildings or trees. Within the more recent literature, many papers have augmented their DEMs with building height data to take account for lines of sight or viewsheds where wind turbines are blocked by buildings Jensen et al. (2014), Dröes & Koster (2016), Heblich et al. (2016). For Jensen et al. (2014) and Dröes & Koster (2016), this does not require restricting the data in their sample as they were able to access building height data covering the entire areas of their analyses. Heblich et al. (2016) restrict their analysis which includes building heights as there is currently not data available for the whole of Scotland. One final augmentation to GIS visibility analyses is the inclusion of land cover such as trees, which may also block views of wind turbines. Currently, only Jensen et al. (2014), Sunak & Madlener (2016), and Sunak & Madlener (2017) do so in their analyses. To achieve this, they make use of Digital Surface Models (DSM) which include this

[11]Figure 4.10 provides a comparison between DEMs at differing resolutions.

information, and are the most refined level of visibility estimation set ups possible within GIS. The key difference between a DSM and a DEM is that DSMs represent the elevation of topographic, flora, and building objects on the Earth's surface. A DEM represents only the topographic elevation.

One may be under the impression that including building heights or flora will improve accuracy of the visibility modeling. This is true under some circumstances, but unfortunately, there is a potential drawback to the utilization of these data. This is due to the potential for a DSM or a DEM augmented with building heights to no longer represent the true state of the landscape. None of the papers in the literature restrict their analyses to be within ± 3 years of the date that the building height data or DSMs are published as is recommended to ensure the accuracy of the data (CEDA 2020).This is simply due to the nature of changes in both the building stock or flora within the data, where outside of this range the data may no longer reflect the past or present state of things. This is particularly problematic for the papers whose analyses contain very long study periods such as Heblich et al. (2016) or Sunak & Madlener (2016).

2.4 Analytical Approaches

In the previous sections, I have outlined the approaches within the data to determine the size or area of the studies, the type of transactions used to estimate prices, as well as the differing approaches to measuring windfarm proximity and visibility. There are a variety of approaches to building the dataset which is ultimately used to estimate price impacts on housing transactions from wind energy developments. Regarding the statistical approaches within the literature, there is a similar level of variety of econometric models to which these datasets are fitted and ultimately generate the discordant findings in the literature. In this section I will discuss these econometric approaches from the literature, but I do focus on those most relevant to the empirical analyses of Chapters 3 and 4.

$$Price = Treatment + Characteristics + Location + Time \tag{2.1}$$

I provide a generalized version of the Hedonic Models applied within the literature in Equation 2.1. This is to highlight the key variables included in the analyses whether they take the form of a standard hedonic regression model or the more complex Difference-in-Difference models. Each of these variables are expected to influence the transaction prices of homes, first and foremost is the treatment by windfarm or wind turbine siting. As described in the previous sections, treatment may include windfarm proximity, visibility, or a combination of the two. When the coefficient for this term is negative and statistically significant, this indicates that wind turbines may be a disamenity - if it is positive and statistically significant, the opposite is true. The property characteristics are also expected to be an important contribution to the transaction price, and these may include construction type or information about the size or number of rooms of a property. Location is also important, firstly to control for local fixed-effects, but also locational features such as crime or school quality, or proximity to recreational activities. Lastly, it is important to incorporate the time of transactions or treatment into the model as prices may fluctuate over time due to shocks such as the Financial Crisis of 2008.

Within the wider hedonic literature, the standard hedonic regression model is the workhorse for revealing the value buyers place on property features or characteristics. However, there are some drawbacks to this approach, mostly regarding the potential for spatial auto-correlation, and the presence of fixed effects that these basic models do not account for. Within the windfarm amenity literature, only Sims et al. (2008), McCarthy & Balli (2014), and Hoen et al. (2015) use this as their econometric approach to valuing the house price effects from windfarm siting. Though the nature of these two papers being essentially case studies, examining the impacts of one windfarm (Sims et al. 2008) and two windfarms on two towns (McCarthy & Balli 2014), the models are unlikely to suffer too much from these issues.

Other papers within the literature augment the basic OLS regression models to account for spatial fixed effects, these include simply including fixed effects within the OLS model (Heintzelman & Tuttle 2012), including a Generalized Method of Moments Estimator and a Spatial Autoregressive Error term Jensen et al. (2014), or applying a Generalized Additive Model with fixed effects (Jensen et al. 2018). These models account for spatial and temporal fixed effects, by including quarter and location dummy variables which should reduce bias in the estimated coefficients, but there are more robust econometric methods within the literature. For example, Vyn & McCullough (2014) apply a Hedonic Box-Cox model with a Spatial Autoregressive Error while others apply a difference-in-differences model.

The fixed effects DID model is the preferred econometric tool for estimating the price impacts of windfarm siting on property values because it generates far more reliable results than the standard OLS models. This is largely due to the DID models simulating a Randomized Control Trial, or natural experiment via comparing a treatment and control groups to measure the differences in outcomes for the two groups (Callaway et al. 2018, Athey & Imbens 2018, 2021). Within the context of the literature, treatment relates to the siting of a nearby windfarm and the control group are those untreated. I will discuss treatment and control groups in more depth in the next section. Ultimately, a DID compares the pre-treatment values of properties in both the treatment and control group to the post treatment values generating estimates of windfarm siting (turbine proximity and visibility)on nearby house prices.

2.4.1 Spatial Fixed-Effects Difference-in-Differences Models

The majority of the most recently published papers have chosen to implement a spatial fixed-effects DID model to estimate the house price impacts from windfarm proximity or visibility (Hoen et al. 2011, Lang et al. 2014, Gibbons 2015, Dröes & Koster 2016, Heblich et al. 2016, Sunak & Madlener 2016, 2017). Within any DID framework, there is an interaction between the treatment - regardless of what the treatment ultimately is defined as - and the time that the treatment becomes active. In this section I will discuss the differences across this subset of the literature and how they have implemented the DID framework, and provide a generalized version of this model below in Equation 2.2.

$$\begin{aligned} Price = (Treatment * Time\, of\, Treatment)\ +\ Characteristics\ + \\ Location\ Fixed - Effects\ +\ Time\, Trends \end{aligned} \tag{2.2}$$

In its generalized form, the variable of interest in the DID framework incorporates both treatment and the time that treatment occurs. Again, treatment could be turbine proximity, visibility, or a combination of the two. The time of the treatment varies across the analyses within the literature, but is generally defined as the date when a wind turbine (or the windfarm it belongs to) becomes operational. However some papers test for robustness to alternative definitions of treatment timing such as after construction is announced but prior to operation. There is also some variety in regards to the chosen control or comparator group within the literature. Characteristics are again incorporated into the DID model. To control for unobserved features which may influence house prices, but which vary across both time and space, these models incorporate Fixed-effects and Time Trends into the model. This is to account for any features or time trends which may be related to the treatment by windfarm proximity or visibility.

Hoen et al. (2011) define treatment as being within 2 miles of a visible wind turbine, but they perform three separate analyses within this framework which differ through the timing of this treatment. In the first model, treatment occurs when it is publicly announced that a windfarm will be built, the second model sets treatment at the time construction work begins, and the final model sets the treatment timing at the date the windfarm becomes operational. This allows for an estimation of anticipation effects (ie, home buyers would not see a wind turbine at the announcement date, but may anticipate an impact on the views from the home). In each model, the set of control properties are those located 3-10 miles from the nearest wind turbine and in the post-operation model the control set includes the properties which will ultimately have a line of sight to a turbine. Lang et al. (2014) take a similar approach, with the treated group being properties within 2 miles of an operational visible wind turbine, though the set of control properties are those between 5 and 10 miles from the nearest turbine. They set the time of treatment to be at the date the wind turbine becomes operational. Both of these analyses allow for the treatment itself to be a categorical variable depending on the level of visual impact as described in Section 2.3.2.

While all other papers in the literature define treatment at the property level, Gibbons (2015) and Dröes & Koster (2016) do so at the postcode level, as do Heblich et al. (2016) under their initial analysis. For Gibbons and Heblich et al., treatment occurs when a nearby windfarm becomes operational within one of the discrete distance bands of the analysis. However, unlike any of the other papers in the literature, these papers perform two separate DID analyses. The first analysis defines the treatment group as postcodes where windfarms are visible, and the second where windfarms are not visible.

Here, the control properties are those which will - by the end of the study period - be treated by an operational windfarm within 14km but have not yet been treated. The time of treatment is the time that the first windfarm becomes operational. The separate DID models are ultimately combined to estimate at triple difference to generate the difference in price impact between the two treatment groups. Dröes & Koster (2016) apply a very similar approach, but the treated postcodes are limited to those within 2km of an operational wind turbine, and the control are those between 2 and 8km from the nearest wind turbine.

Sunak & Madlener (2016) and Sunak & Madlener (2017) also apply a definition of treatment and control properties which which is very similar to that of Gibbons (2015). The treatment group are properties with a direct view of the wind turbine, and are within 2km of it. The timing of the treatment occurs when the wind turbine becomes operational, and the control properties include those which have not yet been treated, but by the end of the study period will receive treatment. Though unlike Gibbons (2015) controls also include properties which will never have a view of the wind turbine.

In their main analysis, Heblich et al. (2016) define treatment at the property level, but maintain the same definition as their and Gibbons (2015) average sales analysis - simply applying this to properties rather than postcodes. Treatment properties are those within a distance band of an operational windfarm, again separating visible and non-visible treatment groups. However, unlike the other studies within the literature the control properties are not limited to those within a given distance. Instead they build a set of control properties based on these having a similar set of characteristics or internal features to those properties which ultimately receive the treatment. The only difference being that these properties are never exposed to an operational windfarm.

2.5 Conclusions

In this chapter, I have reviewed the well-developed literature examining house price impacts from windfarm proximity and visibility. The literature itself has not settled on the direction, nor the size, of any price effects on nearby homes arising from windfarms. As I have outlined within this chapter, there are several key analytical decisions that any research into this area must make, and within the literature there is little consistency regarding any of these steps. This begins with the simple selection of the properties which will be analyzed, whether the analysis will take the form of a repeat sales or average price analysis. The research decisions diverge further with the decisions regarding the classifications of proximity and visibility - though there is agreement within the literature that these will be the key drivers behind any observed price impact, there is no agreement in the approach to modeling either. Lastly, the econometric approaches within the literature vary greatly. Though there does seem to be a converging preference towards a Difference-in-Difference empirical framework, the definitions of both the treatment and control groups differ within each study. As such, it is unsurprising that the findings of these analyses range from a house price decrease of 14% to an increase of %3.

In the next chapter of this thesis, I attempt to better inform the literature regarding the influence of these differences across research decisions and their impact on the findings. I follow the basic framework laid out by Gibbons (2015) and firstly replicate that analysis to generate a baseline for comparison. I then extend this analysis to include additional data unavailable at the time this article was published, and then begin to test the analysis for robustness to alternative assumptions from the wider literature. This analysis is then used as a comparison to the empirical work of Chapter 4 where I apply a repeat sales analysis and test this framework as well. This literature review has served to inform the alternative assumptions tested as well as the best practices for an analysis within the area.

Chapter 3

House price effects of windfarm siting in England and Wales: An average price analysis

3.1 Introduction

Wind energy is one of the key tools in the transition to a low-carbon economy. This due to the increasing maturity of the technology and its cost advantage over other renewable energy generation technologies, and indeed over some non-renewable energy generation technologies (IEA 2020*a*). The need to rapidly transition away from a carbon-based economy has led to substantial investment in wind energy, and the United Kingdom has led the charge. In 2000, 946 GWh of electricity was generated by wind turbines, and by 2018, this had increased by nearly 57 fold to 56,906 GWh (OFGEM 2020*b*). This has been hailed as a successful start to decarbonizing the UK economy. The public has remained broadly supportive of offshore and onshore wind (BEIS 2020*a*), and the government has made increasingly ambitions emissions reductions targets (PMO 2021). Yet, there is some evidence that despite this support, windfarm siting may be negatively affecting property values when properties are located near and have views of nearby wind turbines.

As discussed in the previous chapter, there has been growing academic interest in exploring the price effects of windfarm visibility and proximity on nearby house prices, though the literature itself disagrees on the size and direction of such effects. The peer-reviewed literature examining these effects within the UK has also found differing impacts ranging from price decreases, no price impact, to price increases. Sims & Dent (2007) found no effect in Wales, Gibbons (2015) found significant negative price effects in England and Wales, and Heblich et al. (2016) found evidence of positive and negative effects which were not consistent with distance from the windfarms. In addition to the disagreeing evidence of these analyses, there has been a rapid change in the landscape – for

England and Wales particularly – with a substantial increase in both the number of wind turbines operating, but also the size of such new turbines.

Figure 3.1 below highlights the rapid increase in windfarms between 1997 and 2017, Figure 3.2 shows the growth in the total number of operational wind turbines, and Figure 3.3 shows the increase in generation capacity over the same period. Across these Figures, the extent of wind turbine deployment and generation capacity as it stood during the period analyzed by Gibbons (2015) is shown by the solid line, and shown by the dotted is the period after. At the end of 2011, there were 1,551 operational turbines in England and Wales. By the end of 2017, there were 3,746; an increase of more than 200% over a period of 6 years. Over the same period, there was a 300% increase in generation capacity reflecting the fact that the size of installed turbines also saw a sharp increase both in height and blade diameter. Figure 3.4 shows this development geographically[1]. The changing landscape in England and Wales taken together with the disagreement in the wider windfarm visual amenity literature warrant further investigation into the relationship between wind turbines and house prices.

[1] Appendix A2 presents the geographic distribution of windfarms showing the areas within 14km of each windfarm, highlighting the increased share of locations within England and Wales which are potentially affected by windfarm siting.

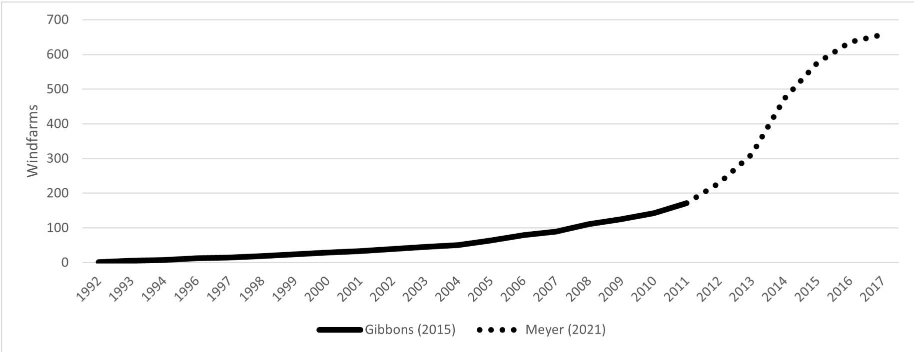

Figure 3.1: Operational Windfarms in England and Wales: 1992-2017

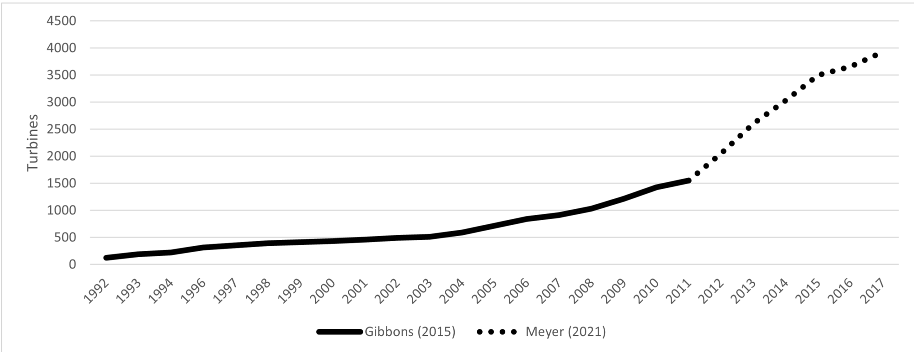

Figure 3.2: Operational Wind Turbines in England and Wales: 1992-2017

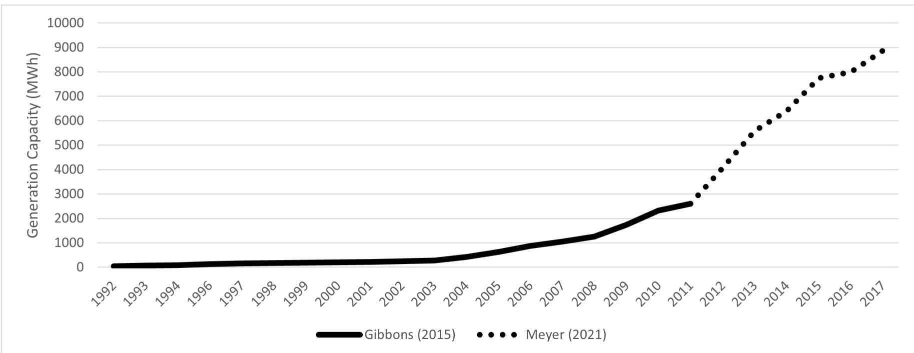

Figure 3.3: Generation Capacity (MWh) in England and Wales: 1992-2017

To this end, this chapter analyzes the interaction between windfarm proximity and visibility on nearby house prices. This is done firstly by replicating Gibbons (2015) which makes use of an average hedonic pricing approach to value the effects of windfarm visibility on house prices. After replicating the paper by Gibbons, I apply the same analytical methods to an extended dataset, which includes 12 additional years of property transaction data, and 6 additional years of windfarm siting data. I then test the results for sensitivity to a variety of alternative specifications or datasets. This analysis is, at the time of writing, the largest average price HPM analysis of its kind. The analyses performed here will contribute clarity to a wider literature which is largely in disagreement as to the direction and scale that windfarm proximity and visibility affect house prices.

This extension of the analyzed period allows for several important contributions to the wider literature on windfarm- visibility and its effects on house prices. Because we assume that it is home buyer's preferences towards windfarm visibility that lead to the reduced postcode-average prices presented in the replication section, it is crucial to test whether these are stable across time. Sunak & Madlener (2017) found that reduced property values arising from windfarm visibility was persistent over at least ten years. However, their analysis reported effects arising from three windfarms. It is therefore important to test if this is the case on a national rather than local scale. This is because not only did the visibility of these windfarms remain stable over time, so did the availability of properties with no view of the windfarms in their study. In the wider environmental amenity valuation research, a meta-analysis of price effects from urban rivers used a treatment dummy to differentiate effects from studies published prior to and after 2000 Chen et al. (2019). They found that willingness to pay estimates were signif-

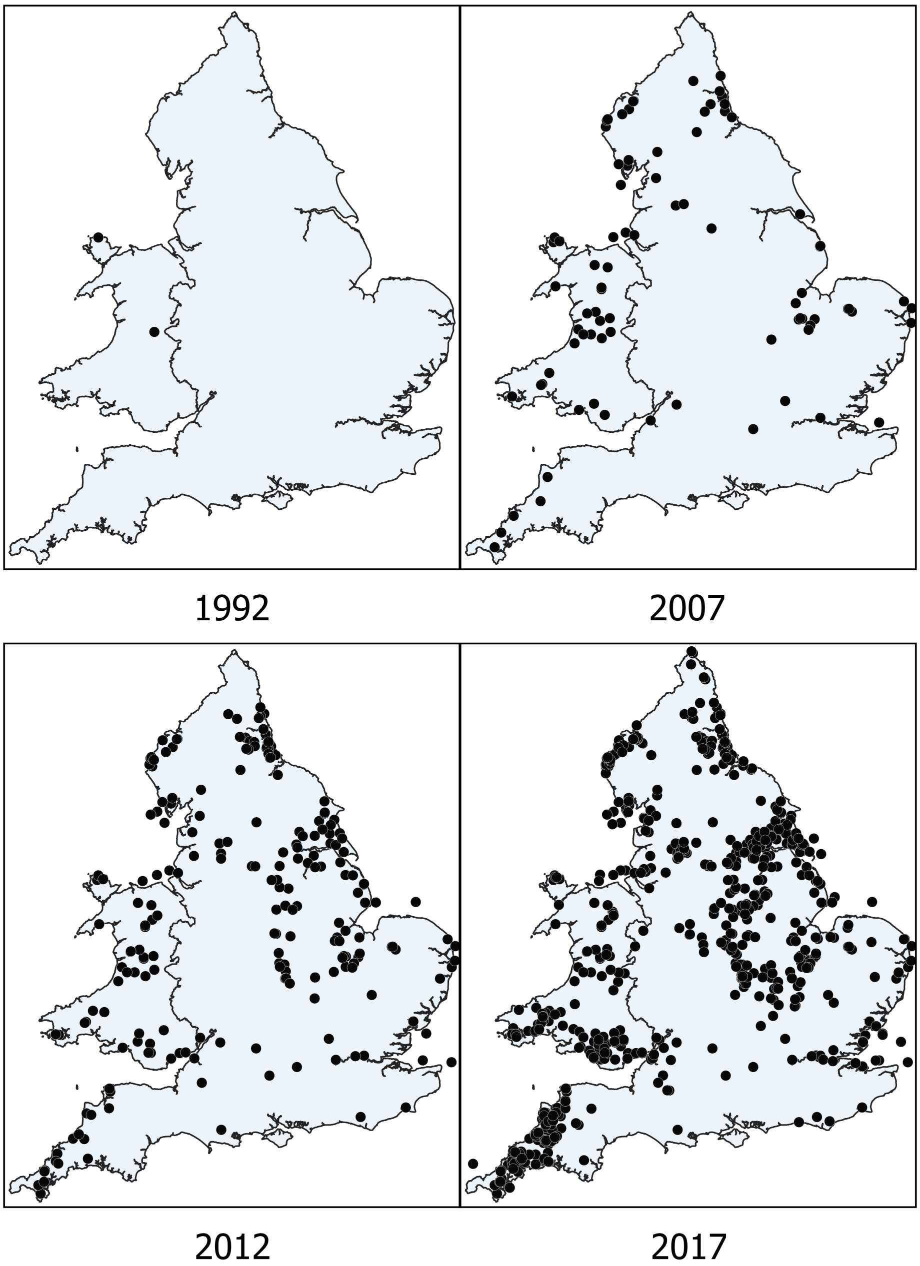

Figure 3.4: Geographic Distribution of Operational Windfarms in England and Wales (1997-2017)

.

icantly higher for studies after 2000 – they attribute this to local and regional governments implementing river cleanup programs becoming more common after 2000. This increased WTP can be interpreted as home buyers showing a preference change when the valued amenity is altered – or their perception of the amenity is altered. As there has been a substantial change in the landscape with postcodes being treated by both visible and non-visible windfarms for the first time, it is important to test if our estimated treatment effect is stable over time. If the effects are stable, this would imply that the persistence in treatment effects reported by Sunak & Madlener (2016) may be generalizable to all windfarms across time. If the effects are found to be unstable, this would reflect that while individual windfarms may have stable impacts, that this is not necessarily the case for properties treated at other points in time.

Through the replication of analysis performed by Gibbons (2015), I confirm the original results. This replication is not the only objective or contribution of this chapter, but it generates a useful baseline set of results which are subsequently tested to alternative analytical assumptions as well as additional data. Gibbons' analysis includes windfarms which became operational between 1992 and 2011. Since 2011, the number of operational wind turbines in the UK has nearly doubled, as has the area of the UK which is within 14km of an operational windfarm. This rapid growth justifies revisiting the effects found in previous research, as more and more housing transactions occur in locations with nearby wind farms. This means that an increasing share of residential properties will be situated near windfarms, and therefore an increasing number of properties will be treated by windfarm proximity and visibility. In fact, the number of postcode-quarter transactions occurring between 1995 and 2018 is over four times larger than the period

2000 – 2012. In addition to analyzing a larger dataset, I determine if the observed effects are consistent over time and robust to alternative analytical assumptions. Beyond the replication and extension, this chapter presents a series of robustness and sensitivity checks which were not presented in the replicated paper.

3.2 Contextual Background

Rapid development of wind power generation since 2000 has generally followed the pattern of the sites most optimally suited to generating wind energy being developed first, and the spread to areas with less wind potential being developed as they become the 'next best' siting location (OECD 2012). This is essentially an interaction of wind energy potential, cost of construction and cost of transmission. Sites which would generate the greatest investment returns are generally developed first, i.e., windy sites, near energy consumers, and which are relatively low-cost to construct windfarms will generally be the preferred sites (OECD 2012).

Of course, there is a technological pattern as well, where increased efficiency will open previously unproductive or unprofitable locations to wind development. Importantly, there is also a policy dimension to windfarm siting. Developments must receive planning permission from local authorities, and in some cases these may be rejected – though rejections can be appealed to and overridden by the central government. Broadly, these developments tend to occur in sparsely populated areas with large wind energy resources, however, the siting of windfarms and turbines have slowly encroached on more densely populated areas. This greater distribution across space implies that wind farms are encroaching on residential areas and therefore increasingly becoming a characteristic of homes that buyers consider when making purchases (Heblich et al. 2016, OECD 2012).

Determining whether wind turbines should be considered an environmental amenity or disamenity has important implications for energy policy, environmental policy, and the housing market. To some extent, this has been recognized by the UK government, which prescribes locations where wind energy development is and is not appropriate. If wind farms cause a substantial decrease in house prices, this implies that part of the siting and planning process of windfarms should entail payments to the affected individuals to compensate for the loss of wealth that this could cause. Similarly, if a positive impact is found on house prices, this may encourage many communities to increase their support for wind energy developments locally and may reduce delays faced in the construction and operation of such farms.

3.2.1 Wind Energy And House Prices

When surveyed, the people of the UK are rather supportive of wind energy development as well as supporting the siting of developments within their 'local area' (BEIS 2020*a*) . However, when considering the evidence from hedonic analyses in England and Wales, there may be a 'but not in my back yard' attitude that accompanies this support. Essentially individuals tend to support the growth and installment of wind turbines but prefer that this development occurs out of sight and therefore out of mind. If we follow the literature and assume that price effects will occur near to windfarms, there is potential for the costs and benefits of windfarm installation to accrue at different spatial scales. The reduction of carbon emissions that wind energy provides accrues globally. It could be argued that increasing domestic wind energy generation increases energy security and affordability, which are benefits which accrue nationally. However, a reduction in home values near windfarms is a cost which accrues only in areas near to windfarms. It is this unequal distribution of costs and benefits which have driven the academic interest in this topic area.

If there are negative price effects, individuals may lose some of the equity in their homes, which could be a significant share of their wealth (ONS n.d.). Of course, reducing C02 and other emissions is of key importance both economically and environmentally, but these individuals would essentially have a global share of the benefits of reduced emissions whilst themselves experiencing the lion's share of the costs. It is therefore highly important to accurately and reliably estimate the effects that windfarms have on their local property values.

As the deployment of windfarms has increased rapidly since 1995, with approximately 30 gigawatts of installed capacity to over 280 gigawatts in 2013 IEA (2020*a*). The impacts of these developments have been studied with increasing scrutiny from economic, environmental and social perspectives OECD (2012). These have considered the wider impacts of wind turbine and windfarm developments on variables such as employment and income, but more recently research has been conducted to test for and estimate price effects from windfarm siting. Although the work of Gibbons (2015) is the focus of this chapter, it is not alone in the application of a revealed price approach to estimate price effects.

Revealed preference methods use the actual choices that individuals make to place a value on the amenity in question Blumenschein et al. (2008). The Hedonic Price Method is one of the most common revealed preference approaches for determining the value of non-market environmental amenities in housing transactions. This is because environmental characteristics can be considered characteristics of the property, and contribute to the price at which buyers are willing to purchase a house Monson (2009). This could include the age, size, number of bedrooms, etc. that the house contains. However, the price may also include aspects such as the surrounding scenery. The Hedonic Pricing Method provides information on the value that individuals place on certain characteristics or amenities, via the price premium that certain amenities have over others Lang et al. (2014). The generic formula for determining house prices is given below:

$$Price = f(intrinsicfeatures, extrinsicfeatures, time)$$

Hedonic pricing approaches allow for the decomposition of property values into a variety of components, each with their own attached value which taken together equal the sales price of a property transaction. These values are revealed through the prices paid by home buyers. Determining the marginal value of individual property characteristics that home buyers are willing to pay or be compensated for is the key value of applying any hedonic analysis. When this is applied to features, those which increase the value of a home are considered amenities, those which decrease the value of the home are considered disamenities Monson (2009).

Intrinsic features refer to characteristics such as the number of bedrooms, bathrooms, lot size, etc. Extrinsic features are characteristics of the location where a property is situated, such as the crime rate, school quality, or as the focus of this chapter – the presence of nearby wind turbines. Lastly, the timing of a sale is an important factor which this chapter as well as much of the literature exploit to generate estimates of price effects from windfarm siting. This is due in part to time-specific price trends across the housing market itself, but we exploit the fact that transactions may occur before and after windfarms become operational. Thus the analysis of this chapter estimates how the change in the characteristics of properties (the novel presence of a windfarm) is factored into the price.

3.2.2 Windfarms, House Prices, and the United Kingdom

I now move on to discuss the two papers which apply similar hedonic approaches to value visual amenities or disamenities from windfarms within the UK. The first paper, by Heblich et al. (2016) evaluated the price effects in Scotland and found a mixed set of results. The second, Gibbons (2015), explores these effects in England and Wales and finds consistent and substantial negative price impacts. These two papers apply very similar methodological approaches, but with several key differences in their definitions of the treated and control groups – which may explain some of the differences that they find. Both papers are unique to the literature insofar as they seek to value both the price of windfarm proximity and visibility by comparing effects from proximity to nearby visible and nearby non-visible windfarms. Additionally, Gibbons (2015) provides a triple difference estimate to compare these two price effects.

Heblich et al. (2016) evaluated the impact that windfarms had on nearby house prices in Scotland. The analysis found results that suggest some areas of Scotland experience price increases from windfarm visibility and decreases from lack of visibility. These findings differ from the rest of the literature which applies similar methodological approaches as well as the wider windfarm-amenity literature. The study used house price paid data from 1990-2014 and windfarm data for windfarms built between 1995 and 2014. The results of this study were robust to a variety of methodological approaches and reliability tests.

The analysis was conducted using a Repeat Sales approach where individual homes were linked to individual turbines or entire windfarms, as well as an average price approach following Gibbons (2015) Their viewshed calculation included natural landscape and built infrastructure which could obscure views. The results, which found that at the 2-3km range, house prices where windfarms were visible increased by 2% compared to similar homes where the windfarms were not visible. Interestingly, at other distance bands, it was found that the lack of windfarm view was associated with price decreases. Additionally, they find that some of the more rural regions of Scotland experienced insignificant decreases in property prices.

Gibbons (2015) examines the price effects of windfarm proximity and visibility in England and Wales, using similar methods but ultimately finds severe negative price impacts from windfarm siting. The analysis includes modeling for proximity alone, visibility, and both proximity and visibility at different spatial scales. Gibbons (2015) defines treatment as occurring once a windfarm is operational within a given distance from a postcode. He models windfarm visibility and lack of visibility separately. Once a postcode is treated, it remains treated through to the end of the study. Here, the control or comparator group are postcodes which have not been treated by operational windfarms – yet. This essentially means that the treatment and control groups are the same sets of postcodes –

by the end of the study all postcodes are treated.

This removes any potential bias from the arbitrary selection of a control group, ensuring that there is comparability in the spatially fixed-effects between the control and treated groups. Gibbons (2015)s defines two DD analyses one which estimates the effect of treatment by visible windfarms, and the other which estimates the effect of treatment by non-visible windfarms. He finds that treatment by visible windfarms lead to negative price effects, whilst treatment by non-visible windfarms leads to positive price effects. These results are then compared by performing a triple-difference (DIDID) analysis. which shows the relative difference between these two treatment groups, and suggests that the DID models alone may underestimate the true price impact from windfarm visibility.

The findings of this paper were that for visible-operational windfarms within 1km was -6.32%, within 2km -6.28%, 4km -3%, 8km – 1% and within 14km found a negative, but insignificant impact. When these were split into distance bands, similar results were found: 0-1km -5.39%, 1-2km, -5.78%, 2-4km, -1.93%, 4-8km, -1.04% and 8-14km -.5%. All of which were statistically significant results. These are some of the largest price effects within the literature, particularly for properties which are relatively distant from windfarms.

In the subsequent sections of this chapter, I follow the methodological approach laid out by Gibbons to assess the hedonic value of windfarm visibility and proximity on house prices. The first analyses are an attempt to replicate the work by Gibbons, and the second set of results extends this analysis to include a considerably larger dataset as well as to investigate how changes to the underlying data impact upon the results. The primary investigations revolve around updating the analysis to include 12 additional years of data, apply alternative visibility estimation techniques, and testing for robustness to key assumptions around data inclusion and exclusion criteria.

3.3 Methods

To determine the effects of windfarm siting on house prices, the analysis performed in this chapter applies a Fixed Effects, Difference-in-Difference (DID) model in a Hedonic Pricing context. The model itself is borrowed from Gibbons, but I test it to a variety of alternative datasets to both improve the estimation and test for sensitivity to alternative visibility estimates, and data exclusion criteria. In this section, I describe the estimation strategy and the data underpinning the analysis. I then apply this DID framework to a replicated dataset and subsequently to an extended dataset. The spatial FE model is applied here because there may be unobserved characteristics of locations suitable for windfarm siting which also impact house prices - this arises from spatial dependence, must be controlled for within the model.

A key requirement in the application of a DID estimation is that both the parallel trends and parallel shocks assumptions be met (Athey & Imbens 2018). The analysis here achieves this by restricting the analysis to property transactions occurring within postcodes that are comparable in their suitability for wind energy developments. Essentially it is assumed that the characteristics that lead to windfarm siting are similar across regions, and the price effect estimated here represent the change in windfarm visibility. Postcode fixed effects are removed in Equation 3.1 by using the within groups estimator and common time effects are removed by using quarter-specific dummy variables. Here I am accounting for the similarities in house prices that we would expect to see as a virtue of a property being near an area suitable to wind energy developments. The time effects account for external factors that are time variant, i.e. the financial crisis of 2008 which caused a significant drop in property.

Figure 3.5 shows the natural log of the average transaction price for properties within 14km of a visible windfarm (dotted line), within 14km of a non-visible windfarm (solid line), and properties more than 14km from a wind turbine (dashed line). The same trends are presented for the mean transaction price in ix A7. The analysis of this chapter makes use of only properties within 14km of a windfarm and these track each other quite well showing that the parallel trends and shocks assumptions are not violated and therefore the data is suitable for the application of the DID model. The set of properties located in postcodes more than 14km from an operational windfarm have a much higher average price, and the yearly growth in price does not follow the same path as those near to windfarms.

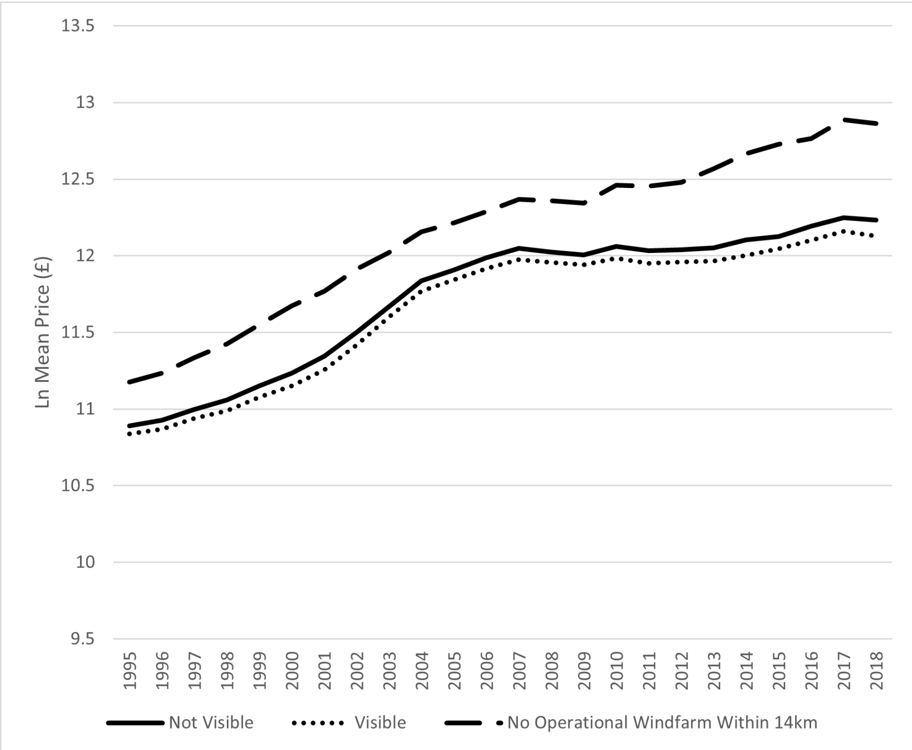

Figure 3.5: Yearly Transaction Prices for Properties within 14km of Visible and Non-visible Windfarms (1995-2018)

3.3.1 Estimation Strategy

I follow the estimation strategy applied by Gibbons (2015). This makes use of a Fixed Effect DID, as well as a difference in difference in difference (DIDID) estimator at the later stage of the analysis. The original work did not have information on planning permission or construction dates for the windfarms analyzed, so it is assumed that the date of operation is taken as the date around which the price effects are expected to take effect. However, I make use of an improved dataset and test for anticipation effects post announcement and prior to operation and find a negligible effect. This serves as an event study and enhances the reliability of the estimated effect arising from operation and the effects are reported in Table A5.

The following basic equation is the starting point for the DID/fixed effects regression specification it is worth noting that this equation is used separately to generate estimated price effects for visible and non-visible windfarms. Equation 3.1:

$$\begin{aligned} ln(price)_{it} = \sum_k \beta_k (visible, j_k < dististance < k, operational)_{it-1} \\ + x'_{it}\gamma + f(i,t) + E_{it} \end{aligned} \tag{3.1}$$

Where:

- $ln(price)_{it}$ is the average housing transaction price in postcode i in quarter t.

- $(visible, j_k < dististance < k, operational)_{it-1}$ is an interaction dummy indicator which captures exposure to windfarm developments. With a value of 1, this indicates that postcode i has at least one visible - or non visible -operational windfarm between j_k and k kilometers in the preceding quarter.

  $visible$ is the visibility indicator. When effects from visible windfarms are estimated this takes the value 1 if a windfarm is visible from postcode i. When effects from non-visible windfarms are estimated, this takes the value of 1 if a windfarm is not visible from postcode i.

  $j_k < dististance < k$ is the distance indicator – under the first specification distance takes the form of radii, and in the second distance is defined in bands to test allow for a more refined estimate of proximity effects.

  $operational$ is a post-policy indicator which indicates whether the turbine(s) have been built and are operational.

- β_k, the coefficient of interest is the average effect of windfarm turbines visible within distance band jk-k on housing prices. The sign (+-) of β_k is unknown a priori because it will be influenced by home buyers' preferences for views of windfarms, the impact of noise or visual disturbance, potential gains or losses from shares in profits, community grants, employment or other impacts related to windfarm proximity. Because this is a DID model, if the coefficient is negative, this means that prices have fallen relative to the comparator group. If the coefficient is positive, prices have risen relative to the comparator group.

- $f(i,t)$ represents unobserved components which may vary over time and space, and are inevitably correlated with the visibility of windfarms due to the fact that windfarms are not distributed randomly, as well as the fact that postcodes near windfarms may not be similar to postcodes that are not near windfarms. Correlation with the time effects is present because the number of windfarms increases over time, and this would create a spurious correlation between any trend in average prices over time with visibility of windfarms.

- E_{it} is the general error term.

For the first stage of analysis, Equation 3.1 compares effects at differing distance radii. Under this specification, the dummy variable $j_k < dististance < k$ is defined as jk = 0, and k = 1km, 2km, 4km, 8km, or 14km. The radius stops at 14 because wind turbines are expected to be only barely visible at this distance. The use of distance radii allow for testing whether any price effects decay with distance. However, the use of radii as the distance measure may contaminate the analysis insofar as the coefficient of interest here reports the average treatment effect for the full sample of affected transactions within that radius. To generate a more refined estimate of the influence of proximity on the size and direction of potential price effects, I refine the distance measure.

I redefine $j_k < dististance < k$ to represent discrete distance bands. By splitting the radii into bands, it is possible to analyze effects at discrete distances from wind turbines. these more closely and to see exactly where impacts are greatest. Here, there are 5 distance bands in the model: 0-1km, 1-2km, 2-4km, 4-8km, 8-14km. These distance bands are a feature of Gibbons (2015) and I maintain them both within the replication and the extended analysis. This allows for a cleaner analysis of the impacts at varying levels of proximity to windfarms. Figure A6 illustrates the distance bands used in this analysis.

Lastly, to obtain clearer estimates of windfarm visibility impacts, 3.1 is augmented with treatment indicators for postcodes within the given distance bands of windfarms, but where the turbines are unlikely to be visible from due to the lay of the land. Cleaner estimates of the impacts of windfarm visibility can be estimated by differencing these two differences. This triple difference is the relative difference in changes to price by treatment by visible windfarms relative to treatment by non-visible windfarms. This is an important test of the estimated effects of windfarm visibility because postcodes because the two sets of treated postcodes are located in geographically very similar locations. These two treated groups are assumed to be very comparable on unobserved characteristics and therefore subject to similar price trends arising through omitted causal channels. I should note here, that it is possible that differences observed at the closest range (0-1km) between visible and non-visible wind farms could arise from sound effects from windfarms and not visibility only. Equation 3.2 is shown below:

$$ln(price)_{it} = \sum_k \beta_k (visible, j_k < dististance < k, operational)_{it-1} + \sum_k \delta_k (non-visible, j_k < dististance < k, operational)_{it-1} + x'_{it}\gamma + f(i,t) + E_{it} \quad (3.2)$$

Equation 3.2 uses a treatment dummy which indicates that there are no visible windfarms at the postcode, that the postcode is at a given distance band, and the turbines are operational. In this equation, δ_k are estimates of the effects on house prices due to proximity to operational wind farms when there is no impact from the turbines being visible in the postcode. Therefore, the DIDID estimate of β_k - δ_k will provide a better estimate of the specific impact of windfarm visibility through the gap between house prices where windfarms are visible and where they are not. This allows for an explicitly inferred willingness to pay through housing expenditure to avoid views of wind turbines, as well as the estimated compensation required for willingness to accept the visual of wind farms disamenity. This gives a clearer estimate of the valuation of the visual impacts of windfarms on house prices.

3.3.2 Data

The data necessary to the analysis of this chapter was obtained from a variety of sources, but mostly match the sources used by Gibbons. All datasets are either freely available to the public or are freely available via an academic use license. I list the data and their sources a Table 3.1 below. I merge windfarm, postcode, and house price data to construct several datasets with observations in each postcode i at quarter t. Firstly, I impute windfarm and postcode geocoordinates into GIS where I then generate a linked dataset which includes distance from windfarm centroids to postcode centroids. I then generate several visibility estimates for each postcode to determine which windfarms it has a line of sight to, or is within the viewshed of. Example viewsheds are located in Appendix A4 In addition, I generate postcode level geographic controls including elevation, slope and aspect. This dataset is then merged with the property transaction dataset and further cleaned and transformed into a usable format.

Ultimately this provides a series of unbalanced panels, with gaps when there are no transactions in postcode i in quarter t. The panel includes the following variables of interest for each postcode-quarter: natural log of the average price, share of transactions by property type, postcode elevation, slope and aspect, year dummies, quarter dummies, postcode distance from windfarms, dummy indicators for postcodes with transactions which occur before or after the windfarm became operational, and visibility dummies. The replicated dataset is restricted to transactions occurring between January 2000 and March 2012 within 14km of windfarms which became operational before January 2012, while the extended analysis uses the full dataset. I make further restrictions and alterations to the underlying panel in order to perform a series of robustness checks. These alterations are described in more detail below, and are primarily related to alternative visibility calculation and the relaxation of windfarm exclusion criteria.

It is important to note that within this application of the HPM, it is only possible to measure marginal changes in price without the application of a further sorting model which is beyond the scope of this work.

Table 3.1: Data and Data Sources

Data	Period	Published	Source	Application	Access
Postcode Data	1995-2018	2018	ONS	Geocoordinates of postcode centroids used to determine distance and visibility to windfarms.	Free
Housing Transactions	1995-2018	10-Jul	HM Land Registry	Housing Type, price, and address	Free
Wind Energy					
Wind Farms	1992-2017	10-Jul	Renewable Energy Planning Database, Renewable UK	Windfarm location, turbine height, generation capacity, status, dates of operation	Academic
Digital Elevation Models					
200m DEM		2015	Digimaps	Visibility analysis and geographic controls	Academic
90m DEM		2015	Digimaps	Visibility analysis and geographic controls	Academic
Landcover					
Landcover Map 2000		2000		Windfarm Exclusion based on urban landcover	Academic
Landcover Map 2010		2010			Academic

3.3.3 Descriptive Figures and Statistics

In this section, I provide the summary statistics of the windfarms, postcodes, and transactions used to generate the results presented in subsequent section of this chapter. The replicated analysis was limited to the inclusion of windfarm siting data between 1992 and the final quarter of 2011, and housing transactions occurring over the period 2000-2012. I now extend the analysis to include windfarm siting data between 1992 and the final quarter of 2017 and housing transactions occurring between 1995 and the first quarter of 2018. In this section, I provide the summary statistics of the windfarms, postcodes, and transactions used to generate the results presented in subsequent section of this chapter. The replicated analysis was limited to the inclusion of windfarm siting data between 1992 and the final quarter of 2011, and housing transactions occurring over the period 2000-2012. I now extend the analysis to include windfarm siting data between 1992 and the final quarter of 2017 and housing transactions occurring between 1995 and the first quarter of 2018. I also show here visibility estimates under a variety of alternative specifications. Price effects are then are reported in Table 3.4 for the following datasets:

- Results as presented by Gibbons (2015)
- Replicated Analysis: Restrict to the period 2000-2012, maintain Gibbons' exclusion criteria
- Extended Analysis: Using all available data, but maintain Gibbons' exclusion criteria
- Full Analysis: No restriction Criteria
- Robustness to: House Price Restrictions, treatment intensity, and the study period.

The characteristics of the windfarms included in the subsequent analyses are presented in Table 3.2. The turbines used in more recent windfarm construction tend to be larger, and have a higher generation capacity, and more numerous than in older windfarm sites. The extended analysis of this chapter considers these changes in windfarm characteristics. Although the replicated dataset is using the process outlined by Gibbons (2015), there are some important differences between his reported statistics and those of the replicated dataset which I note here. These include the average height to tip, MW capacity, and the number of turbines.

It is also worth noting that when using the image of windfarm locations that Gibbons provides, I find several differences in the windfarm locations. If it is assumed that the image presented by Gibbons reflects the windfarms in the analysis, there are a few violations of the data restrictions: one Scottish windfarm is included[2], two windfarms which are more than 14km from any English or Welsh postcodes is included[3]. Additionally, several windfarm centroids in the North Midlands do not match the locations presented by Gibbons (2015). A side-by-side comparison of centroid locations is presented in Appendix A4.

The use of a slightly different windfarm dataset feeds into the rest of the panel building process, and therefore the results. Despite this, the replicated dataset and results remain broadly similar to those reported by Gibbons (2015). The number of postcodes included in the analysis are also broadly similar as reported in Table 3.3. Table 3.4 compares visibility estimations for the postcodes included in the analysis under several specifications. The largest discrepancies arise at the greatest distances from the windfarm centroids, which is to be expected. These

[2] Robin Rigg
[3] Walney I and II

Table 3.2: Windfarm Summary Statistics

	Gibbons: 1992-2011	Replicated: 1992-2011	Extended: 1992-2018	All: 1992-2018
		Turbines		
Mean	11.2	10.76	6.25	5.74
SD	15.4	15.2	15.9	14.63
Min	1	1	1	1
Max	103	103	175	175
		MW Capacity		
Mean	18.6	18.9	14.25	12.58
SD	39.2	35.63	53.95	48.93
Min	0.22	0.2	0.09	0.09
Max	300	300	630	630
		Height to Tip		
Mean	90.9	86.2	85.11	85.94
SD	29.2	25.57	27.76	27.99
Min	42	42	28	28
Max	150	150	195	195
Total Windfarms	148	148	589	653

are also the areas expected to have the least impact due to the visible size of a wind turbine at such distances. This table shows that particularly at the closest distances where price effects are expected to be the most felt, the visibility estimates are quite consistent regardless of the DEM and visibility calculation applied. Lastly, Table 3.5 reports the postcode-quarter transaction summary statistics. The largest discrepancy between the data as reported by Gibbons and the replication is the share property types of transactions included in the analysis. The replicated dataset includes substantially fewer flats/maisonette and substantially more semidetached properties, which are presented in table 3.5 It is worth noting that according to the 2011 and 2018 English Housing Survey, roughly 20% of UK household stock consisted of flats. This is higher in Urban areas and lower in rural areas. It is possible that the replicated data includes more rural postcodes relative to the selection by Gibbons.

Table 3.3: Postcode Counts by Distance Radius

Distance from wind-farm	Gibbons: 2000-2011	Replicated: 2000-2011	Extended: 1995-2018	All: 1995-2012
0-1km	1,142	1,149	3,146	3,751
0-2km	5,350	5,898	18,971	22,353
0-4km	20,838	22,416	79,277	105,550
0-8km	81,820	85,573	260,219	327,155
0-14km	220,669	233,473	597,846	726,356

Table 3.4: Visibility Estimates

Distance from windfarm	Extended	200m DEM	90m DEM	200m DEM	90m DEM
0-1km	3,146	3,095	3,092	3,104	3,099
0-2km	18,971	18,049	17,919	18,222	18,234
0-4km	79,277	66,593	65,800	69,748	69,177
0-8km	260,219	172,577	167,138	191,599	189,361
0-14km	597,846	295,037	282,721	351,414	344,539
Visible	X	X	X	X	X
Not Visible	X				
Viewshed Gradient	X	X	X		
Horizon Depth	X			X	X

Table 3.5: Postcode Quarter Transaction Summary

	Gibbons	Replication	Extended	All Windfarms
Observations	1,710,293	1,830,664	5,512,092	6,944,185
Mean Log Price	11.56	11.63	11.64	11.56
Detached	0.25	0.2122	0.24	0.25
Semidetached	0.07	0.3002	0.31	0.32
Flat	0.361	0.1133	0.11	0.08
Terraced	0.32	0.3743	0.34	0.34
Newbuild	0.041	0.0393	0.04	0.05
Freehold	0.849	0.8081	0.83	0.83

3.4 Results

The following section reports a series of results arising from the empirical analyses described in the previous sections. These results are organized as follows: First, I present the results of my replication of the work by Gibbons which becomes the baseline for the additional analyses tested here. The replicated results are restricted to include only postcode-quarter transactions from the period 2000-2012, and windfarms which become operational in the period 1992-2011. I then report the results of an expanded analysis which makes use of the full dataset and alternative model specifications. This extends the period of analysis to include transactions from 1995-2018, and windfarms which became operational in the period 1992-2017. I then report a series of robustness checks to test the key assumptions of the model, as well as test for sensitivity to alternative restrictions in the dataset. These extended results and their contributions are discussed in the context of the wider windfarm-amenity literature.

The replicated distance radii results are reported in Table 3.6 where I compare the effects of both visible and non-visible windfarms on postcode average prices. These results are grouped by distance and the replicated results are compared to the original findings for both effects from visibility and lack of visibility. I then move on to focus on the distance band specification which forms the backbone of the extended analysis and separately report the results of the distance radii analysis in Table 3.7. I then present the robustness tests in the subsequent section.

3.4.1 Replication Distance Radii

This section provides a set of results estimating the treatment effect, or the average change in postcode average prices after being treated by an operational visible or non-visible windfarm at a given radius. The replicated results are well-aligned with those reported by Gibbons. I find negative price effects which are generally larger at the most proximate distances to visible windfarms and become insignificant when postcodes at the greatest distances from windfarm centroids are included in the analysis.

Both the replicated results and those reported by Gibbons show negative and statistically significant price effects arising from non-visible wind farms at the 0-8km radius. In the 0-1km distance radius the results are remarkably similar - both Gibbons's and the replicated estimate that a visible windfarm at this distance will decrease property values by about 6.6%. However, the similarity in estimated impact begins to decrease for each subsequent distance radius. At the 0-2km radius, I find a nearly 1.3% larger negative impact than Gibbons. This patterns continue at the 0-4km radius, where I find a 1.4% larger decrease in transaction price. At the 0-8km radius this grows to a 3.2% difference between the replicated estimate and Gibbons'. Although the estimated impact is not statistically significant in the 0-14km radius the replicated results are still much larger and more negative. This is consistent across the replication.Broadly speaking, The Properties nearest visible windfarms sell at lower prices than those further away, and these differences are often statistically significant. These reduced prices can be attributed to the presence of operational and visible windfarms. Lack of visibility is associated with insignificant negative impacts, and a statistically significant positive impact at the 0-8km radius.

Table 3.6: Replication: Distance Radii

Radius	Visible		Not Visible	
0-1km	-0.0666*	-.0667***	-	-
(RSE)	(0.0221)	(0.0019)	-	-
Obs.	8,052	8,659	-	-
0-2km	-0.0558***	-0.0689***	-0.0611	-0.0697
(RSE)	(0.0095)	(0.0074)	(0.0609)	(0.0560)
Obs.	37,998	45,575	37,998	45,575
0-4km	-0.0244***	-0.0426***	-0.0018	0.0030
(RSE)	(0.0054)	(0.0044)	(0.0125)	(0.0031)
Obs.	150,907	182,077	150,907	182,077
0-8km	-0.0035	-0.0358***	0.0165***	0.0113***
(RSE)	(0.0029)	(0.0024)	(0.0041)	(0.0043)
Obs.	621,395	702,191	621,395	702,191
0-14km	-0.0027	-0.0252	-0.0024	-0.0014
(RSE)	(0.0017)	(0.0154)	(0.0020)	(0.0010)
Obs.	1,710,293	1,830,644	1,710,293	1,830,644
Gibbons	X		X	
Replication		X		X

Notes:
$***p < 0.001$
$**p < 0.01$
$*p < 0.05$
RSE: Robust Standard Errors Clustered at the Output Census Area
Control Vars: postcode slope-by-year, elevation-by-year, aspect by-year dummies, proportions of sales of detached, semi-detached, terraced, flat/maisonette;quarterly dummies. Results are reported to four decimal places.

3.4.2 Replication Distance Band Analysis

The results of the distance band analysis are presented in Table 3.7 which provides a side-by-side comparison of the replicated results and those from Gibbons. I present effects at each distance band for visible and non-visible windfarms. I have arranged Table 3.7 to present results of the replication beside the comparable analysis from Gibbons. Effects from visible windfarms are reported in Columns 1 and 2; non-visible windfarms in Columns 3 and 4; and the triple difference estimates in Columns 5 and 6. The replicated results again agree with those reported by Gibbons.

Table 3.7 Column (2) reports the statistically significant estimated price effect which ranges from a 6.42% to a 3.13% decrease in price for properties within the first four distance bands, up to 8km, from visible and operational windfarms before becoming statistically insignificant at the 8-14km band. Again, it was expected that the largest effect on house prices would occur at the closest distance to windfarms where a larger share of a property's viewshed would be interrupted by wind turbines. However, much like the radii specification of the model, it is the second distance group (the 1-2km distance band) where the largest effect is detected. Here, I find that property transactions in the 0-1km band sell on average, for 6.02% less after visible windfarms become operational. This effect then increases to -6.42% at the 1-2km distance band, then the effect begins to decrease with an average reduction of 3.8% in the 2-4km distance band, dropping to a 3.13% decrease in the average price at the 4-8km band before further reducing to a statistically insignificant 2.26% at the 8-14km band.

Table 3.7: Replication: Distance Bands

	Price Effect Of Presence of an Operational Windfarm					
Distance Band	1	2	3	4	5	6
0-1km	-0.0539***	-0.0602***	-	-0.0660	-	-
(RSE)	(0.0164)	(0.0017)	-	(0.0770)	-	-
1-2km	-0.0578***	-0.0642***	0.0268	0.0012	-0.0847	-0.0654
(RSE)	(0.0092)	(0.0073)	(0.0498)	(0.0491)	(0.0501)	(0.0850)
2-4km	-0.0193***	-0.0379***	0.0152	0.0113	-0.0345**	-0.0492*
(RSE)	(0.0052)	(0.0042)	(0.0105)	(0.0114)	(0.0106)	(0.0238)
4-8km	-0.0104***	-0.0313***	0.0223***	0.0091**	-0.0327***	-0.0404**
(RSE)	(0.0028)	(0.0022)	(0.0040)	(0.0041)	(0.0046)	(0.0141)
8-14km	-0.0050**	-0.0226	0.0018	-0.00241	-0.0068*	-0.0202
(RSE)	(0.0019)	(0.0153)	(0.0021)	(0.0021)	(0.0027)	(0.0492)
Geographic Controls	X	X	X	X	X	X
Fixed Effects	X	X	X	X	X	X
Gibbons	X		X		X	
Replication		X		X		X
Visible	X	X				
Not Visible			X	X		
Tripple Difference					X	X

Notes:
$***p < 0.001$
$**p < 0.01$
$*p < 0.05$
RSE: Robust Standard Errors Clustered at the Output Census Area
Control Vars: postcode slope-by-year, elevation-by-year, aspect by-year dummies, proportions of sales of detached, semi-detached, terraced, flat/maisonette;quarterly dummies. Results are reported to four decimal places.

Relative to the analysis by Gibbons, column (1) the replicated results show a larger negative impact arising from windfarm siting when these windfarms are visible. Gibbons results range from -5.78% to -0.5% price impacts and find statistically significant effects at each distance band – even the 8-14km band (though this is roughly ½ of one percent). I find that the replicated results are closest to those reported by Gibbons at the 0-1 and 1-2km distance bands, though these both differ by approximately 0.7%. The difference between the estimates again increases as the analysis moves to distance bands which are located farther from the windfarm centroids.

The replication's estimated effects of non-visible windfarms on postcode average house prices, reported in column (4) statistically significant for only the 4-8km range, which is again much larger than that reported by gibbons, which also finds statistically significant effects at the 4-8km distance band. The replicated estimate is less than half that reported by Gibbons, but both are positive. The replication presents an estimated effect at the 0-1km, which Gibbons does not report (1). This is negative, but not statistically significant. It should be noted that at the 0-1km radius, very few postcodes are near to non-visible windfarms. Again, this implies that there is a small price premium for properties in postcodes where non-visible windfarms are nearby which should be taken to mean that vistas without wind turbines become more valuable after windfarms became operational.

Column 6 present the replicated triple difference estimates, which are statistically significant at the 2-4 and 4-8km distance bands, though only at the 5% level. Comparing these estimates to those reported in Column 2 show that the DID estimates may actually undervalue the discount that home buyers demand to compensate for windfarm visibility, as the average treatment effect from visible windfarms is between 4 and 5% lower postcode-average prices relative to postcode average prices where windfarms are not visible at the same distances. At the 1-2km band, the triple difference is smaller than that reported by Gibbons, but larger at all other distance bands.

3.4.3 Replication Summary

Despite replicating the dataset and analytical approach of Gibbons (2015), the replication finds consistently larger negative price effects. There are statistically significant price effects arising from windfarm visibility, and that generally these price effects decrease as distance to the windfarm increases. Although the coefficient estimates are larger - sometimes a difference of 3%, both analyses do find negative price effects from windfarm visibility, and that these are larger at the closest proximities to windfarms. This implies that the results are robust to differences in the underlying dataset, which is in of itself a robustness check of the model. To summarize, the distance band specification of Equation 3.1 provides a cleaner estimation of the change in postcode-average prices after windfarms become operational than is produced under the distance radii specification.

An interesting difference between the replicated results and those of Gibbons (2015) is that the negative impacts of visible windfarms in the replication are considerably larger than those reported by Gibbons. A similar, though less consistent difference is also present for the impacts arising from non-visible windfarms. at the 2-4km the replicated results are very close to Gibbons', but at the other distance bands, the impact of a non-visible windfarm is considerably smaller than those reported by Gibbons.

By classifying postcodes into several discrete distance bands which don't overlap, I am able to estimate the degradation of the visual impacts of windfarm visibility as proximity decreases. This is because the main expected effect of windfarm siting on nearby properties is to alter the views from these locations, and it is assumed that there will be larger visual interference with natural views at the closest proximities. This approach is superior to the use of a continuous distance measure which would only serve to estimate the treatment effect at the average distance between a postcode and windfarm centroid. The results reported in Table 3.7 shows that the impacts are greatest in postcodes located in the nearest distance bands. Apart from the 1-2km distance band reporting a slightly larger negative effect, postcodes treated by visible windfarms saw the average treatment effect reduce as distance from the visible windfarms increased. Although the estimation of treatment effects from non-visible operational windfarms was only found to be statistically significant at the 4-8km distance band, the triple difference effects show potential underestimation of the treatment effect for windfarm visibility. I now move on to the extended analysis of this chapter.

3.4.4 Extended Analysis

In the previous subsection, I presented a series of results from the replication of Gibbons (2015) using both distance bands and radii to estimate house price effects from windfarm proximity and visibility. In this subsection I apply the same methodological framework from the analyses from the replication section to the full dataset available at the time of writing, an extension of 12 years of property transactions (6 years prior to and 6 years after the study period of the replication) and 6 additional years of windfarm data. I then perform a series of robustness checks applying alternative treatment definitions, and data restrictions.

These tests for sensitivity to alternative models and differences in the underlying data and how this feeds through to the results, and is used to better inform the analysis itself.

Table 3.8: Extension - Distance Bands Analysis

Distance Band	Price Effect Of Presence of an Operational Windfarm					
	1	2	3	4	5	6
0-1km	-0.0602***	-0.0660	0.0549***	0.0302	-0.0497***	-0.0082
(RSE)	(0.0017)	(0.0770)	(0.0183)	(0.0296)	(0.0126)	(0.0693)
1-2km	-0.0642***	0.0012	0.0290**	0.0329*	-0.0171	-0.0211
(RSE)	(0.0073)	(0.0491)	(0.0120)	(0.0173)	(0.0504)	(0.0168)
2-4km	-0.0379***	0.0113	0.0251***	0.0773***	-0.0039	-0.0054
(RSE)	(0.0042)	(0.0114)	(0.0048)	(0.0076)	(0.0028)	(0.0059)
4-8km	-0.0313***	0.0091**	0.0616***	0.0574***	0.0027*	0.0092***
(RSE)	(0.0022)	(0.0041)	(0.0023)	(0.0061)	(0.0016)	(0.0024)
8-14km	-0.0226	-0.0024	-0.0189	0.0841	0.0006	0.0008
(RSE)	(0.0153)	(0.0021)	(0.0147)	(0.0453)	(0.0012)	(0.0015)
Geographic Controls	X	X	X	X	X	X
Fixed Effects	X	X	X	X	X	X
Replicated	X	X				
Replicated Exclusion Criteria	X	X	X	X		
House Price Restrictions						
Full Dataset					X	X
Visible	X		X		X	
Not Visible		X		X		X

Notes:
$***p < 0.001$
$**p < 0.01$
$*p < 0.05$
RSE: Robust Standard Errors Clustered at the Output Census Area
Control Vars: postcode slope-by-year, elevation-by-year, aspect by-year dummies, proportions of sales of detached, semi-detached, terraced, flat/maisonette;quarterly dummies. Results are reported to four decimal places.

Table 3.8 reports the effects found in the replicated analysis to those found utilizing the entire sample available, although I do follow the same methodological approach, so here we are only seeing differences which arise from the inclusion of additional data rather than changes in the analytical approach itself. Here I present the results from the replication period of analysis beside the results under the extended period. The difference between the two sets of results is substantial. I then test for robustness to house price restrictions while maintaining the Gibbons' exclusion criteria, and lastly relax the exclusion criteria to include all windfarms and nearby postcodes including urban sited windfarms. I present the estimated coefficients for effects in postcodes with at least one visible-operational windfarm in the odd numbered columns, and those with non-visible operational windfarms are located in the even numbered columns. The replicated analysis is presented in columns 1 and 2; the extended analysis in columns 3 and 4; and the analysis which makes use of the full dataset is presented in columns 5 and 6.

When the extended set of postcode-quarter transactions which occur between 1995 Q1 and 2018 Q1, and within 14m of a non-urban windfarm is analyzed the results diverge substantially from the replicated results. Columns 3 and 4 show that when the dataset is extended to the maximum period while maintaining the same exclusion criteria (non-urban windfarms only), the estimated price effects from visible windfarms become positive rather than negative. The positive increases in price should be interpreted as the average change in price in postcode-quarterly average transaction for properties sold after a visible windfarm becomes operational relative to those which sold prior to windfarm operation. Impacts are positive at all distance bands but increase in size and significance with distance. At the 0-1km range, I find that visibility increases the average transaction price by roughly 5%, and 2.9% at the

1-2km range, 2.5% at the 2-4km band, and 6.2% at the 408km band. The effect is negative, but not statistically significant at the 8-14km band.

I find statistically significant and positive impacts at the 1-2km, 2-4km and 4-8km bands. The estimate price effects are increases of 3.3%, 7.7% and 5.7% respectively for postcodes treated by operational but not-visible wind turbines. Under the extended analysis the estimated price effects are broadly positive from treatment by both visible and non-visible windfarms.

3.4.5 Triple Difference Analysis

As a further test for price effects from windfarm visibility, I perform a triple difference analysis. Here, the estimated impacts of non-visible windfarms are subtracted from the impacts of visible windfarms at each distance band, and the results are reported in Columns 3 and 6 of Table 3.9 for the Extended Analysis and Full Dataset Analysis respectively. This essentially estimates the willingness to pay to avoid windfarm visibility. Under the extended analysis, I find that there is a statistically significant triple difference at the 2-4km and 8-14km bands. At the 2-4km band, both visible and non-visible windfarms are associated with positive price impacts, but the effect from non-visible windfarms is larger. At the 8-14km band, the triple difference is substantial - roughly 10% lower price effect from visible windfarms relative to non-visible windfarms. When triple differencing the results of the Full Dataset Analysis, I find that statistically significant negative relative price impact is observed at the 0-1km band (-4.15%) and at the 4-8km band (-0.65%). Therefore, it would seem that there is some evidence that home buyers are willing to pay roughly 4% less for properties located in postcodes within 1km of a visible windfarm, and about 1% less for properties within 4-8km of a visible windfarm.

Table 3.9: Triple Differences: Extension and Full Dataset

Distance Band	1	2	3	4	5	6
0-1km	0.0549***	0.0302	0.0247	-0.0497***	-0.0082	-0.0415*
(RSE)	(0.0183)	(0.0296)	(0.0185)	(0.0126)	(0.0693)	(0.0207)
1-2km	0.0290**	0.0329*	-0.0039	-0.0171	-0.0211	0.0040
(RSE)	(0.0120)	(0.0173)	(0.0122)	(0.0504)	(0.0168)	(0.0280)
2-4km	0.0251***	0.0773***	-0.0522***	-0.0039	-0.0054	0.0015
(RSE)	(0.0048)	(0.0076)	(0.0051)	(0.0028)	(0.0059)	(0.0043)
4-8km	0.0616***	0.0574***	0.0042	0.0027*	0.0092***	-0.0065*
(RSE)	(0.0023)	(0.0061)	(0.0034)	(0.0016)	(0.0024)	(0.0030)
8-14km	-0.0189	0.0841	-0.1030***	0.0006	0.0008	-0.0002
(RSE)	(0.0147)	(0.0453)	(0.0299)	(0.0012)	(0.0015)	(0.0012)
Visible	X					
Not Visible		X				
Triple Difference			X			
Extension	X	X	X			
Full Dataset				X	X	X

Notes:
***p < 0.001
**p < 0.01
*p < 0.05

3.4.6 Robustness to Alternative Assumptions

Columns 5 and 6 of Table 3.8 report effects at each distance band when I relax the windfarm exclusion criteria to allow for windfarms sited in predominantly urban areas to be included in the analysis. This increases the number of windfarms in the analysis substantially and is expected to alter the results from the baseline due to issues of visibility estimation. Recall that we assume visibility will be the main determinant of price effects. This test is an important robustness check, as the inclusion of urban-sited windfarms increases the possibility of incorrect visibility estimation. The potential for miscoded visibility estimates arises from the fact that, urban areas tend to be more built up. These areas are more likely to contain buildings which may obscure a windfarm from view. However, the DEM utilized for the visibility calculation does not take buildings into account, so this will lead to an increased number of postcodes with misclassified visibility.

Table 3.10 reports a series of series of results organized by visibility and compared to the extended analysis (Columns 1 and 4). These include transaction exclusion criteria where the transactions below the 5th and above the 95th percentiles are dropped from the analysis (Columns 2 and 6); intensity of treatment where postcodes near only a single windfarm (Columns 3 and 7) are compared to postcodes multiple windfarms (Columns 4 and 8). I test for visibility intensity by segregating the dataset into two groups-where the first group contains postcodes within 14km of only one windfarm and the second group contains only postcodes within 14km of multiple windfarms (at least two).

Next, I test the analysis for sensitivity to the study period and present results in Table 3.11. Here I have split the full dataset into 4 discrete study periods to test if effects are consistent across the full dataset. The periods consist of three, six-year groupings and one five-year grouping: 1995- 2000; 2001-2006; 2007-2012; 2013-2018. These groupings were chosen to estimate the evolution of windfarm-visibility impacts over the full period of data. An advantage of the current division is that it is split into roughly equal time periods which capture effects at the pre-replication period, an early replication period, a late replication period, and a post replication period. This allows for a reporting of the evolution of price impacts at useful points. All visible effects are grouped to the left of the table, and all non-visible effects to the right.

Table 3.10: Extended Analysis: Robustness Checks

Distance Band	1	2	3	4	5	6	7	8
0-1km	0.0302	-0.0398***	0.113***	-0.115***	0.0549***	-0.0346	0.0670	-0.0432
(RSE)	(0.0296)	(0.0112)	(0.0262)	(0.0068)	(0.0183)	(0.0572)	(0.0349)	(0.0716)
1-2km	0.0329*	-0.0131***	0.0218	-0.0287**	0.0290**	-0.0050	0.0216	-0.0096
(RSE)	(0.0173)	(0.0046)	(0.0398)	(0.0141)	(0.0120)	(0.0140)	(0.0441)	(0.0239)
2-4km	0.0773***	0.00254	0.0492***	-0.0073	0.0251***	-0.0133***	-0.0087	0.0097
(RSE)	(0.0076)	(0.0025)	(0.0178)	(0.0073)	(0.0049)	(0.0051)	(0.0202)	(0.0063)
4-8km	0.0574***	0.0037***	0.0387***	0.0070	0.0062***	0.0063***	-0.0010	0.0213***
(RSE)	(0.0061)	(0.0014)	(0.0096)	(0.0064)	(0.00234)	(0.00201)	(0.0100)	(0.0027)
8-14km	0.0841***	0.0052***	0.0422***	0.0642*	-0.0189***	0.0028**	0.0763***	-0.0094***
(RSE)	(0.0045)	(0.0010)	(0.0059)	(0.0368)	(0.0015)	(0.0024)	(0.0070)	(0.0027)
Extension	X				X			
House Price Restrictions		X				X		
Single Windfarm Only			X				X	
Multiple Windfarm Only				X				X
Visible	X	X	X	X				
Not Visible					X	X	X	X

Notes:
$***p < 0.001$
$**p < 0.01$
$*p < 0.05$
RSE: Robust Standard Errors Clustered at the Output Census Area
Control Vars: postcode slope-by-year, elevation-by-year, aspect by-year dummies, proportions of sales of detached, semi-detached, terraced, flat/maisonette;quarterly dummies. Results are reported to four decimal places.

Table 3.11: Robustness to Study Period

Distance Band	Price Effect Of Presence of an Operational Windfarm									
	1	2	3	4	5	6	7	8	9	10
0-1km	0.0302	-0.0167	-0.0980****	-0.0079***	0.0191	0.0549***	-	0.0720***	0.0753	0.0643**
(RSE)	(0.0296)	(0.0477)	(0.00275)	(0.00170)	(0.0168)	(0.0183)	-	(0.00585)	(0.0665)	(0.0288)
1-2km	0.0329*	-0.0473*	-0.0870***	-0.0032	0.0063	0.0290**	0.2530	0.0723***	0.0799***	-0.0015
(RSE)	(0.0173)	(0.0283)	(0.00969)	(0.00717)	(0.00738)	(0.0120)	(0.322)	(0.0146)	(0.0288)	(0.0158)
2-4km	0.0773***	-0.0887***	-0.104***	0.0383***	-0.0091**	0.0251***	-0.0451*	-0.0119**	0.0151	0.0117*
(RSE)	(0.00759)	(0.0101)	(0.00475)	(0.00376)	(0.00363)	(0.00485)	(0.0266)	(0.00367)	(0.00115)	(0.00641)
4-8km	0.0574***	-0.0683***	-0.110***	0.0114***	0.0017	0.0062***	-0.0299***	-0.0260	0.0271***	0.0065
(RSE)	(0.00610)	(0.00520)	(0.00293)	(0.0024)	(0.00229)	(0.00234)	(0.00911)	(0.0238)	(0.00449)	(0.00351)
8-14km	0.0841***	-0.0026	-0.0233	-0.00722	-0.0080***	-0.0189***	-0.0202**	0.0109	-0.0011*	-0.0079***
(RSE)	(0.00453)	(0.00259)	(0.0181)	(0.0165)	(0.00199)	(0.00147)	(0.00823)	(0.0416)	(0.00245)	(0.00236)
Extension	X					X				
1995-2000		X					X			
2001-2006			X					X		
2007-2012				X					X	
2013-2018					X					X
Visible	X	X	X	X	X					
Not Visible						X	X	X	X	X

Notes:
$***p < 0.001$
$**p < 0.01$
$*p < 0.05$
RSE: Robust Standard Errors Clustered at the Output Census Area
Control Vars: postcode slope-by-year, elevation-by-year, aspect by-year dummies, proportions of sales of detached, semi-detached, terraced, flat/maisonette;quarterly dummies. Results are reported to four decimal places. There were not enough observations to generate results for Column 7 at the 0-1km DB.

The effects from visible windfarms are found to be largest and most negative in the period 2001-2006, where the statistically significant decrease in postcode average price ranges from -11% to -8.7%. The 1995-2000 period sees similar trends but range from -8.9% to -4.7%. Although these effects are smaller on average, during this period, windfarms generally consisted of smaller and fewer turbines. In the 2007-2012 period, there remains a strong negative impact from windfarm visibility at the 0-1km band, of 7.9%, but the more distant treatment bands become statistically significant and positive at the 2-4 and 4-8km bands. (3.8% and 1.1% respectively). The final period exhibits statistically significant treatment effects at only the 2-4km and 4-8km bands, but these are both less than one percent decreases in price. Generally speaking, the visual disamenity of windfarms appears to degrade over time. However, it is important to contextualize this degradation with non-visible treatment effects. Non-visible windfarm treatment is generally positive across the different study periods. These positive effects appear to be consistently significant and sizable at the closest distance bands, ranging from an 8% increase in 2007-2012. The positive effects decrease with time, and distance.

3.5 Discussion

The replicated results find negative impacts from windfarm visibility and proximity and though this is in agreement with what has been reported by Gibbons (2015), though the replicated results are considerably larger - in particular at the greater distance bands. At the same time, the impacts arising from non-visible windfarms is generally about 2% smaller than found by Gibbons. There is evidence that nearby, visible windfarms did decrease the average transaction price nearby postcodes over the study period of 2000-2011. This implies that home buyers were willing to accept the presence of an operational and visible windfarm near their home, but only when the price of that property is reduced by between 6.4 and 3.8% if it lies within 8km of such a windfarm. Conversely, these results imply that home buyers are not willing to pay a statistically significant higher price for properties in postcodes within 4km of a windfarm, or further than 8km of a windfarm. There is evidence of a willingness to pay a slight premium for lack of visibility at the 4-8km distance band. Overall, I find evidence that in England and Wales between 2000 and 2012, among home buyers there was a willingness to accept a discount for windfarm disability and only a small willingness to pay for properties in similar locales but without a line of sight to a windfarm. Therefore, I can conclude that under these assumptions, and during this period in time, windfarm visibility was an environmental disamenity.

When the expanded dataset is analyzed relaxing only the time period restriction, I find that the negative impacts arising from nearby windfarms disappear. Even at the closest distance, the effects are no longer negative - they are positive, though insignificant. Interestingly, the positive price effect mostly increases with distance from a windfarm and reaches a maximum of +8.41% at the 8-14km distance group. This does still follow Gibbons' results in the sense that the average price of properties increases with distance from a windfarm, though when including all transactions and windfarms, the results are increasingly positive rather than decreasingly negative. An examination of the changes in price effects over time has not yet been included within the literature, however there are two potential avenues that may explain this divergence in price effect between the early period of the replication and Gibbons (2015). Firstly, it is possible that home buyers are now willing to pay a premium for the view of a windfarm relative to the lack thereof. Secondly, and potentially more plausible is the issue of changes to the characteristics of windfarms themselves. While the installed turbine sizes has increased over time, the inclusion of recreational amenities within windfarm cites has also increased substantially since the Replication period of 2000-2011. Such amenities include payments to local community funds, mountain biking or hiking trails and other outdoor activity sites. It may be that windfarm visibility remains a disamenity, and preferences towards it remains negative as research by Parmeter & Pope (2013) suggests; but the included amenities over compensate for the decrease in value driven by visibility.

The results are sensitive to the inclusion of urban-sited windfarms. This is expected due to the increased ambiguity of visibility estimates for these urban windfarms, and as we expect that visibility is a key determinant of the price effect this sensitivity is reasonable. I can report that the results are sensitive to the inclusion of urban windfarms, which drastically change the results relative to the baseline estimate. This essentially validates the comparison made between treatment by proximity versus an interaction between proximity and visibility as argued by several of the key papers in the literature (Sunak & Madlener 2016, 2017, Heblich et al. 2016, Gibbons 2015). As a building height dataset for the entire land area of England and Wales was not available, this also further justifies the restriction of the baseline analysis to incorporate rural windfarms only. Lang et al. (2014) and Hoen et al. (2011) are the only papers in the literature to exclusively model effects from urban windfarms alone, though they do not find statistically significant impacts.

Next, I tested for consistency in price effects across time by subdividing the panel in to four subgroups. Overall, the treatment effects of both windfarm visibility and proximity appear to diminish over time. This evolution of windfarm visibility preferences explains the differences between the results over the replicated period and the full period of analysis. It seems that with time, individuals demand smaller and smaller compensation for windfarm visibility. At the same time, their willingness to pay a premium for properties treated by non-visible windfarms has decreased. There are a few factors which could explain this temporal variation in coefficient estimates.

This may reflect the fact that there are fewer and fewer properties on the market which are greater than 14km from a windfarm centroid due to the expansion of the UK's wind energy generation. This could imply that as there are fewer and fewer homes without a view of a windfarm, home buyers may place less and less importance on the presence of nearby windfarm visibility. Additionally, as proposed previously, there may be increasing amenities associated with windfarm sites which compensate for the reduction in natural views of the landscape. There could also be a general trend in preferences towards views of windfarms where individuals generally appreciate the sight of turbines. Under the assumptions of the extended analysis, windfarm visibility is associated with positive price effects at all but the 8-14km distance band. However, when the full dataset is analyzed, there is a statistically significant price impact from visibility at the 0-1km band. This is also reflected in the Triple Difference Analysis.

I have shown that results are, in a manner, sensitive to the period of the analysis performed, but more importantly that by splitting the long-term analysis into period subgroups, it is possible to track that sensitivity. This may show that changes in attitudes towards windfarm visibility which have shifted from largely negative, to now largely positive. However, it is also possible that the preferences towards windfarm visibility has remained stable, but the characteristics of windfarms themselves has changed and the inclusion of recreational amenities within windfarm developments overcompensates for the visual disamenity of the windfarm.

This analysis has thus provided a generalized context to the wider literature by potentially explaining the variation in reported results throughout. While papers such as those by Sunak & Madlener (2016, 2017) find stable price effects on attitudes towards three windfarms, I find that attitudes do indeed change, and become more positive. Though it is out with the scope of this Chapter to determine exactly what has caused these new attitudes.

To summarize, the study period analysis showed that the negative effects from windfarm visibility are largest in the earlier periods and become positive or insignificant by the 2013-2018 period. Effects from non-visible windfarms are largest and most positive in the 2001-2006 group but are more mixed than the effects from visibility. This supplies evidence that preferences towards windfarm visibility are not constant over time or that changes to windfarm characteristics have overcompensated for the visual disamenity of the windfarms. . To further test the robustness of these results, I altered the treatment definition to account for treatment intensity. Under this specification, I found that there is a considerably larger negative effect from visibility of multiple windfarms relative to postcodes with only a single visible windfarm.

In the Headline results, I define a postcode as treated when: one quarter after at least one visible or not-visible windfarm within a given distance band becomes operational. This leads to postcodes receiving treatments from different windfarms at different distance bands, which is arguably a very different to a treatment with only one windfarm at any distance band. One would expect the visual impacts arising from several windfarms to be different from the effects arising from a single windfarm. This is of particular importance in the extended analysis due to the significant increase in the number of operation windfarms. Additionally, the larger the number of nearby visible windfarms, the less likely it is that a postcode is grouped into the wrong treatment group (visible or non-visible treatment).

To test this, I split the panel in two groups-where the first group contains postcodes within 14km of only one windfarm and the second group contains only postcodes within 14km of multiple windfarms (at least two). This is essentially to test for differences arising from treatment intensities. I find that there is indeed a difference in the estimated coefficients when accounting for treatment intensity, following expectations, price effects are much more negative for visibility of multiple windfarms, whereas treatment by single visible windfarms are broadly positive. This is in line with expectations Heblich et al. (2016), Gibbons (2015), Sunak & Madlener (2016), Jensen et al. (2014). This is not unexpected, as it is assumed the main driver of any effect will be visibility and therefore, visibility of multiple windfarms increases the likelihood that a view from a postcode centroid will contain a larger share of wind turbines.

Lastly, I made a restriction to the house prices to remove outlier transactions. This was partly due to the fact that some transactions in the dataset were as low as £0, whilst others were as high as £780,000,000. This specification showed that the results are sensitive to restriction of very low and very high value transactions, though still makes use of the majority of the available data. The only other analysis which mentions the removal of outlier transactions is that by Sims & Dent (2007), who restrict their analysis to include transactions above £40,000 and below £400,000.

3.6 Conclusions

In the first portion of this Chapter I replicate the analysis of Gibbons (2015). The replication found substantially larger negative impacts arising from windfarm visibility than the estimates reported by Gibbons (2015). I find evidence to support the claim that visible windfarms did have a statistically significant negative impact on postcode-average property transactions in England and Wales between 2000 and 2012. In the subsequent section, I applied Gibbons' methodological approach to a much larger and more current dataset. This expanded the postcode-quarter transaction of the analysis from 1.8 million observations to nearly 7 million observations. Additionally, 505 more windfarms were included in the analysis, many of which are structurally different than those analyzed by Gibbons. This extended analysis was necessary due to the substantial increase in the number of operational windfarms in England and Wales, as well as the substantial increase in the number of housing transactions occurring within 14km of these windfarms.

The results of this extended analysis show firstly, that when simply applying Gibbons' analysis the results show a complete reversal of the effects found in the replicated analysis. Windfarm visibility is associated with a large statistically significant positive effect on house prices, with this being greatest at the largest distances from the windfarm. I find that the lack of visibility is no longer associated with a consistently positive effect on house prices at the closest proximities. However, I also find that the analysis is highly sensitive to a number of assumptions regarding the underlying data applied within the model. These include sensitivity to the inclusion of all windfarms, both urban and rural sited; house price restriction to exclude the least and most valuable 5% of transactions, the intensity of the view of nearby windfarms, and the study period itself.

The findings of this chapter have several important implications for policymakers. Firstly, there may be justification to compensate homeowners whose properties decreased in value - though this may be difficult as the results imply this was the case for transactions the farthest in the past. Going forward, it may be advisable to account for the fact that postcodes near to multiple windfarms do experience larger and more negative price impacts. This should not necessarily discourage approval of multiple windfarms near to each other, rather clustering should be encouraged where there are relatively few nearby homes. In addition, as there is now evidence that windfarm visibility positively impacts home values there may be scope to further increase the number of windfarms in England and Wales. However, there should be additional research examining whether the positive effects found in the extended analysis arise from windfarm visibility or rather access to amenities which are constructed at windfarm sites. Policymakers should share this information with communities near potential windfarm locations to build greater support for such developments, as well as ensuring that new or existing developments include recreational amenities. This would have the benefit of both increasing electricity generation from a low-carbon source, but may also increase nearby house prices allowing for home owners to reap both environmental and financial windfalls.

There are limitations of the analysis presented within this chapter.Firstly, the Average Price approach to Hedonic Valuation assumes that the properties within a postcode are very similar, that the average transaction price in a postcode is representative of the average value of properties within that postcode. Therefore, any changes or differences in price that are observed can be explained via the presence of a nearby windfarm. However, this may not necessarily be the case.

Many postcode-quarter averages are generated via a single transaction. This could lead to averages below or above the truly average transaction price leading to biases in the estimation, as the composition of average transactions will vary over time. Additionally, by restricting to windfarms which are not sited in urban areas, I increase the likelihood that the properties in a postcode are less similar – consider a block of flats or terraced homes, at least in regard to their layout and size are very similar. Whereas in more rural areas or postcodes with a larger share of detached homes, there will be greater variation in the housing stock. My analysis does include the property type, but has no information on more detailed property characteristics which could be utilized to ensure the housing stock in postcodes are as similar as possible.

A potential weakness of the approach taken in this analysis is the definitions of both windfarm proximity and visibility. In the analysis of this chapter I tested for sensitivity to alternative means of visibility estimation, achieved in two ways: dropping postcodes which have a high visibility gradient from the 200m or 90m grid cells surrounding the postcode, or testing for changes in visibility estimates by increasing the viewpoint (postcode centroid) by 5 additional meters. Postcode visibility estimates were relatively stable across these alternative methods. However, the use of postcode and windfarm centroids as visibility points is crude – though useful.

This ultimately leads to an estimation of the effects on house prices in neighborhoods where windfarms are likely visible. This could potentially bias estimates if preferences about windfarm visibility from ones' own property is what drives behavior rather than simply living in an area where they may be visible. Though this has been little discussed within the windfarm amenity literature, there is some evidence from other visual amenity analyses that there are differences in value placed on properties in neighborhoods with good views, compared to properties with good views (Bourassa et al. 2004).

Further, the application of the Fixed Effects DID to the panel assumes that by controlling generally for geographical effects as well as time effects, the omitted variable bias is avoided. However, it this may not be the case for such a large national sample as used within this chapter. There may be local economic effects which the model cannot capture such as changes in employment or connectivity to larger job markets. Unfortunately it is not feasible to include all variables which may control for all possible biases - within the wider hedonic literature, there is even evidence of micro markets within individual cities, therefore this will be a limitation of any national scale analysis which applies fixed effects.

As the analysis makes use of a Hedonic Valuation approach which is built on the seminal work by Rosen (1974) it is worth noting the limitation inherent to this approach. Firstly, there is a risk of misspecification of the hedonic price schedule (HPS) which would lead to unreliable estimation of the marginal price impact of windfarm visibility on nearby house prices. However, this is less likely to be an issue within quasi-experimental contexts, as well as in cases where changes in characteristics (windfarm visibility) occur in the past Greenstone (2017). Both of these are key features of the research design used here. To fully measure non-marginal price impacts would require the application of a sorting model, which was beyond the scope of this research.

Finally, there is the potential for selection bias influencing the results. As the price effects become more positive and larger with time, there is potential that these are driven by the fact that as the number of sites where windfarms become operational increases the characteristics of homes and buyers may change. Because of this, it is unclear whether the changes between the early and late periods are due to changes in preferences, characteristics of windfarms, or simply because new windfarms are being located in areas where house prices are already increasing. This is an important limitation of the analysis, and is something that future research should examine in more depth.

To account for the limitations of the analysis presented in this chapter, I continue the analysis in the next chapter, but apply a Repeat Sales specification of the Hedonic Pricing Method. In this analysis, individual properties located within 14km of operational windfarms are analyzed rather than postcode-averages. This means that the analysis compares the price of property i in quarter t to property i in quarter t+x. This analysis allows for further testing of sensitivity to the underlying data under an alternative, but complementary analytical framework. I am able to exploit the characteristics of a repeat sales analysis to provide additional insight to the results within this chapter.

I address the potentially flawed assumption that all transactions within a given postcode are representative of the other properties within that postcode by comparing properties to themselves over time. Additionally, I am able to refine visibility estimates from neighborhood visibility to property visibility. This is a potentially important distinction if buyers value the views from their homes more than the views from their neighborhoods. I then make use of additional data such as building height information and individual turbine locations to better model visibility. The issue of comparability of properties within a given postcode no longer becomes relevant because properties are being compared to themselves. In addition, I am able to account for the intrinsic characteristics of the properties themselves, in as well as the property type.

Chapter 4

House price effects of windfarm siting in England and Wales: A repeat sales analysis

4.1 Introduction

Over recent years, there has been a substantial shift in attitudes towards the use of carbon emitting electricity generation. This is largely due to ever growing concerns regarding the economic and environmental impacts arising from substantial increases to the global average temperature and the dangers that this poses. As such, there has been wide-scale and rapid deployment of renewable energy developments in many industrialized countries. The UK has been one of the leading adopters of renewable, low-carbon energy generation, and in fact its domestic carbon emissions peaked in and have reduced since 1972, or in 2007 if imported emissions are included ONS (n.d.). This reduction has arisen from several sources, but between 1990 and 2017, the carbon intensity of electricity generation within the UK has fallen by 67% ONS (n.d.).

This reduction in the carbon intensity of electricity generation has largely occurred due to the decommissioning of coal fired plants which have been replaced with renewable generation. Wind energy generation in the UK has developed rapidly leading to both praise for the amount of carbon emissions not generated due to these installations, as well as concerns about externalities that such a wide-scale and rapid deployment of new energy infrastructures may entail. One of the most common concerns voiced by the public are the effects of windfarm siting on nearby property values as other energy generation infrastructure has been shown to have detrimental effects for properties sited nearby.

This has led to a growing academic interest in the area, with several papers being published between 2005 and 2019 which attempt to determine if there is a link between property values and windfarm siting. Of these papers, only a select few have explored the price effects of windfarm siting within the UK, the two largest and most comprehensive being Gibbons (2015) and Heblich et al. (2016). These papers made use of property transactions which occurred up to 2012 and 2014 respectively. Because the British public is largely supportive of wind energy, as well as the siting of large-scale renewable infrastructure 'in their area' it is possible to compare stated and revealed preferences towards wind energy developments within England and Wales.

Gibbons (2015) was replicated and extended in the previous chapter due to the substantial change in the landscape since the time of publication. Indeed, with a continued rapid development of windfarms more and more areas of the country are located nearby wind farms and the turbines installed in the latest windfarms are considerably larger than older turbines within the replication period. The analysis of the previous chapter applies an average sales hedonic model to estimate price impacts of windfarm proximity and visibility, with observations at the quarterly postcode level.

In this chapter, I build on the work in Chapter 3 in several ways. First, I perform a repeat sales analysis to determine if the effects found on postcode-average prices are similar for the repeat sales analysis applying a property-windfarm centroid visibility analysis. This is potentially an important comparison for two reasons: By comparing postcode level to individual property level effects, I am able to decompose visibility into neighborhood level and property level visibility while also testing whether the estimated effects are consistent across an alternative analytical approach. Here, I compare effects from windfarm centroid visibility and non-visibility at several different distance bands.

As a further refinement of the approach taken in the previous chapter, I then enhance the analysis to include property-turbine visibility analysis rather than estimating visibility from the center of the windfarm. This does restrict the analysis somewhat due to lack of turbine locations for some windfarms, but it is still a useful exercise. I also perform a further refinement of the visibility estimation on a subset of the data where building height data can be included in visibility estimation, which accounts for intervening buildings which block views of wind turbines. Because the analysis occurs at the property level, I make use of detailed property characteristics data to control for differing attributes of the properties in the analysis. While the analysis of the previous chapter controls for property type, here I perform the analysis on a subset of the transactions which were paired with detailed property characteristics information.

The empirical analyses performed in both this chapter and Chapter 3 is underpinned by the assumption that the visibility of windfarms/turbines the key treatment that will affect property values, and that proximity is essentially a measure of how prominent the windfarm/turbines are within a viewshed. As such, I model visibility in several ways, making the trade-off that for the most accurate visibility estimates requires restricting the analysis to a very small subset of the data. However, even when the most accurate visibility estimates are not available, I control for fixed effects and apply a staggered difference-in-difference analysis, where the treated group of properties are those which sell after a nearby windfarm has become operational and the control properties are those selling at the same time, but where nearby turbines have not become operational yet. This analysis is then followed by a difference-in-difference-in-difference analysis comparing effects between properties treated by visible and non-visible wind turbines.

4.2 Contextual Background: A Changing Landscape

The United Kingdom has made substantial progress in reducing the emissions of its electricity generation, decreasing the carbon intensity of the electricity it generates by nearly 70% between 1990 and 2017 (ONS n.d.). One of the key drivers of this reduction has been a rapid expansion of renewable energy technology within the country which has gradually replaced decommissioned coal fired plants. The share of electricity generated from coal within the UK has seen a steady decline whilst the contribution of renewable energy has seen a sustained increase over the period 1998-2020. In fact, as of the first quarter of 2020, fossil fuels supplied just under 30% of UK electricity and the largest single source of electricity was generated from wind turbines (27.585%) (OFGEM 2020*b*). The rapid advancement of the share of total electricity by wind has been fueled by the rapid deployment of wind turbines within the UK. In this chapter I will only discuss these advancements within England and Wales as the other UK regions are not included in the Repeat Sales Analysis performed below. Figure 4.1 highlights how the stock of UK wind energy generation has evolved over time, displaying cumulative installed capacity, number of windfarms, and number of wind turbines respectively. Figure 4.2 displays the stock on and offshore turbines over the same period and highlights the more recent and perhaps more rapid growth in size of offshore wind generation. As I will show in the results section of this chapter, the changes in windfarm/turbine stock over time is an important driver of results, as preferences towards windfarms is not stable over time, as well as the number of homes lacking visible windfarms decreases as windfarms become more widespread.

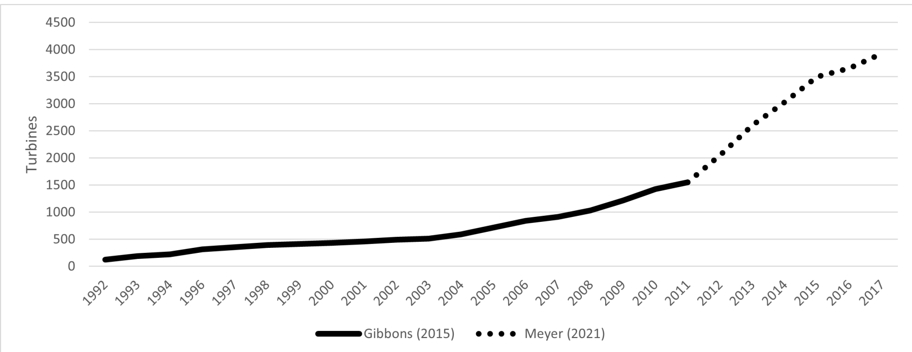

Figure 4.1: Operational Wind Turbines in England and Wales: 1992-2020

The figures above highlight the rapid deployment of wind energy de-

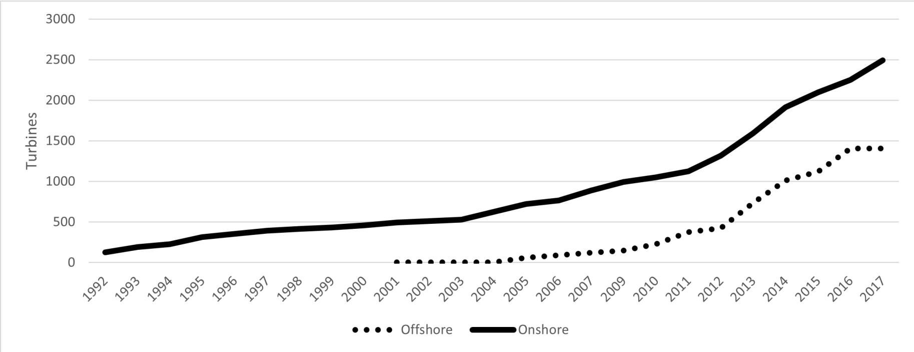

Figure 4.2: On and Offshore Turbines in England and Wales: 1992-2020

velopments within England and Wales. As can be seen, not only has the UK's stock of wind energy generation greatly increased since the first windfarm became active in 1992, but the current stock has greatly changed since the time period that the paper by Gibbons (2015) explored the effects of windfarm siting – namely the interaction of windfarm proximity and visibility – on house prices. The results presented in the previous chapter show that this change in preferences towards windfarm proximity and visibility over time substantially affect the price effect and this is also found for the repeat analysis of this chapter.

These changes in preferences towards windfarms are assumed to be the key driver of any price difference arising between properties with and without a visible windfarm nearby. As reported in the Department for Business, Energy and Industrial Strategy Public Attitudes Tracker, there has been increasing support both for wind energy generation in the abstract, as well as support for large scale renewable energy infrastructure 'in their area.' Of course, for a true comparison between the stated preferences of the wider public and their preferences as revealed through hedonic pricing, it would be necessary to survey actual home buyers, which is beyond the scope of this thesis. However, changes in surveyed preferences coupled with changes in revealed preferences taken together is a strong indicator of the important role time plays in the analysis of this chapter.

This shift in attitudes is also borne out in the BEIS surveys on energy and climate change BEIS (2020*a*). The surveys cover the period 2012-2020, though unfortunately not all waves of the survey contain the questions presented in Figures 4.3 and 4.4. Figure 4.3 displays the percent of respondents who agree that they "Would be happy to have a large-scale renewable energy development in my area" over the period 2012-2018. The responses show that though there is a majority of respondents who agree, though this covers all renewable developments and does not define the size of the 'area.' Taken together with survey questions about preferences towards wind energy, it would appear that there is support both for wind energy as well as other renewable energy developments local to the respondents.

Figure 4.4 presents the survey results for On and Offshore Wind Energy: "Q13. Generally speaking, do you support or oppose the use of the following renewable energy developments." Between 2012 and 2020, there has been an increase in the percent of respondents who claim to support offshore wind developments, rising from 76% in 2012 to 81% in 2020, with peak support of 83% in 2018 and 2019. Broadly speaking, the support for offshore wind has broadly been consistently high over the period of the survey. Onshore wind has seen a much larger gain in support, rising from only 66% of respondents in 2012 to 78% in 2020, after peaking at 79% in 2019. This increase in support has tracked with the increased presence of onshore wind over the period.

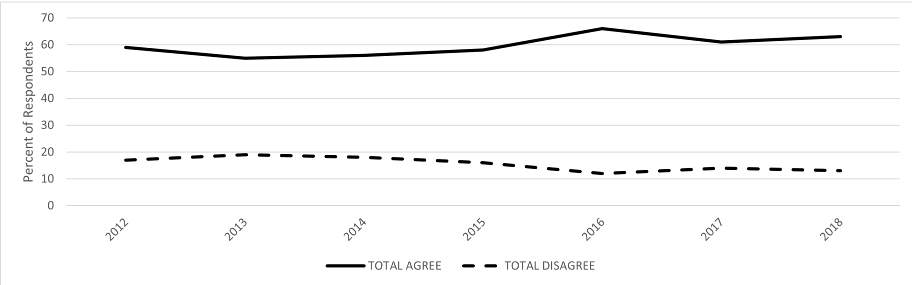

Figure 4.3: Agreement with the statement “I would be happy to have a large-scale renewable energy development in my area"

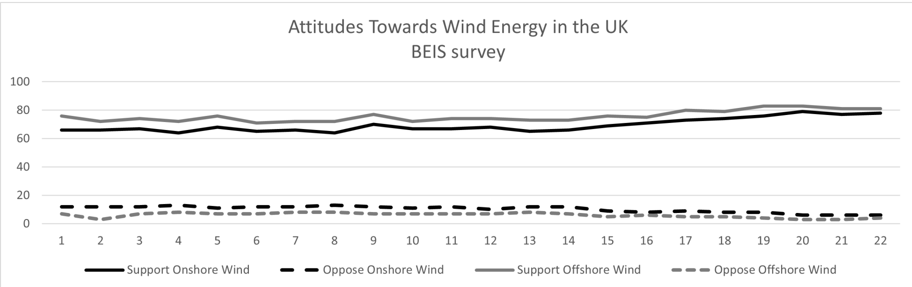

Figure 4.4: “Generally speaking, do you support or oppose the use of the following renewable energy developments?”

4.3 Relevant Literature

In this section, I will present the literature relevant to this chapter and the empirical analysis performed within it organized into three general categories. Firstly, I will discuss the concepts of environmental amenities and disamenities, secondly how these can be valued via the Hedonic Pricing Method, and lastly I will discuss how other works within the literature have applied these within the context of house price effects arising from windfarm proximity-visibility. There is some duplication of the discussions of Chapters 2 and 3, though here I focus on the application of repeat sales analyses within the literature.

4.3.1 Amenities and Disamenities

The purpose of the empirical analysis in this chapter is to apply a repeat sales analysis which supports the classification of wind energy developments as either environmental amenities or disamenities. There is not a standard definition of either environmental amenities nor environmental disamenities (Schaeffer & Dissart 2018) and in fact the exact definition tends to differ from paper to paper. However, in general terms, an environmental amenity is some environmental feature which provides a benefit. This means that if there is a measurable benefit arising from the presence or loss of a feature or characteristic of nature, this is an environmental amenity. Conversely if there is a measurable economic cost arising from the presence or loss of a feature of nature, this is an environmental disamenity.

In the context of this and the previous chapter, I seek to determine if the visibility of a windfarm or wind turbine is an environmental amenity or disamenity. To do so, this requires that there be a measurable economic cost or economic benefit arising due to the presence of a visible turbine or turbines. Or, if we assume that natural vistas have an economic benefit which can be measured, this can be rephrased to the research question, 'does windfarm or wind turbine presence affect the economic benefit derived from natural vistas?'. If the answer is yes, and they increase the economic benefits, windfarms and wind turbines can be classified as an environmental amenity. If the answer is yes, and they decrease the economic benefits of natural vistas, windfarms and wind turbines can be classified as environmental disamenities.

4.3.2 Hedonic Valuation of Environmental Features

To determine whether visual impacts of windfarms are environmental amenities or disamenities, it is necessary that the economic costs or benefits associated with them be measurable. This presents a problem, as there is no explicit market for natural vistas with or without wind turbines, from which a value can be quantified. It is, therefore, why many studies which attempt to quantify the economic value of environmental amenities or disamenities apply some form of Hedonic Valuation. Hedonic Valuation or Hedonic Pricing is the application of a model which assumes that the price of a good is determined by both internal and external features or characteristic which affect that good. Because the price or value of the good is determined by the combination of its characteristics, if these characteristics are disaggregated, it becomes possible to estimate the values of each characteristic (Blumenschein et al. 2008, Lang et al. 2014, Monson 2009).

Following as defined in Chapter 4 the following is the generic form of the hedonic model applied:

$$Price = f(intrinsic features, extrinsic features, time)$$

Intrinsic features refer to characteristics such as the number of bedrooms, bathrooms, lot size, etc. Extrinsic features are characteristics of the location where a property is situated, such as the crime rate, school quality, or as the focus of this and the previous chapter – the presence of nearby wind turbines. When a hedonic analysis is performed for a large sample of properties, it becomes possible to estimate the average value of each component of properties in the analysis. Hedonic Valuation has been applied to classify many environmental features as either amenities or disamenities based on how they affect the value of homes.

There are two types of Hedonic Analyses that have been widely applied in the environmental amenity literature – Average Price analyses and Repeat Sales analyses. The key difference is that an average price approach simply compares the average transaction price of a geographical area – most commonly a postcode or zip code – before and after treatment. This was performed in the previous chapter. A repeat sales analysis compares the same property at transactions which occur before and after treatment - in the case of this chapter, when nearby wind turbines become operational.

4.3.3 Valuing the Visual Impacts of Windfarms

There is a developed literature around the classification of windfarms and wind turbines as environmental amenities, but here I will only address those most relevant to the analysis of this chapter. Most literature discussed here have applied a repeat sales analysis, for a wider review of the full literature see Chapter 2. I have structured this section around the key assumptions which underpin the analyses within the literature and within this chapter. The first key assumption is that the magnitude of a price effect from an environmental feature will be correlated with distance – i.e., a home located closer to an amenity will receive a larger price premium than a home located further away. Conversely, a home located closer to a disamenity will receive a higher price discount than a home located further away.

The second key assumption is that visibility, paired with distance is a key determinant of any price effect arising from wind turbines. This implies that a property with a line of sight to a wind turbine will experience a price effect differing from a property without a line of sight to the wind turbine. It is therefore crucial to accurately estimate whether a property has a line of sight to any turbines. The third key assumption is that preferences towards windfarm proximity-visibility will change over time. Lastly, how the control group is defined will directly affect the estimated effects of any analysis – as the comparison between the treated and control group is what allows for the estimation of any price effects. Therefore, it is key to ensure that treatment and control groups are suitable for such an analysis.

Proximity

Because proximity is of key importance, every paper exploring the price effects of windfarm siting utilize some measure of distance as a key variable in the analyses that they perform. There are two papers which do not make use of visibility estimation in their analyses. Firstly is the paper by Sims et al. (2008) which applies a cross-sectional analysis of 900 properties in Cornwall but find no statistically significant effects.Heintzelman & Tuttle (2012) who are the first to apply a repeat sales analysis in the literature use only proximity to estimate the price effect.

They analyze the effects using transaction data on 11,369 properties over 9 years in Northern New York, USA and find statistically significant negative effect ranging from -7.73% to -14.87%. This is one of the larger price effects in the literature. The main specification of Dröes & Koster (2016) uses proximity alone as the treatment effect, as the authors assume that proximity to turbines is an adequate proxy for visibility, and that because the Netherlands is a very flat country it is likely that nearby wind turbines will be visible. They find a small negative impact of about 1% decreased prices within 2km of wind turbines relative to properties located farther than 2km away from the turbines of the study. They do perform a visibility analysis with a subset of their data and find no impact from direct views.

All other literature which estimates price effects from windfarms make use of some type of visibility analysis, and this is always paired with a measure of distance. This is due to the fact that despite wind turbines growing to ever increasing heights and being fitted with increasingly large blades, even large objects become smaller when viewed from large distances. In all other research, therefore, the treatment effect from windfarm siting is an interaction of both proximity and visibility together.

In all studies aside from those mentioned above, visibility estimates are combined with proximity estimates to define treatment of properties in the literature. Lang et al. (2014), Vyn & McCullough (2014), Sunak & Madlener (2016), Heblich et al. (2016), Dröes & Koster (2016) all make use of distance bands as either the main measure of distance, or as an alternative to additional analyses which use exact distance measures. Distance bands are essentially a means to group properties, or postcodes into discrete treatment groups. An example of the distance bands used in the analysis of this chapter overlaid on land elevation and building height information is shown in Figure 4.5. Highlighted in yellow is the Dagenham single turbine windfarm, the green triangles represent other wind turbines in the area. Each ring represents one of the distance bands used to group properties within the analysis of this chapter.

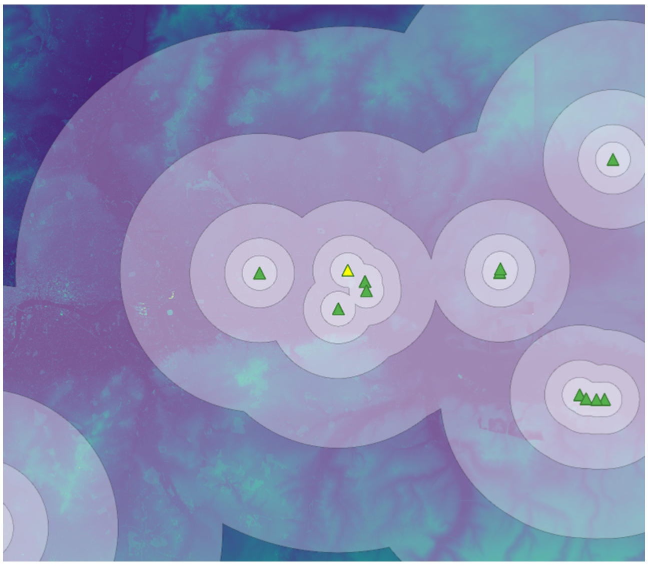

Figure 4.5: Example of Distance Bands Surrounding Wind Turbines

Visibility

Because visibility, paired with distance, is assumed to be the key driver of price effects arising from windfarm siting, the way visibility is defined is crucial to the results. Generally, this is done by using GIS software to estimate whether properties or postcodes have a line of sight, or view of a windfarm or turbine. This is achieved by plotting observers and windfarms or turbines by their geographic location, overlaid on some form of digital elevation or surface map into GIS software to calculate if there are obstructions or a clear view between the two points.

This is broadly the approach taken by the vast majority of the literature. Vyn & McCullough (2014), Sunak & Madlener (2016), Heblich et al. (2016), Dröes & Koster (2016) all use this approach. Heblich et al. (2016) make use of a Digital Elevation Map (DEM) comprised of 5m grid squares to account for topographic features which block views of windfarms, and as a further refinement to their analysis, add building height data to account for intervening buildings for the subset of their transactions occurring in locations where this data is available.

Sunak & Madlener (2016) follow a similar approach, but make use of a Digital Surface Map (DSM) at a resolution of 1m grids. This accounts for topographic features, building heights, and even vegetation. However, it should be noted that neither paper restrict their analysis to a period based on the publication date of the building height or DSM data used. This may reduce the reliability of these highly refined visibility estimates simply due to the fact that the landscape (particularly vegetation and structures) change frequently and the study periods cover 24 and 26 years for Heblich et al. (2016) and Sunak & Madlener (2016) respectively.

Two studies performed in-person field visits to inspect properties and generate visibility and intensity categories. Lang et al. (2014) visited 1,354 properties which were located within 2 miles of urban windfarms in Rhode Island. While Hoen et al. (2011) visited 6,194 properties across 5 US states to generate visibility measures, as well as to categorize the quality of vistas surrounding the properties. They then created an estimate of the degree of dominance the windfarm had on the landscape, which fell into five categories: no view, minor, moderate, high, and extreme. McCarthy & Balli (2014) made use of both GIS analysis and field visits, though they visited only properties which were classified as having visible turbines. Comparing between their estimated and actual visibility classification, there was roughly 5% of cases where GIS and field visitors disagreed.

Site visits were not feasible for the empirical analysis performed in this chapter due the sheer number of properties included. As such, I follow the approach taken by the other literature and model visibility using

GIS software, a digital elevation map and the locations of residential properties, windfarms, and turbines, and for a subset of the data include the locations and heights of intervening buildings to refine the visibility analysis of this chapter.

Difference-in-differences Estimation

The empirical analysis performed in this chapter applies a Spatial Fixed Effects Difference-in-difference (DID) model to estimate the price effects arising from windfarm proximity and visibility. This is the approach taken by Lang et al. (2014), Gibbons (2015), Heblich et al. (2016), Sunak & Madlener (2016), Dröes & Koster (2016). A more detailed discussion of the application of a DID model and controlling for spatial and time effects can be found in the previous chapter of this thesis. Here I will discuss how the treatment and control groups have been defined by the relevant papers, and how these compare to the definitions of treatment and control within this chapter.

Treatment of a property has been defined as the presence of a nearby operational windfarm or wind turbine, though I do estimate effects arising from visible and non-visible windfarms and turbines separately. However, equally important is the definition of the control, or comparator group from which the change in price resulting from treatment is estimated. The literature has broadly taken three approaches for defining the control groups. Two of these use somewhat arbitrary means of defining the control group. The first and most common is the creation of a control group for comparison using a set of properties which are located near windfarms, but where these windfarms or turbines are not visible. Vyn & McCullough (2014) compare properties at the same distance to the treatment group, but where the windfarms are not visible.

Sunak & Madlener (2016) generate their control groups using properties which have not yet been treated by visible windfarms, as well as those without a view of operational windfarms. Hoen et al. (2011), Lang et al. (2014), and Heintzelman & Tuttle (2012) create their control group using properties which are farther from the wind turbines than the treated group. For Hoen et al. (2011) These are properties 4-6miles from wind turbines, for Lan et al. (2020) the control properties are between 3 and 5 miles from the turbines, and for Heintzelman & Tuttle (2012) control properties are those more than 1/2, but within 10miles from an operational turbine.

Heblich et al. (2016) take a different approach. They explore price effects from both windfarm visibility and lack of visibility by comparing these properties to a control group which closely resemble the property characteristics of dwellings in the treated group, but which are not near a windfarm. Their use of detailed property characteristics enhances the suitability of the control group for comparison with the treated group, however there is still some scope for bias. The list of property characteristics is not exhaustive, so there may be some unobserved property features which differ between treatment and control groups.

For these reasons, this empirical analysis follows the DID framework as used by Gibbons (2015), Dröes & Koster (2016) and that of Chapter 3. Here, the treatment and control groups are one in the same – the key difference is that the control group will eventually be treated by windfarm siting. See Figure 4.6. Therefore, the two groups are equally suitable for wind developments, and because this is a repeat sales analysis, I compare the property values both before and after treatment by a windfarm. Therefore, the only real difference between the treatment and control groups is that while a property is in the control group, it

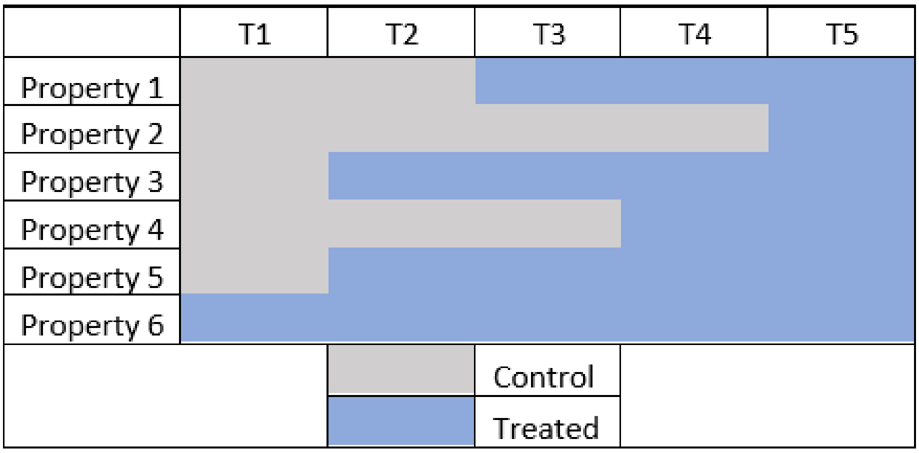

Figure 4.6: Treatment Visualized

simply has not been treated yet or at the same distance band. The analysis presented subsequently presents a difference-in-difference-in-difference (DIDID) estimation. Here, the average treatment effect from visible windfarms is compared to the average treatment effect from non-visible windfarms. This therefore will provide a more reliable estimate of the specific price effect of windfarm visibility in terms of WTP.

4.4 Estimation Strategy

The research design applied within this chapter follows from that applied in Chapter 3 and includes both fixed effects and regression-based difference-in-difference methods. Here I apply the model to individual properties as opposed to postcode averages. At all times, I am comparing pre and post treatment effects arising from windfarm visibility and lack of visibility. This means that I am estimating the difference in house prices arising from nearby and visible windfarms by comparing the price of homes prior to windfarm operation to prices of the same homes post-operation. I also perform the analysis comparing the pre and post treatment effects for properties sold before a non-visible windfarm is operations to the observed prices of those same properties after the windfarm has become operational. Ultimately the two sets of results are compared through a difference-in-difference-in-difference model.

The underpinning logic for comparing pre and post treatment prices is that the properties in the analysis will be suitable for comparing as 1) this removes the concern that observed price effects may be the result of differences in housing characteristics as the properties are compared to themselves and 2) regional effects will be similar as all properties in the control group eventually become treated. Equation 2.1 is the baseline analysis, which models price effects of proximity to windfarms alone.

Here, the treatment groups are those property transactions occurring after a windfarm has become operational at a given distance band – regardless of whether that windfarm is visible or not. Whereas the control group is the set of properties for which transactions occurred prior to a windfarm becoming operational but will eventually be located within 14km of an operational windfarm. This equation is subsequently augmented to include a visibility indicator as well as a vector of property specific characteristics, and it takes the same form as 3.1.

$$ln(price)_{it} = \sum_{k} \beta_k j_k < distance < k, operational)_{it-1} + x'_{it}\gamma + f(i,t) + E_{it} \tag{4.1}$$

Where:

- $ln(price)_{it}$ is the transaction price of property i in quarter t.
- $(j_k < dististance < k, operational)_{it-1}$ is an interaction dummy indicator which captures exposure to windfarm developments. With

a value of 1, this indicates that property i has at least one operational windfarm between j_k and k kilometers in the preceding quarter. This is an interaction of two dummy variables:

$j_k < dististance < k$ indicates whether the turbine(s) are within a given distance from a property.

$operational$ is a post-policy indicator which indicates whether the turbine(s) have been built and are operational.

- β_k, the coefficient of interest is the average effect of operational windfarm turbines within distance band $j_k - k$ on property prices. β_k captures home buyer's preferences for proximity to windfarms. Factors influencing the coefficient will include: noise or visual pollution, community grants, employment or other impacts related to turbine proximity. If β_k is positive, this implies that the mean transaction price of treated properties have risen faster than that of the control properties. If it is negative, the opposite is true.

- $f(i, t)$ represents unobserved characteristics which may vary over time and space, and are likely correlated with the visibility of windfarms. This potential correlation may arise because windfarms are not distributed randomly. Correlation with the time effects is present because the number of windfarms increases over time, and this would create a spurious correlation between any trend in prices over time with proximity to windfarms.

- E_{it} is the general error term.

As in Chapter 3 I have tested for anticipation effects, but using the sample of repeat sales properties I find no statistically significant effect. Property fixed effects are removed in 4.1 by using the within groups estimator and common time effects are removed by using quarter-specific dummy variables. The time effects account for external factors that are time variant, i.e. the financial crisis of 2008 which caused a significant drop in property values.

This restriction of properties creates a property groups which are similar in the following ways: they are close to sites which are amenable to windfarm developments, planning and construction have been completed, and it is likely that windfarms will be visible from the postcode. Further, properties are being compared to themselves which should ensure that property specific fixed-effects are accounted for. In one specification I also include a set of property specific intrinsic characteristics.

It is important to note that within this application of the HPM, it is only possible to measure marginal changes in price without the application of a further sorting model which is beyond the scope of this work.

4.4.1 Visibility

I then refine Equation 4.1 to include a visibility indicator, as well as property specific characteristics. I estimate treatment effects for properties treated by visible and non-visible wind turbines separately. Equation 4.2 is the model which generates the headline results for this chapter. Comparisons of the proportion of properties with views of windfarms, and the differences between visibility estimation at the postcode vs property level are presented in Appendix A8 and A9.

$$\begin{aligned} ln(price)_{it} = \sum_{k} \beta_k (visible, j_k < dististance < k, operational)_{it-1} \\ + x'_{it}\gamma + f(i,t) + E_{it} \end{aligned} \tag{4.2}$$

Where:

- $(visible, j_k < dististance < k, operational)_{it-1}$ is an interaction dummy indicator which captures exposure to windfarm developments. With a value of 1, this indicates that property i has at least one visible - or non visible -operational windfarm between j_k and k kilometers in the preceding quarter.

  $visible$ is the visibility indicator. When effects from visible windfarms are estimated this takes the value 1 if a windfarm is visible from property i. When effects from non-visible windfarms are estimated, this takes the value of 1 if a windfarm is not visible from property i.

- $x'_{it}\gamma$ is a vector of property specific characteristics - ie, number of rooms, floor area, etc.

The refinement of the model from Equation 4.1 to 4.2 will account for price effects arising from visible and non-visible windfarms, as well as allowing for control of property specific features which may affect price. This should provide a cleaner estimate of the price effects arising from windfarm proximity and visibility. Lastly, I apply a triple difference to estimate the true price premium or discount that home buyers demand as a result of the presence of a visible windfarm at a given distance. This is modeled in Equation 4.3:

$$ln(price)_{it} = \sum_k \beta_k (visible, j_k < dististance < k, operational)_{it-1} + \sum_k \delta_k (non-visible, j_k < dististance < k, operational)_{it-1} + x'_{it}\gamma + f(i,t) + E_{it} \tag{4.3}$$

Where:

- $(non-visible, j_k < dististance < k, operational)_{it-1}$ is the interaction dummy for lack of visibility, distance band, and operation of a windfarm.

- δ_k reports the average treatment effect arising from proximity to operational, but not-visible windfarms.

Here I compare the average price effect from windfarm visibility on post operation transaction to the average price effect from windfarm non-visibility on post operation transactions. If the resulting DIDID is negative, this essentially estimates the WTA for windfarm visibility.

4.5 Data, Descriptive Figures, and Statistics

Here I present the current state of wind energy generation within England and Wales, the housing market, and highlight the key features of the data relating to the analyses performed herein.

4.5.1 Windfarms and Wind Turbines

Figure 4.1 shows the significant and rapid changes in the wind turbine stock of England and Wales between 1992 and 2020, as well as the increasingly rapid growth after the analysis of Gibbons (2015). Figure 4.7 shows the turbine locations for all of England and Wales as well as the turbines included within this study. Not all operational windfarms or turbines are included in the analysis, as there are several offshore windfarms entirely or partially more than 14km from land, and therefore will are more than 14km from any residential properties. Table 4.1provides summary statistics on the size and capacity of the turbines included in the analysis. Note that there is a substantial range in both heights of turbines as well as their generation capacity. The smallest turbines, as well as those with the lowest generation capacity tend to be the earliest installed, and the largest tend to be more recent developments, with the largest being offshore developments.

Table 4.1: Wind Turbine Summary Statistics

	Mean	SD	Min	Max
Height to Tip	169.58	53.73	34	305
Rotor Diameter	94	36.12	-	180
Turbine Capacity	3.55	14.38	0.22	630

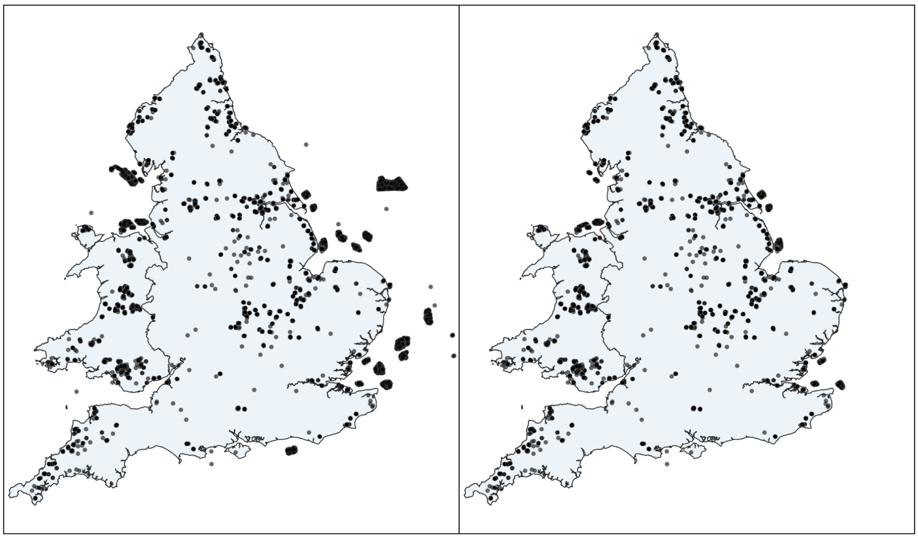

Figure 4.7: Comparison of all English and Welsh turbine locations (Left) to those within the study (Right).

4.5.2 Properties

The next set of figures highlight the locations of the properties used in the Analysis. I restrict this dataset in several ways pertaining to the analyses undertaken here. Firstly, only properties which were sold repeatedly over the period 1995-2018 are included, secondly, repeat transactions must occur within 14km of an operational wind turbine. This reduces the number of properties in the dataset to 2,299,352 – which have on average been sold 2.05 times over the period. The geographic distribution of these properties is displayed in Figure 4.8a.

One final geography-based restriction is for the analysis which applies the visibility estimation using information on the locations of and heights of buildings to account for visual obstructions due to man-made objects. These properties are shown in Figure 4.8b. Due to the nature of this data, there are both geographic and time-based restrictions due to the accuracy of the data becoming questionable 3 years prior and after its publication. This reduces the properties to 1,481,338 in total based

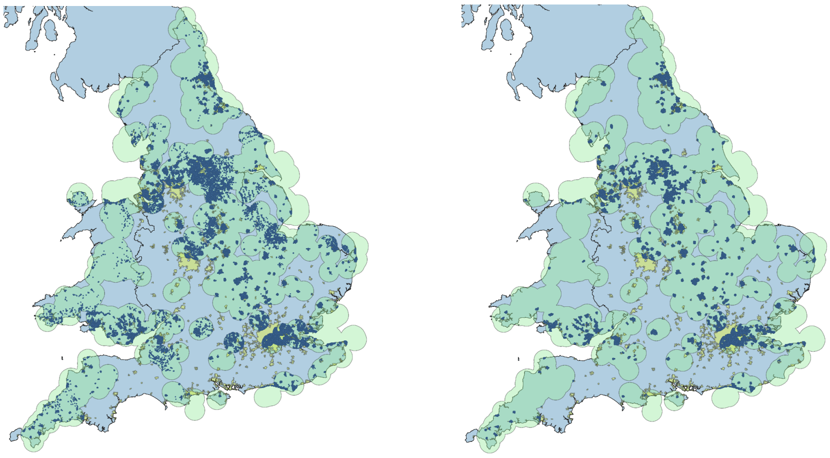

(a) Full Analysis (b) Building Height Restriction

Figure 4.8: Repeat Property Locations

on the Geographic cutoff, and further reduces it to 665,121 due to the time constraint. There is also a reduction in the number of properties to 895,261 under the specification where I restrict to repeat properties with detailed characteristics available. Summary information is presented in Table 4.2.

Table 4.2: Property Summary Statistics

	Main 1995-2018	Building Height 2011-2017	Characteristics 1995-2018
Observations	4,717,888	1,735,966	1,844,238
Properties	2,299,352	665,121	895,261
Mean Log Price	11.59	11.7	11.58
Mean Transactions	2.05	2.61	2.06
Detached	0.19	0.05	0.2
Semidetached	0.31	0.18	0.29
Flat	0.08	0.58	0.11
Terraced	0.42	0.19	0.4
Freehold	0.84	0.86	0.83

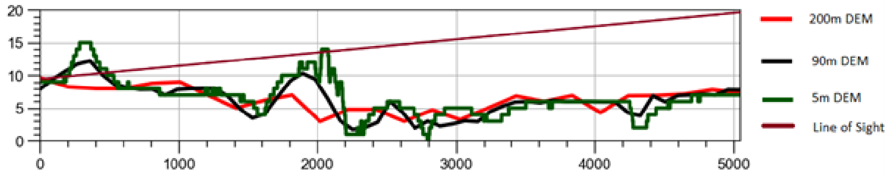

Figure 4.9: Comparison of line of sight using 200m, 90m, and 5m DEMs

4.5.3 Digital Elevation and Digital Surface Maps

Because Visibility together with distance is considered to be the main treatment effect, it is necessary to ensure accurate estimation of whether a nearby windfarm or wind turbine is visible from a given property – as misclassification of visibility will lead to mis-grouping properties into the appropriate group (treatment or control). Gibbons (2015) calculated visibility using a 200m grid elevation – elevation is estimated for a 200m x 200m square. In Chapter 3 I calculated visibility on a 200m grid and 90m grid. This provided a more granular and more accurate model of the land. However, for the repeat sales analysis, I calculate visibility using the 5m grid, and ultimately include building height information. To highlight the differences in accuracy, I show an example of all three elevation models at the same location for comparison in Figure 4.10 and Figure 4.9 gives an example of the differences in obstructions to a line of sight under each of the DEMs. Using the 200m elevation model, the line of sight is not obstructed, however for both higher resolutions, the line of sight is obstructed by natural features of the land. I have also generated a line of sight with a land profile which includes building height data in Figure 4.11. Again, differences in visibility estimates are presented in Appendices A8 and A9.

The analysis of this chapter also includes a comparison of the Elevation Map with building height data. This is data covering central and east London. For this, I convert the 5m grid to a one-meter grid (the values of the natural elevation still use the same data, but in 25 1m

(a) 200m Resolution

(b) 90m Resolution

(c) 5m Resolution

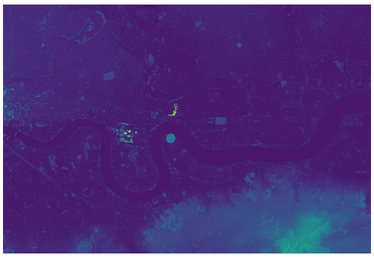

(d) 5m Resolution and Building Heights

Figure 4.10: DEM Resolutions and Inclusion of Building Height Data

squares rather than one 5m square). I then added the building height data to this natural elevation and calculated visibility for the dataset where building height information was available. This is illustrated in Figure **??**. The top right image shows the elevation model on a 200m grid, Top left displays the elevation model on a 90m grid, Bottom Right shows the elevation on a 5m Grid, and Bottom Left shows the elevation with added building height data at a 1m grid. I have included larger versions of these images in Appendix A10 Figure 4.11 shows the differences in observed elevation as a profile, with an example line of sight to highlight where different DEM resolutions could affect visibility estimates. Elevation and horizontal distance are shown in meters. Table 4.3 summarizes the source and use of the data which is used in the analysis of this chapter.

Figure 4.11 highlights the fact that even when using the DEM at the highest resolution, failure to account for intervening buildings may still lead to errors in visibility estimation. Unfortunately it was not possible

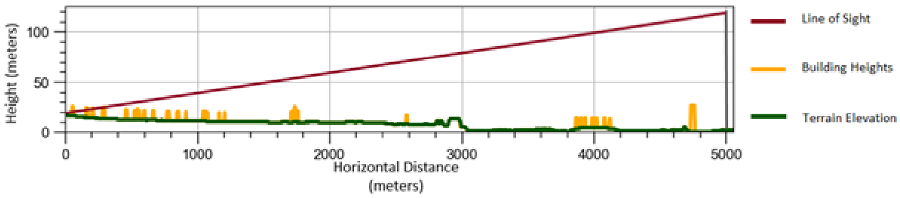

Figure 4.11: Line of Sight to Dagenham II Windfarm

to obtain a Digital Surface Map which included vegetation such as trees, but despite this the analysis here is considerably more robust than any other within England and Wales, as well as much of the wider literature.

Table 4.3: Data Sources

Data		Period	Published	Source	Application	Free Access
Housing Transactions		1995-2018	19-Mar	HM Land Registry Price Paid Data	Housing Type, price, and address	Yes
Wind Energy						
	Wind Farms	1992-2020	20-Jul	Renewable Energy Planning Database	Windfarm location, turbine height, generation capacity, status, dates of operation	Academic
	Wind Turbines	1992-2020	20-Jul	Renewables-MapUK	Turbine locations	Academic
Digital Elevation Models						
	200m DEM		2015	Digimaps	Visibility analysis and geographic controls	Academic
	90m DEM		2015	Digimaps	Visibility analysis and geographic controls	Academic
	5m DEM		2015	Centre for Environmental Data Analysis	Visibility analysis and geographic controls	Academic
	Building Heights		2014	Centre for Environmental Data Analysis	Visibility analysis	Academic
Energy Performance Certificates		1995-2018	2019	Ministry of Housing, Communities & Local Government	Property Characteristics	Academic

4.6 Results

The following section presents a series of results of the analyses described in the preceding sections of this chapter. Firstly, I present the headline results in Table 4.4, where they are presented across 10 columns. These results include a proximity alone analysis, a proximity-visibility analysis I then summarize these results, which are followed by a series of robustness checks which test key assumptions of the analysis.

4.6.1 Proximity Alone

The first two columns of Table 4.4 report the results of a proximity only analysis, column one reports the results when no geographic controls are included in the analysis and column two reports results with controls. Although I have established that the key indicator is an interaction of windfarm proximity and visibility, it is still useful to compare the reported price effects with and without a visibility indicator to highlight the fact that visibility interacted with distance is a far better proxy for determining siting effects on house prices. Comparing the first two columns, it is possible to see the importance that geographic controls play in the results, although these are broadly similar. There are statistically significant and negative impacts reported at the first two distance bands in column two of about 1%, however at the 2-4km and 8-14km range the effects are statistically significant and positive.

4.6.2 Proximity and Visibility

Columns 3, 4, 5 and 6 of Table 4.4 report results of the full dataset using a proximity and visibility interaction as the basis of the treatment. Columns 3 and 4 show the effects from visible (3) and non-visible (4) windfarms when there are no geographic controls included in the model. I find very similar results from windfarm visibility and non-visibility. Columns 5 and 6 report effects from visible (5) and non-visible (6) windfarms respectively and include geographic controls. At the 0-1km distance, there are statistically insignificant and positive effects for both visible and non-visible windfarms. However, properties within 1-2km of both visible and non-visible windfarms show a decrease in price of roughly 2%, with non-visible windfarms associated with a slightly larger drop in value. At the 2-4km band, Both results are a statistically significant and positive, close to 1% increase in price. Though non-visible windfarms are associated with a roughly 0.15% larger increase. At the 4-8km band, both visible and non-visible windfarms are again associated with a slight increase in home value of between 1.7% and 1.6% respectively. At the farthest distance band of 8-14km, it was found that both locations are associated with a slightly larger than 3% increase in price after windfarms became operational. See Table 4.5 for the DIDID results comparing these results to quantify the price effect of windfarm proximity and visibility more clearly.

4.6.3 Building Heights

Because it is assumed that visibility is the main cause of any price effect arising from windfarm siting, it is important to test the results using the most accurate visibility estimates possible. However, this required the exclusion of a substantial portion of the dataset. Nonetheless, in Table 4.4 I report the estimated price effect of windfarm proximity and visibility when building height data is included for visible and non-visible windfarms in columns 7 and 8, respectively. Under this restriction, I find only two statistically significant price effects. These occur at the 8-14km band for properties with visible windfarms, and the effect is a roughly 1.5% increase post operation. For properties with non-visible windfarms, only those located at the 4-8km band have a statistically significant price effect of +1.08%. All other results under this restriction range from positive to negative, but all are insignificant.

4.6.4 Detailed Property Characteristics

To ensure the robustness of the results, I performed the analysis for the subset of property transactions which were able to be paired with the property information from the English and Welsh EPC registry. This also led to a substantial restriction in the transactions included in the data. I report these results in columns 9 and 10 of Table 4.4 for visible and non-visible windfarms, respectively. When characteristics are included, statistically significant effects are found at the 0-1km range for properties with both visible and non-visible windfarms. For those with visible windfarms, there was a decrease of 0.75% after windfarm operation, and for those with non-visible windfarms, there was an increase of 2.24% after the windfarms became operational. At the 2-4km band, properties with nearby visible and non-visible windfarms both experience an increase in price after the windfarm has become operational. For properties with visible windfarms, there was an increase of 1.34% while non-visible windfarms increased the price by 1.98%.

Table 4.4: Headline Results

Distance Band	Price Effect Of Presence of an Operational Windfarm									
	1	2	3	4	5	6	7	8	9	10
0-1km	-0.0069	-0.0125**	0.0116***	0.0409	0.0090***	0.0446	0.0086	-0.0187	-0.0075*	0.0224**
(RSE)	(0.0059)	(0.0057)	(0.0037)	(0.0348)	(0.0035)	(0.0324)	(0.0104)	(0.0105)	(0.0042)	(0.0712)
1-2km	-0.0094***	-0.0104***	0.0081***	-0.0193*	0.0066**	-0.0232**	0.0075	-0.0015	0.0078	0.0085
(RSE)	(0.003)	(0.0029)	-0.0030	(0.0105)	(0.0028)	(0.0097)	(0.0100)	(0.0093)	(0.0048)	(0.0120)
2-4km	0.0080***	0.0081***	-0.0089***	0.0075***	-0.0081***	0.0090***	0.0042	0.0115	0.0134***	0.0198*
(RSE)	(0.0014)	(0.0014)	-0.0015	(0.0034)	(0.0014)	(0.0032)	(0.0099)	(0.0089)	(0.0029)	(0.0120)
4-8km	0.0002	-0.0009	0.0034***	0.0172***	0.0049***	0.0159***	-0.0039	0.0108*	-0.0186	-0.0181
(RSE)	(0.0008)	(0.0008)	-0.0009	(0.0014)	(0.0008)	(0.0013)	(0.0098)	(0.0061)	(0.0124)	(0.0130)
8-14km	0.0120***	0.0116***	-0.0029***	0.0320***	-0.0023***	0.0312***	0.0147**	0.0099	0.0037	0.0053
(RSE)	(0.0006)	(0.0006)	-0.0007	(0.0009)	(0.0063)	(0.00081)	(0.0062)	(0.0070)	(0.0121)	(0.0051)
Obs.	4,717,888	4,717,888	4,717,888	4,717,888	4,717,888	4,717,888	1,735,966	1,735,966	1,844,238	1,844,238
R-Squared	0.888	0.896	0.889	0.889	0.903	0.903	0.884	0.913	0.921	0.924
Geographic Controls		X			X	X	X	X	X	X
Property Controls									X	X
Proximity Alone	X	X								
Proximity and Visibility			X	X	X	X	X	X	X	X
Visible			X		X		X		X	
Not Visible				X		X		X		X
Building Height Data							X	X		

*** p<0.001
**p<0.01
*p<0.05
RSE: Robust Standard Errors Clustered at the Output Census Area
Control Vars: property slope-by-year, elevation-by-year, aspect by-year dummies. Results are reported to four decimal places.

4.6.5 DIDID

To provide a clearer estimation of the price premium or discount associated with windfarm visibility, I perform a triple difference estimation comparing the average effect from windfarm visibility relative to windfarm non-visibility. The results presented in Table 4.5 on the following page show the average price effect from non-visible windfarms subtracted from the average price effect from visible windfarms. Here I find statistically significant relative price effects at all distance bands except the 0-1km range.

There is no statistically significant effect at the closest band, and interesting the difference is not consistent. These results imply that within 1-2km, home buyers are willing to spend just under 3% more for a view WITH a windfarm relative to a view without one. The same is true at the 4-8km band, though the premium for windfarm visibility is only 0.5%. At the 2-4 and 8-14km bands, home buyers seem to be willing to pay more for a lack of view f windfarms, though these premiums are very small, both being close to 0.2%.

Table 4.5: Difference-in-Difference-in-Differences

Distance Band	1	2	3
	Visible	Not Visible	DIDID
0-1km	0.0090***	0.0446	-0.0356
(RSE)	(0.00345)	(0.0324)	(0.0325)
1-2km	0.0066*	-0.0232**	0.0298***
(RSE)	(0.00282)	(0.0097)	(0.0106)
2-4km	-0.0081***	0.0090***	-0.0017***
(RSE)	(0.0014)	(0.0032)	(0.0004)
4-8km	0.0049***	0.0159***	0.0049***
(RSE)	(0.0008)	(0.00130)	(0.0008)
8-14km	-0.00228***	0.0031***	-0.0023***
(RSE)	(0.0063)	0.0009)	(0.00063)
Obs.	4,717,888	4,717,888	4,717,888

*** p<0.001
**p<0.01
*p<0.05
RSE: Robust Standard Errors Clustered at the Output Census Area
Control Vars: property slope-by-year, elevation-by-year, aspect by-year dummies. Results are reported to four decimal places.

4.6.6 Additional Robustness Checks

Here I present additional results testing for robustness to the study period and visibility intensity.

Study Period

To test for robustness to the study period, I break the analysis into 4 smaller time periods which provide the average price effect over the periods 1995-2000, 2001-2006, 2007-2012, and 2012-2018. This is presented in Table 4.6 and shows how the preferences towards windfarms have changed over time. There is considerable variation in the reported results across all periods for both visible and non-visible windfarms. When disaggregated across smaller time periods, there does appear to be rather more negative effects arising from visible windfarms than from non-visible windfarms, Though by the final period, I show that windfarm visibility is associated with some very large increases in prices.

View Intensity

In Table 4.7 I present the price effects arising from differences in the number of visible and non-visible wind turbines, or wind turbine view intensity. Accounting for intensity of treatment by visible or not visible turbines is a potentially important robustness check as the presence of several turbines very close to a property would arguably result in a much larger impact on the views from any property than the presence of relatively few turbines at larger distances. This analysis is restricted to include only repeatedly sold properties which are within 5km of an operational wind turbine. I further restrict the dataset to include properties with turbines at only a single distance band to avoid any spillover effects arising from treatment at several distance bands. The intensity has been grouped into four intensity categories. These intensity categories are: presence of a single turbine, 2-4 turbines, 5-10 turbines, and more than ten turbines. These are then grouped by distance from the property to determine if there are differing price effects arising from a larger number of turbines at a given distance from a property.

Table 4.6: Study Period Subdivisions

Distance Band	1	2	3	4	5	6	7	8	9	10
0-1km	0.009***	-0.0293***	-0.0154***	0.0106	0.0302***	0.0446	-0.0209***	0.0245*	0.0763***	0.0458
(RSE)	(0.0035)	(0.0028)	(0.0008)	(0.0829)	(0.0077)	(0.0324)	(0.0058)	(0.0147)	(0.0091)	(0.0745)
1-2km	0.0066*	0.0337	-0.0165***	0.0094	-0.0123**	-0.0232**	0.0726***	0.0676***	-0.0205***	-0.0428*
(RSE)	(0.0028)	(0.0243)	(0.0023)	(0.0407)	(0.0062)	(0.0097)	(0.0018)	(0.0067)	(0.0076)	(0.0225)
2-4km	-0.0081***	-0.0363	-0.0113***	0.0278***	-0.0008	0.0090***	0.0853***	0.0131	0.0296***	0.0022
(RSE)	(0.0014)	(0.0445)	(0.0011)	(0.0036)	(0.0029)	(0.0032)	(0.0125)	(0.0033)	(0.0090)	(0.0073)
4-8km	0.0049***	0.0237**	0.0123*	-0.0187***	0.0199***	0.0159***	0.0912***	0.0451***	0.0261***	-0.0039**
(RSE)	(0.0008)	(0.0105)	(0.0043)	(0.0020)	(0.0031)	(0.0013)	(0.0072)	(0.0019)	(0.0039)	(0.0020)
8-14km	-0.0023***	-0.0012	-0.0455	0.0185***	-0.0056***	0.0031***	0.0737***	0.024***	-0.0021	0.0073***
(RSE)	(0.0063)	(0.0067)	(0.259)	(0.0020)	(0.0016)	(0.0009)	(0.0060)	(0.0016)	(0.0017)	(0.0024)
Obs.	4,717,888					4,717,888				
R-Squared	0.847	0.866	0.898	0.911	0.904	0.850	0.866	0.898	0.911	0.904
Extension	X					X				
1995-2000		X					X			
2001-2006			X					X		
2007-2012				X					X	
2013-2018					X					X
Visible	X	X	X	X	X					
Not Visible						X	X	X	X	X

*** p<0.001
**p<0.01
*p<0.05
RSE: Robust Standard Errors Clustered at the Output Census Area
Control Vars: postcode slope-by-year, elevation-by-year, aspect by-year dummies, proportions of sales of detached, semi-detached, terraced, flat/maisonette;quarterly dummies. Results are reported to four decimal places.

Table 4.7: Turbine Visibility Intensity

Distance Band	1	2	3	4	5	6	7	8
0-1km	-0.0074	-	-	-	0.0170	-	-	-
(RSE)	(0.0730)				(0.0180)			
1-2km	0.0339	0.0064	-	-	-0.0053	-	-	-
(RSE)	(0.0618)	(0.0082)			(0.0529)			
2-3km	0.0404	0.0260	-0.0176***	-0.046***	0.0523***	0.0751	-	-
(RSE)	(0.0560)	(0.0197)	(0.0041)	(0.0009)	(0.0105)	(0.0483)		
3-4km	0.0342	-0.0115	0.0166	-0.0222	0.0410***	0.0188*	0.0038	-
(RSE)	(0.0544)	(0.0237)	(0.0207)	(0.0198)	(0.0110)	(0.0095)	(0.0574)	
4-5km	0.0308	0.0083	-0.0357	0.0056	-0.0158	0.0307	0.0041	0.0560
(RSE)	(0.1580)	(0.0920)	(0.1240)	(0.0847)	(0.0178)	(0.0181)	(0.0692)	(0.0250)
Number of Turbines								
1	X				X			
2-4		X				X		
5-10			X				X	
10+				X				X
Visible	X	X	X	X				
Not Visible					X	X	X	X

*** p<0.001
**p<0.01
*p<0.05
RSE: Robust Standard Errors Clustered at the Output Census Area
Control Vars: postcode slope-by-year, elevation-by-year, aspect by-year dummies, proportions of sales of detached, semi-detached, terraced, flat/maisonette;quarterly dummies. Results are reported to four decimal places.

4.7 Discussion

The Results presented provide evidence that wind farm visibility is both an amenity and disamenity through the application of a fixed effects hedonic model of valuation. The Results as presented in table 4.5 show that home buyers will pay a premium for homes where there is a view of a nearby windfarm – but only at certain distances (1-2km and 4-8km). Buyers will also pay a premium for properties with no view of a windfarm if the windfarm is sited between 2-4km or 8-14km from the property. This is interesting as it implies there is not a consistent attitude towards wind turbine visibility and proximity. However, it must be restated that this analysis has observed only properties which have sold repeatedly over the period, though we assume any unobserved effects have been captured by the fixed effects model. This inconsistency could be explained by several factors, such as the time the transactions occur, as well as the number of nearby turbines which I test and presented results for in tables 4.6 and 4.7 respectively.

The results of the headline analysis fit well with those found by Heblich et al. (2016) insofar as there is no consistent negative effect arising from either windfarms which are visible or those which are not. The statistically significant effects found here are smaller, and though do not agree with the direction of the effect, are more in line with Dröes & Koster (2016) who find a decrease of about 1% but only within 2k of a wind turbine than studies such as Lang et al. (2014), Hoen et al. (2011), or McCarthy & Balli (2014) all of which found no statistically significant impact on house prices arising from windfarm visibility. I find no evidence of the severe statistically significant negative impacts reported by Heintzelman & Tuttle (2012) or Sunak & Madlener (2016).

Interestingly, the results of the repeat sales analysis within this chapter don't match up with the results reported in Chapter 3 as closely as had been expected. Under the extended analysis (presented in Table 3.8) I found consistently positive price effects from windfarm visibility at all distance bands, with significant increases of 2.51%, 2.9%, and 6.16% at the 2-4km, 1-2km, and 4-8km bands respectively. The observed effects here while statistically significant at all distance bands were near zero and negative at the 2-4km and 8-14km bands. I also found consistently positive, though not necessarily statistically significant effects from non-visible windfarms within the average sales analysis.

This is evidence that the analyses in the literature may be sensitive to the application of repeat vs average sales analyses. This is sensible when one considers what is actually being estimated across the two analyses. In the average sales model, visibility is estimated using the postcode centroid. Put in another way, this is a measure of whether a windfarm is visible from the neighborhood where a property is located. The repeat sales approach is instead estimating price effects from visibility at the property level itself. Finding a different effect under the two approaches is not unexpected as this is the case for other amenities Bourassa et al. (2004)

It should be noted here that selection may be an issue and it is not clear what is driving the difference in estimates between the Average and Repeat Sales headline results - as this could arise from either the sample itself or the inclusion of fixed effects at different scales.

In addition to the hedonic approach, visibility estimation is highly important to the analysis. When accounting for intervening buildings, the statistical significance of the results disappear for all but the 8-

14km distance band with visible turbines increasing transaction prices by about 1.5%, and non-visible turbines increasing transaction prices by about 1%. This highlights the importance of an accurate visibility estimate. Although Sunak & Madlener (2016) are able to make use of a higher resolution DSM which includes both building heights and vegetation, it is unlikely that this is accurate for the entirety of their study period which covers 1992-2010. They do not list the year that their DSM was published. Similarly, Heblich et al. (2016) also make use of a high resolution DEM and building height data, but do not restrict their study period under this analysis, making it unlikely that the visibility estimates are accurate for the study period of 1992 - 2014.

Therefore, despite the restriction to transactions occurring only three years before or after the publishing of the Building Height data used within this chapter's analysis, I provide what is likely a far more accurate visibility estimation for this specification. I should note of course, that the time restrictions on the building height analysis lead to a study period which overlaps with a time when negative impacts of windfarm siting are diminishing as shown in the Study Period analysis.

Because the analysis of this chapter applies a revealed choice method, where we assume the preferences towards wind turbine proximity and visibility are revealed through property transaction prices. As in Chapter 3 I test for changes in preferences over time. I have found that attitudes towards windfarm visibility have been changing, with a larger abundance of negative price effects from wind turbine visibility in the earlier transactions and most recently it would seem that home buyers are less sensitive to the presence of wind turbines. This may imply that preferences towards wind turbine proximity and visibility are not stable over time, which could be explained individuals simply becoming accustomed to the sight of windfarms and simply accepting these as part of the landscape. An alternative explanation is that preferences towards windfarm visibility do remain stable over time as has been found previously by Parmeter & Pope (2013), but the increasing inclusion of recreational amenities and community benefit programs over-compensates for the visual disamenity that windfarms generate. Again, the impacts do not match exactly to those of the average sales analysis of the preceding chapter, but they are in agreement insofar as both find more evidence of negative effects in the earliest study periods and largely positive effects in the most recent ones. This may partly explain some of the negative findings within the wider literature as many studies are skewed towards the earlier ends of the analysis performed here and in Chapter 3. The majority of the papers in the wider literature, only Heblich et al. (2016) includes data more recent than 2012.

In addition to exploring changes in the effects over time, I test whether there are differences arising from the number of wind turbines at a given distance band. Since the presence of several turbines may affect vistas differently than presence of a single turbine or relatively few turbines, this may be an important robustness check. Particularly because many windfarms consist of a single turbine and may be very different from windfarms which consist of a few hundred turbines in terms of their impact on property vistas. Table 4.7 indicates that there is a difference in effects arising from the presence of a single turbine versus several turbines at the same distance from the property. In regards to policy implications, this intensity effect could be further explored under a variety of assumptions to determine if there is an optimal wind farm size and distance from residential properties, as well as accounting for the density of properties, that minimizes property effects while still providing the benefits of wind energy development.

4.7.1 Limitations of the Analysis

By definition, the properties included in a repeat sales analysis are restricted to those sold multiple times over the period. This reduces substantially the size of the housing transactions dataset which is used to value the presence of nearby and visible windfarms. Secondly, by making this restriction, it is possible that there is a sample selection bias due to the fact that all properties were sold multiple times which may imply that they are generally more desirable than properties which did not experience repeated sales over the period. However, these are limitations of all repeat sales analyses and not unique to the present analysis. Indeed, the limitations of this analysis are partly compensated by the strengths of the average sales analysis of 3.

A limitation of all analyses within the literature suffer from is the visibility estimation. This analysis is no different. Despite generating the most detailed visibility model for England and Wales, and one of the most detailed within the wider literature there are still limits to the application of GIS modeling. Consider firstly that here, I use the center of a property as the view point to estimate visibility of turbines and windfarm centroids. However, it is possible that this does not accurately reflect visibility from windows of the property or from their yards/gardens. Secondly, McCarthy & Balli (2014) find differences between their visibility estimation using GIS and field visits, which implies that even with a very accurate model there may be misclassification. Again, it was not feasible to physically visit individual properties due to the size of the dataset utilized here.

4.8 Conclusion

This chapter estimated the values of windfarm proximity and visibility on house prices through the application of a repeat sales hedonic pricing analysis, within a spatial fixed-effects DID model. Broadly speaking, there was no consistent effect of windfarm visibility or lack of visibility on nearby house prices. The analysis of this chapter and of 3 contribute to the literature by updating the windfarm amenity literature within England and Wales to include a period where there has been substantial changes to the landscape since the analysis by Gibbons (2015). This chapter also presents the first repeat sales analysis using data from England and Wales, and allows for a testing of the estimated impacts robustness to repeat vs average sales techniques. Further, I find evidence across both chapters that preferences towards windfarm visibility is not stable over time, and that broadly speaking acceptance towards and even a preference towards windfarm visibility exists. In addition, I find some evidence that home buyers may place different amenity value on owning a property in a neighborhood with windfarm visibility as opposed to living in a property with its own direct views of wind turbines.

The findings of this chapter have several implications for policy makers. First, I find that early on, there were negative price effects from windfarm viability and proximity. There may be justification in compensating affected property owners, though as mentioned in the previous chapter this may be difficult. Second, there is no consistent negative impact from windfarm proximity and visibility. This should be taken into account when choosing the sites for windfarms. There is evidence that windfarm visibility may lead to increased home values for properties located within 2km of an operational windfarm. This information should be shared with communities where development applications are being lodged, as it may bolster local support for such developments. Third, I do find evidence that properties located near multiple visible wind turbines doe experience statistically significant price decreases. Policy makers should therefore continue to ensure that new wind developments are at least 3km from properties to avoid these negative effects. At this distance, effects from a single turbine are statistically insignificant though positive - but when 5 or more wind turbines are within this distance, statistically significant price reductions are observed.

Chapter 5

A hedonic analysis of residential PV systems in England and Wales

5.1 Introduction

Over the past 30 years, there has been a sustained interest in reducing carbon emissions through transitioning electricity generation to renewable sources (BEIS 2020*a*). In many countries, and particularly within the UK, this transition has been largely composed of increasing the supply of electricity generated by both on and offshore wind by large-scale commercial wind farms. At the same time, the UK has committed to retiring its current stock of coal fired generation by 2025, and is replacing these with mostly renewable generation technologies (Littlecott 2016). Wind energy is currently the cheapest form of renewable energy and the UK is endowed with considerable wind energy generation potential (Gibbons 2015, Heblich et al. 2016, OFGEM 2020*b*). The rapid deployment of wind turbines across England and Wales was the primary focus of the previous two chapters, which explored the effects of large-scale, commercial wind energy generation on house prices. I found that there were price effects, which were significant and negative during the earliest periods of the analysis. By the end of the study period, these had shifted to either insignificant or significant and positive effects for all but properties closest to the wind turbines. Although these wind farms required approval, planning refusal at the local authority level could be overridden by the Secretary of State. And even with a refusal by the local planning authority, there is no discretion at the household level regarding the decision for wind farm siting. Therefore, any house price effects arising from windfarm proximity were caused by factors external to the preferences of any individual property owner. And any effects observed in price effects were the result of the preferences of prospective home buyers.

It is worth noting that although residential solar has been consistently found to capitalize positively into house prices, this is not necessarily the case for commercial-scale solar farms. Dröes & Koster (2021) explored whether commercial scale solar farms have any impacts on nearby house prices and find that they decrease transaction prices for properties within 1km by about 2.6%. This may imply that for the UK it is more sensible to support commercial-scale wind generation. As was reported in Chapters 3 and 4 windfarms are currently associated with positive price impacts on nearby properties. Even if a similar impact is found for commercial-scale solar farms, the disparity between the wind energy and solar energy endowments for the UK would imply that given the opportunity costs of investing in solar, wind will generate greater returns. However, this chapter explores the impacts of residential solar on house prices within the UK.

At the same time, there has also been a rapid and growing deployment of domestic renewable electricity generation within both England and Wales. The boom in residential solar installations takes a similar trajectory to commercial wind shown in Chapters 3 and 4, with a spike occurring in 2010. Unlike commercial developments, the decision to install residential solar or wind generation is entirely at the discretion of property owners. The decision to install solar panels is likely based upon a financial and personal preference basis (Black 2004, Dastrup et al. 2012). Here the financial incentives will arise from costs of installation, savings from domestic generation and the presence and size of any green subsidies.

Regarding the non-pecuniary factors which influence installation, there is some evidence for a 'keeping up with the Jones' effect within the literature (Dastrup et al. 2012). When estimating the price effects of a residential PV system on a property, it is the personal preferences of prospective home buyers that will determine if any solar premium exists. Such a premium may arise from a few avenues, such as the desire to reduce their own carbon footprint, or to support the transition to low-carbon energy generation at the individual level. The price premium may even be a means to signal an individual's commitment to reducing emissions or the importance placed on environmental protection to others (Dastrup et al. 2012).

The growing adoption of residential solar generation has sparked academic interest in evaluating the extent to which solar photovoltaic (PV) systems are capitalized into home values. Though the literature is small, all research into the capitalization of residential PV has been performed in 'sunny' countries or regions which have substantially higher PV generation potential compared to the UK. The literature include works by Hoen et al. (2013) and Dastrup et al. (2012) which found that the solar property premium in California was large enough to recoup installation costs or generate profits, respectively. More recently, Ma et al. (2016), Wee (2016) and Lan et al. (2020) found solar installations costs are fully recovered through increased sales prices in Western Australia, Hawaii, and Queensland, Australia respectively. Lastly, a paper by Qiu et al. (2017) reports that in Phoenix Arizona, the solar premium leads to increased property values of up to $45,000 with a substantial profit over the costs of installation.

The United Kingdom has also experienced a boom in solar PV installations, with reported installations in just over 800,000 properties in England and Wales presently OFGEM (2020*b*). The largest increase was between 2010 and 2014, but growth has continued to the present day. In this chapter, I extend the literature on the solar premium and capitalization by analyzing price effects through 3 statistical models under a variety of assumptions. The first two models reflect approaches taken within the current literature to value the solar premium and the third applies a propensity score matching approach and takes a quasi-experimental design – the first to be applied within this literature. The matching analysis is the largest performed to date, and unlike Qiu et al. (2017) and Lan et al. (2020), I use a national dataset rather than a single city case study. In addition, this is the first analysis to estimate the extent to which solar installations within the UK are capitalized into house prices.

I find evidence that a solar premium exists in England and Wales. Under two of the empirical models, I find the size of the premium is large enough that costs of installations will be fully recovered and may lead to a profit for the average installation. These findings are robust to a number of sensitivity analyses.

5.2 Solar PV Capitalization in the Literature

Here, I will present a discussion of the papers which have reported values of solar premiums or capitalization rates from the literature. These papers have broadly informed the structure of the analysis outlined in Section 5.4 and the results presented in Section 5.5. Firstly, I will discuss the general methodologies applied within the literature and how they have influenced the methods applied in the analysis of this chapter. Subsequently, I discuss the key findings of the papers and lastly I highlight how the current analysis fits into the literature.

Table 5.1 highlights the key features and findings of the analyses from the literature. First, the have thus far included a very small selection of residential transactions where solar PV have been installed. Second, there are three broad methodological approaches taken to value the solar premium/capitalization into house prices. Five of the six papers apply some form of the basic hedonic regression model, half go on to apply a repeat sales analysis, while two utilize control matching techniques. Third, despite the different econometric approaches taken and the broad geographic distribution of study areas, there is agreement within the literature: solar panels are capitalized into house prices and the solar premium is large enough to fully recover or profit from the cost of installing a residential PV system.

Table 5.1: Residential PV Capitalization Literature

Author(s)	Period	Transactions	PV Homes	Study Area	Hedonic Regression	Repeat Sales	Control-Matching	Solar Premium	Recover Installation Costs	Profit Over Installation Costs
Dastrup et al.	1997-2010	364,992	329	California	Y	Y	N	3.50%	Y	Y
Hoen et al.	2000-2009	72,319	894	California	Y	N	N	3.60%	Y	-
Ma et al.	2009-2012	25,970	413	Western Australia	Y	Y	N	2.3-3.2%	Y	-
Wee	2000-2013	47,696	259	Hawaii	Y	Y	N	5%	Y	-
Qiu et al.	2014	246	123	Arizona	N	N	Y	15-17%	Y	Y
Lan et al.	2008-2018	315	228	Queensland, Australia	Y	N	Y	$21,403 AUD	Y	Y

5.2.1 Hedonic Regression

A key assumption of hedonic valuation approach is that the price of a given good is determined by the values a buyer or seller ascribe to both the internal and external characteristics of that good (Monson 2009). As outlined in the previous chapters of this thesis, the price is essentially a sum of the values of the features of the good, it is possible to generate an estimate of each characteristic by disaggregating them (Rosen 1974, Blumenschein et al. 2008). In the context of this chapter and the literature reviewed in this section, the 'good' is a residential property. The intrinsic characteristics are the features of the property itself such as its construction type, floor area, number of rooms, lot size, etc. Extrinsic characteristics are features of the area where the property is located, such as the school catchment area, local crime rate, proximity to local amenities, etc. Equation 5.1 is the generic model applied to estimate the disaggregated values of internal and external characteristics which taken together are the price of a property and will be a familiar feature of this thesis.

$$Price = f(intrinsicfeatures, extrinsicfeatures, time) \tag{5.1}$$

In its simplest form, the hedonic valuation is based on a regression equation where price is the dependent variable, and the characteristics are independent variables. The majority of studies applying a hedonic analysis have used a semi-logarithmic function to estimate the value of property characteristics under a hedonic regression model (Lan et al. 2020). Studies valuing solar premiums have taken this approach, either in their main analysis (Hoen et al. 2013) or as a benchmark for subsequent analyses (Dastrup et al. 2012, Ma et al. 2016, Wee 2016, Lan et al. 2020). Qiu et al. (2017) are unique within the literature in that they do not first make a logarithmic transformation of their dependent variable (price), but rather perform their hedonic regression using the raw values of sales prices from their dataset.

The general approach of these analyses, as well as the benchmark analysis described in the subsequent section of this chapter is to disaggregate and control for the various influences of property characteristics, as well as time on house prices. This is done by using the regression coefficients to reflect the contributions of each characteristics, and time to the total price of a given property, and by taking the coefficients of the regression on an entire dataset, estimate an average price or value of each component. Under this approach, in addition to controlling for property specific characteristics, it is imperative to control for time and fixed effects (Gibbons 2015) to avoid biases in the results.

Applying this approach as the main analysis of their paper, Hoen et al. (2013) report an average solar premium of 3.6% of the average property value in their dataset, which translates to a price premium of $17-38,000 relative to non-solar properties in the same zip code within California. The other papers which performed a hedonic approach used it to benchmark their headline results – either against a Repeat Sales Analysis or a Control-Matching analysis and these approaches are discussed in the following two sections.

5.2.2 Repeat Sales Analyses

An alternative to the standard hedonic regression model is a repeat sales model. A common problem with the basic hedonic regression is the omitted variable bias, which in this case, arise from unobserved property-specific factors. The inclusion of fixed effects in the model can address the potentially omitted variables. Though when this is performed at the zip code level, as was the case for Hoen et al. (2013), there is still a reasonable possibility that some potentially important factors are omitted. The repeat sales approach addresses this issue by comparing properties to themselves, rather than to nearby properties.

This basic idea underpinning this approach is that it compares the change in price between two sales of a given property, while controlling for intrinsic and extrinsic characteristics, time and fixed effects. Because this approach compares a property to itself, it is imperative that both the timing of and the time between sales is accounted for to ensure that price effects are not simply a reflection of house price inflation (Dastrup et al. 2012, Ma et al. 2016, Wee 2016). A repeat sales analysis essentially exploits the fact that the only difference between two transactions is the presence of a new feature of that property. This is then added to the regression model (usually as a dummy indicator) to allow for the estimation of of how much of the change in price between sales is caused by installing solar panels. For present purposes, the feature of interest which is allowed to change between transactions of a given property is the presence of residential solar installations.

A repeat sales model will measure the average additional appreciation of properties which have installed PV systems between consecutive sales and properties located in the same area which have not installed a PV system in the same period. An advantage of this approach is that by comparing a property to itself, a repeat sales analysis eliminates the potential biases in selecting an appropriate control group faced by the standard approach. Under the assumptions of a repeat sales analysis, the transactions are restricted to properties which have sold multiple times, eventually install residential solar PV and have sales during the study period.

This approach was used to generate the headline results for Dastrup et al. (2012) who report and average solar capitalization of 3.5% for properties located near Sacramento and San Diego California. As is the case with all other repeat sales within the literature, Dastrup et al. (2012) define their dependent variable as the natural log of the change in price between sales of a given property. The coefficient of their PV installation dummy variable, which takes a 1 if PV was installed between transactions and 0 otherwise is the percent of the sales price attributable to the presence of the solar panels. This translated into a solar premium of $22,554. This is slightly below the cost of installing the average PV system within the analysis. However, state and federal subsidies towards the purchase, financing, and installation lowered the cost to the extent that home sellers would fully recoup and profit from their installations. Ma et al. (2016) found a similar result for properties in and around Perth, Western Australia where they estimated the capitalization rate to be 2.3-3.2%, which would fully recover the costs of installations analyzed in their study. Lastly, Wee (2016) finds a capitalization rate of 5% of the property value which again, leads to a full recovery of the installation costs through the sale of the property.

The repeat sales approach has the potential to address the omitted variable bias relative to the basic hedonic regression models due essentially controlling for fixed effects at the property level rather than the neighborhood level (if controlled for at all). One potential source of bias in any repeat sales model is the assumption that implicit prices are constant. By implicit prices, I mean the value of each property characteristics which combined equal the transaction price of the property. This assumption likely holds true in the short term, but this assumption becomes weaker as the time between transactions increases Case & Quigley (1991). One method to address this issue within a repeat sales model is to control for the time of the sequential sales, as well as the time between these sales. As an alternative, some papers in the literature have performed a control-matching analysis, which I describe subsequently.

5.2.3 Control Matching

The final approach to estimating the capitalization of solar panels into property values that will be discussed is the control matching approach. As this technique forms the basis for the headline results of this chapter and has not been discussed in the previous two chapters of this dissertation, I discuss the literature and methods used more thoroughly than the recaps on hedonic regressions and repeat sales approaches above.

In a general sense, control matching imposes a quasi-experimental design on the analysis by comparing a treatment group to a control group. In the strictest set up, the two groups are identical on all control variables with the exception that the treatment group receive a treatment and the control group does not (Austin 2011). However, within the literature, there has not been an analysis which has a large enough dataset to perform an exact matching approach - the papers in the literature instead apply inexact matching techniques. In the context of the literature and this chapter, the treated group are properties which install a residential PV system and the control group are properties which do not. Control matching essentially creates a set of control properties which are identical to the treated properties, prior to the installation of a solar panel (Lan et al. 2020). The key assumption of this approach is that: if two properties are sold in the same location, at the same time and have identical sets of intrinsic and extrinsic features - with the exception of the presence of a residential PV system, then any observed difference in price can be attributed to the PV system.

Within the literature on solar PV capitalization, only Qiu et al. (2017) and Lan et al. (2020) use a matching approach to select controls and apply this research design. Both papers apply a fuzzy matching design – their datasets limited them to matching treatment properties to controls using inexact matching on characteristics. For example, when matching on the number of bedrooms, Lan et al. (2020) allowed for matches of ±1 – so a treated property with 4 bedrooms could be matched to a control property with 3, 4 or 5 bedrooms. Qiu et al. (2017) apply Coarsened Exact Matching (CEM) which uses an algorithm to transform variables from discrete values into categorical ranges for both the treated and control groups. These categories were then matched exactly to generate a matched control group. In Table 5.2 I list the matching variables for Qiu et al. (2017), Lan et al. (2020), and the matching performed in Model 3 of this chapter.

The matching approach used in the analysis of this chapter varies from that of the literature. Rather than applying techniques such as fuzzy matching following Lan et al. (2020) or coarsened exact matching asQiu et al. (2017) do, I apply a Propensity Score Matching (PSM) approach. PSM estimates the house price effect of installing residential PV by accounting for the property characteristics which predict installation of a residential PV system. Ultimately, PSM creates a sample of control properties (no solar PV) which are comparable on all observed characteristics to the sample of treatment properties (solar PV). The analysis herein makes use of a national dataset of property transactions and characteristics which allows for exact matching on characteristics between the set of properties which install solar PV and those which do not. This is a significant improvement over the single-city case studies within the current literature.

Table 5.2: Control-Matching variables from the literature and Model 3

Match variables		
Fuzzy Matching		Exact Matching
Lan et al.	Qiu et al.	Model 3
Number of Bedrooms	Number of Bedrooms	Number of Rooms
Number of Bathrooms	Number of Bathrooms	Number of Fireplaces
Number of Parking Spaces	Square Footage	Floor area (m2)
Land size	Lot Square Footage	Property Type
Swimming Pool	Pool	New Build
Grage	Year Built	Tenure Type
Undercover Parking	Good Views	Construction Age Band
Ensuite		Sale Year
Study		Postcode
Built in Robes		
Alarm System		
Gym		
Rumpus Room		
Workshop		
Air Conditioner		
Solar Hot Water		

5.2.4 Propensity Score Matching

Randomized controlled trials (RCT) are considered to be the best approach for estimating treatment effects because random assignment to the treatment or control groups ensures that there is no confounding of treatment effects with known or unknown baseline characteristics and this allows for the effects of a treatment on an outcome to be measured by comparing outcomes between treated and untreated cases (Austin 2011). However, most observational studies do not allow for a random assignment to either a treatment or control group, leading to potential bias in the selection and assignment into treatment or control groups.

This is the case for the research discussed above, as well as for the analysis presented in the subsequent sections of this chapter – residential solar PV systems are not installed at random. Installation is a choice made by homeowners or developers based on their own preferences, incentives and situations. Although an RCT model is the gold-standard for estimating treatment effects, there alternative methods which can simulate an RCT or make adjustments to correct for potential selection biases.

Propensity score matching (PSM) is a statistical technique which simulates an RCT where a well-defined treatment case is matched with one or more control cases based on each case's propensity score (Randolph & Falbe 2014). A propensity score, which forms the basis for matching treatment to control cases, is the probability of receiving a treatment conditional on a vector of baseline characteristics (Austin 2011) and the estimated score is the predicted probability of treatment derived from a fitted regression model, often a logistic regression. Here, a treatment case is a property which installs a residential solar PV system, and a control case is a property which does not install such a system. The baseline characteristics are the intrinsic and extrinsic features of a given property – number of rooms, local school quality, etc. A key assumption of the PSM approach is that treatment assignment is strongly ignorable (Rosenbaum & Rubin 1983). Ignorability refers to what extent cases assigned to treatment or control groups is irrelevant for the data analysis (Rubin 1978). Strong ignorability implies that non-random selection for treatment will not bias the results, and for an unbiased analysis using the PSM, the model must be strongly ignorable (Rosenbaum & Rubin 1983).

For a model to have strong ignorability, there are two key conditions which must be met. Firstly that treatment assignment is independent of the potential outcomes, conditional on the observed baseline covariates (Austin 2011, Rosenbaum & Rubin 1983). In the context of this chapter, this means that the characteristics of the properties in the dataset do not affect the properties' probability of installing a PV system. The second condition requires that all cases – treatment or control – have a non-zero probability of receiving or not receiving the treatment. Rosenbaum & Rubin (1983) show that if assignment to the treatment group is strongly ignorable then conditioning via the propensity score allows for unbiased estimates of the average treatment effect. This estimate is referred to as the average treatment effect of treatment on the treatment group (ATT) (Imbens 2004). Because installation of residential PV systems is voluntary, the ATT reports the effect of installation for properties where owners have chosen to install such a system.

To summarize, the PSM approach estimates the effects of treatment by accounting for the covariates which predict receiving the treatment. Within the context of this chapter, the covariates are the observed property characteristics, and the treatment is the installation of a residential PV system. The strength of the PSM approach is its ability to reduce the confounding variable bias inherent to an estimate of the PV installation price effect in a comparison of outcomes between the solar and non-solar property groups (Rosenbaum & Rubin 1983). These biases arise because installation of residential PV systems is not random, so there is potential that the difference in outcomes between the two groups is caused by some factor which predicts installation of a residential PV system, rather than installation itself (Rosenbaum & Rubin 1983, Imbens 2004). PSM addresses the non-random installation of PV by creating a sample of non-solar properties which are comparable to the solar properties with similar or the same observed characteristics.

5.3 Contextual Background

In the next section, I discuss the current state of the literature on the capitalization of residential solar PV systems into house prices and the analytical approaches taken therein. The literature reports evidence not only that there exists a solar property premium – a premium above that of comparable properties lacking solar PV systems, but also that this is of a value large enough to compensate for the installation costs associated with such systems. However, it should be noted that the current literature examines property transactions which occur in Arizona, California, and Hawaii in the United states and Western Australia and Queensland in Australia – locations which are relatively well known for warm and sunny weather. The United Kingdom, though it has a substantial wind energy endowment, is relatively sparse for solar energy (SolarGIS 2020). Figure 5.1, on the following page, presents the differences in solar photovoltaic potential of California (Left) and England and Wales (Right). Cities where studies on residential solar PV capitalization have been labeled. Hoen et al. (2013) reports capitalization from across California, whilst Dastrup et al. (2012)) reports capitalizations for properties in the Sacramento and San Diego areas.

The sizes of these locations are not shown to scale, but the average annual PV generation potential is consistently displayed across each region. It is immediately clear that the regions which have been the focus of the research within the Residential Solar Capitalization literature has been on locations with considerably larger solar energy resources than England and Wales. In fact for the majority of these regions, locations with the poorest PV are roughly on par with the locations in England

and Wales with the highest PV potential.

Within the literature, the solar premium has been largely explained through energy savings, profits from export of excess generation (Dastrup et al. 2012, Hoen et al. 2013, Ma et al. 2016, Wee 2016, Qiu et al. 2017, Lan et al. 2020). It is reasonable to assume, that given the same technology and the same PV installation size, one located in an area with higher generation potential would produce larger savings and more excess electricity for export. Therefore, the differences in potential solar generation will lead to differences in the total energy savings or profits from export of excess generation to the grid. Of course, local electricity costs and export rates will also influence these values. Regardless, despite the lack of a solar resource endowment this has not prevented England and Wales from experiencing a relatively rapid deployment of residential solar PV systems. Comparisons between the UK PV potential and other study areas from the literature are presented in Appendices A11, A12, A13, and A14.

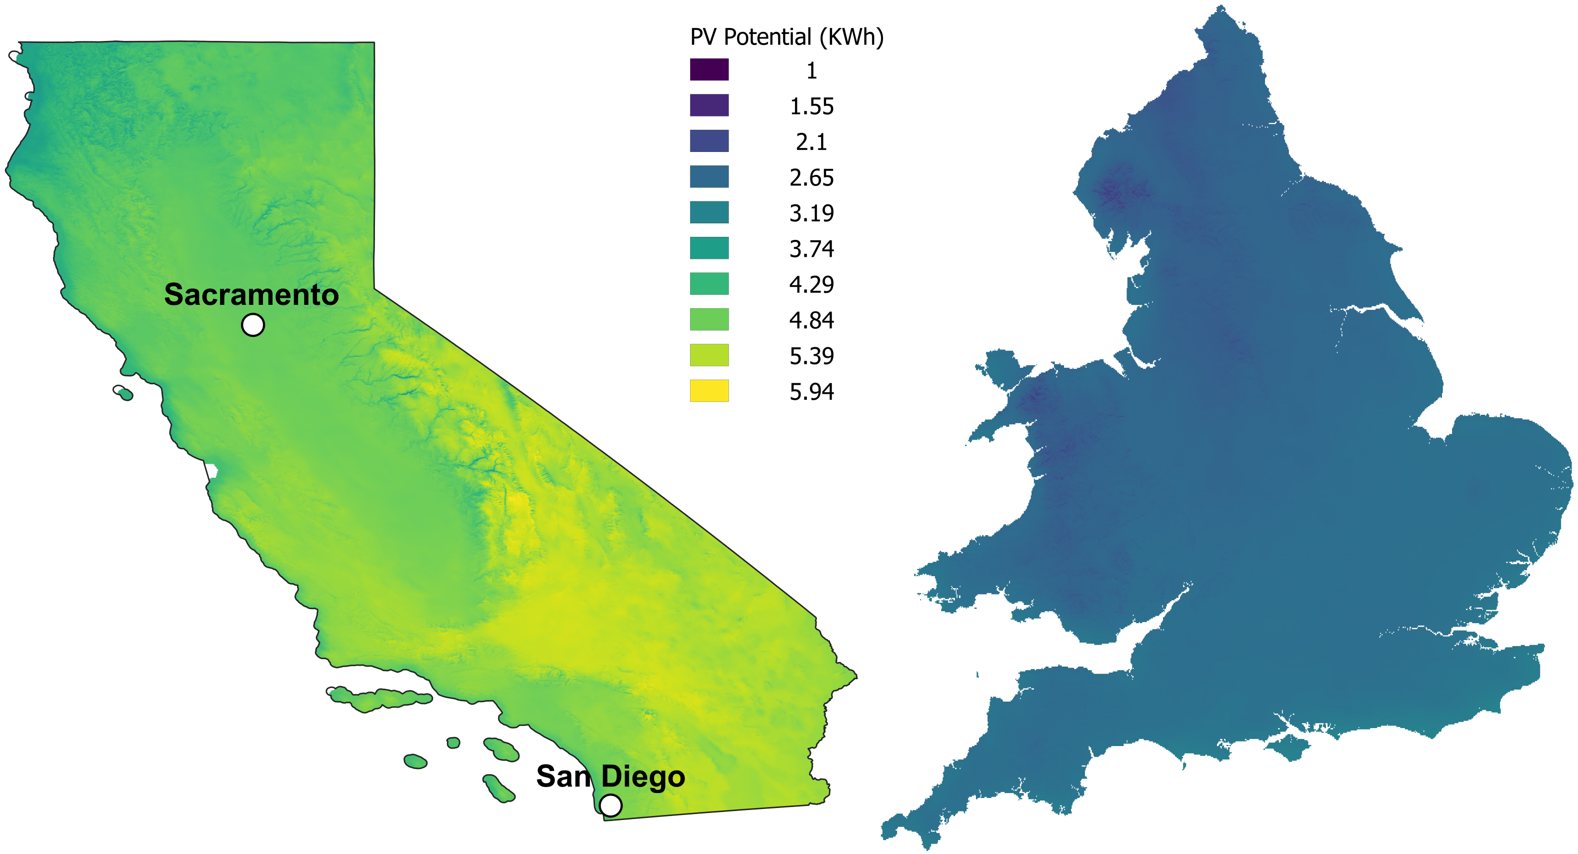

Figure 5.1: Average Annual PV potential: California, England and Wales.
Data: SolarGIS (2020); created in QGIS

5.3.1 The Rise of Residential PV in England and Wales

The increased demand for renewable energy generation over the past few decades has been met with increased efficiency of these new generation technologies which has led to ever increasing shares of national energy profiles coming from renewable energy generation. The increase in generation efficiency from renewable energy sources has been due to a combination of increased capacity as well as improvements to the technologies themselves (Xu et al. 2018). Although wind generation contributes the greatest share of renewable electricity, solar generation has also grown rapidly in recent years. Figure 5.2 shows the share of total electricity generated in the UK between 2000 and 2019 by all non-fossil fuel sources, wind, and solar.

Although solar represents a small share of the total and non-fossil generation, it is increasingly becoming a meaningful electricity source as PV costs decline and efficiency increases. Of course, much of the solar electricity generation is produced by large-scale commercial solar farms (OFGEM 2020*c*), residential solar installations have also seen widespread adoption and currently stand at just under 800,000 for the UK as a whole (OFGEM 2020*c*). In Figure 5.3 I have plotted the total number of residential solar installations in England and Wales and Figure 5.4 plots the total residential solar capacity over the same period. The first residential PV systems became active in 1995 in Southeast England, and residential PV was extremely rare until 2010, when systems became much more widespread. Figure 5.3 shows that from 2010 to 2015 there was a rapid deployment in residential PV systems and cumulative residential capacity. Initially the two values grew at roughly the same rate. This reflects the increase in the average capacity of newer systems relative to older systems as shown in Figure 5.5 below.

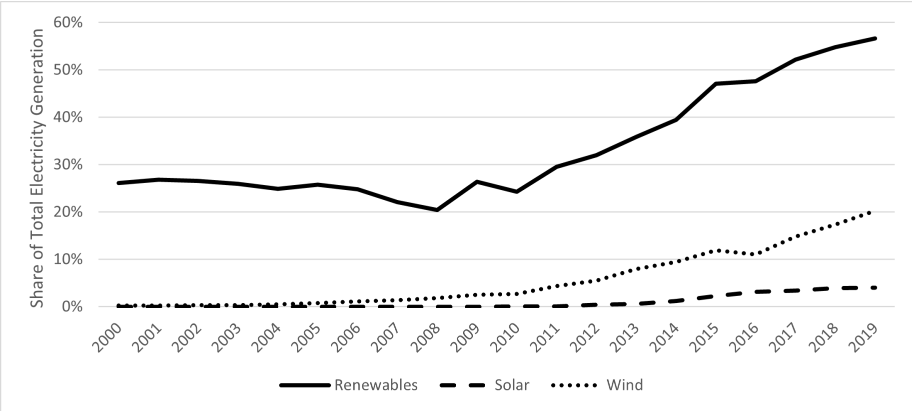

Figure 5.2: UK electricity generation by source, 2000-2020. Data: EMBER (2020)

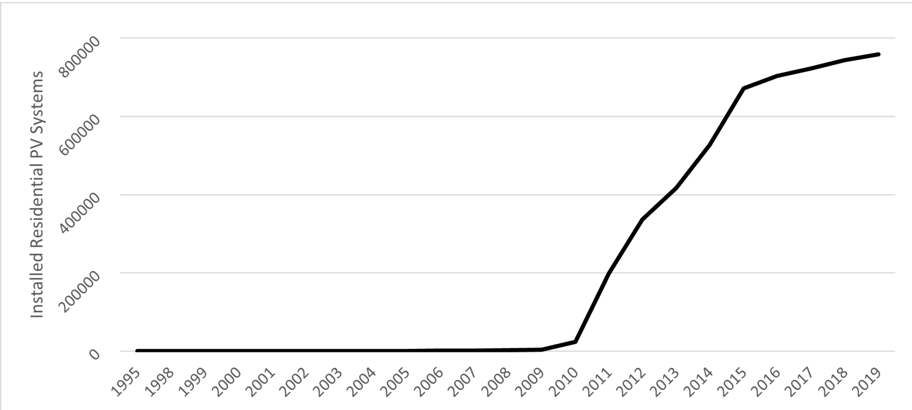

Figure 5.3: Cumulative Installed Residential PV systems in England and Wales: 1995-2020

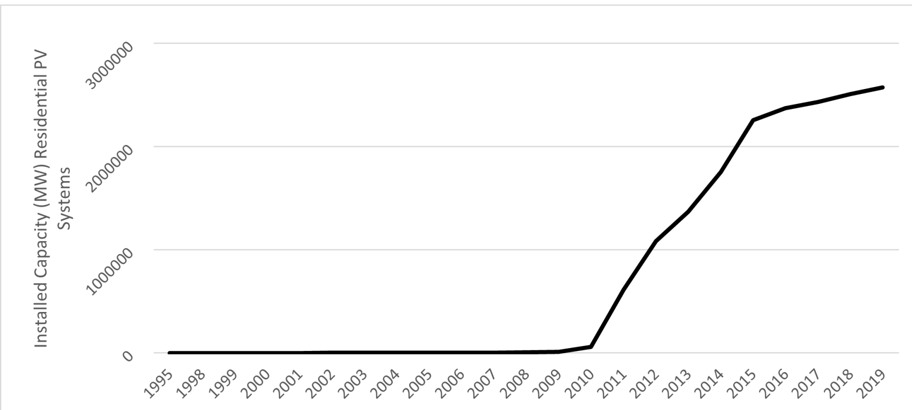

Figure 5.4: Cumulative Installed Residential PV systems in England and Wales: 1995-2020

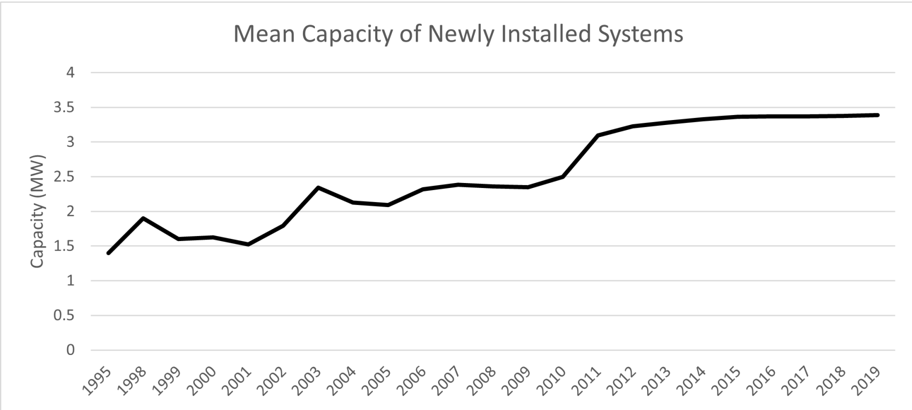

Figure 5.5: Mean Residential PV Generation Capacity

As discussed briefly in the previous section, the solar premium is expected to arise through a combination of energy savings, energy exports, and non-pecuniary preferences towards solar panels. Within the UK, the Feed-in Tariffs Scheme (FIT) was a government policy to promote the uptake of renewable micro-generation technologies including residential PV systems (OFGEM 2020*a*). The FIT Scheme mandated that electricity providers must pay generators for all electricity generated from a qualifying system, as well as any excess electricity exported to the grid. Within the context of this chapter, a generator would be any homeowner who installed residential solar PV and accepted into the scheme. The FIT scheme was opened to applicants on April 1, 2010 and closed to new applicants on April 1, 2019. The Feed-in Tariff was guaranteed for 25 years for solar PV, and the size of the tariff (per KWhr) was determined by the year of entry to the scheme and is constant for the 25 year period. The size of the tariff decreased each year for new entrants. From April 1 2019, the FIT was replaced by the Smart Export Guarantee, which guarantees a non-zero payment (per Kwh) for exported electricity only (OFGEM 2019).

Evidence from the literature shows that not only are solar PV systems capitalized into property values, but this solar premium will compensate for the costs of installation and potentially even lead to profit above the installation costs. However, this evidence has been gathered from regions with considerable larger PV generation potential than that of England and Wales. Despite the lower generation potential within the study region, the FIT scheme was quite generous. The subsequent analysis of this chapter performs a series of statistical analyses to determine if 1) residential PV systems are capitalized into property values in England and Wales, 2) if that capitalization will compensate for the installation costs despite the relative lack of electricity generation potential and, 3) how the capitalization and solar premium compare to those reported in the literature.

5.4 Data and Estimation Strategy

The research design applied to the data is as follows: Firstly, I estimate a baseline hedonic regression model using pooled OLS to estimate the capitalization of residential PV systems into property transaction prices in England and Wales. However, as there are limitations and potential biases associated with this model, I then take a repeat sales approach to generate a second set of results where capitalization is estimated by comparing the price between two neighboring transactions of the same property. Neighboring transactions refer to transactions of the same property which are sequentially neighbors - ie the fist and second transaction or second and third transactions of a given property. Lastly, I perform a propensity score matching analysis to generate a matched control group and the average treatment effect on the treatment group is reported. The goal of all three models is to estimate the average value added to a property arising from the installation of a residential PV system. I perform Moran's I test for spatial auto-correlation within Models 1 and 2, and the results are reported in Appendix A16.

5.4.1 Model 1: Baseline Hedonic Regression

Model 1 is the benchmark against which I compare the results of the subsequent models. This model is a pooled OLS model and applies a semi-log function which allows for the interpretation of the regression coefficients as a percentage of the property transaction price. (β_1) is therefore interpreted as the average price effect from the average sized PV system. Here, the sample analyzed includes property transactions which occur in postcodes where residential PV systems are present and include a mix of properties where the systems have been installed and neighboring properties which have not installed. Model 1 is represented in Equation 5.2 where year-by-quarter dummies (σ_t) and postcode dummies (γ_j) control for any unobserved factors across time and postcodes which may affect transaction prices. The time dummies will adjust for seasonal fluctuations within the housing market and the neighborhood dummies (postcodes) capture locational features such as crime rates, school quality, recreational opportunities etc.

$$ln(price)_{ijt} = \alpha + \beta(solar)_{it-1} + \sigma_t + \gamma_j + x'_{it}y + E_{it} \tag{5.2}$$

Where:

- α is the constant.
- $price_{ijt}$ is the price of property i in postcode j at quarter t.
- β reports the estimated value of the average sized residential PV system as a percentage of the transaction price.

- $Solar_{it-1}$ is a dummy indicator which takes the value of 1 if a property has an installed PV system at least one quarter prior to a sale, and is 0 otherwise.
- σ_t is the quarter and year in which the property was sold.
- γ_j is the postcode where property i is located.
- $x'_{it}y$ is a vector of property specific characteristics including: number of rooms, floor area, (m^2), number of fireplaces, property type, construction age band, and an indicator for whether the property will install a PV system in the future.
- E_{it} is the error term.

Hedonic regression models are the workhorse of many econometric analyses valuing property characteristics. Though it is a more simple analysis than the techniques applied subsequently, its main strength is that it can be used to compare the prices of all PV homes in the sample to all non-PV homes within the dataset. A key assumption of this model is that any unobserved property-specific characteristics are not correlated with PV installation. Any omitted variables which are correlated with the presence of a residential PV system, and affect house prices, may bias the estimated price effect from installing PV. Within this model, I attempt to reduce this bias by controlling for fixed effects at the postcode level. There is also potential for selection bias, or the possibility that the distribution of PV homes may differ systematically from the set of non-PV homes. To further address potential biases in the estimated price effect, I perform two subsequent analyses which address or reduce the selection and omitted variable biases.

5.4.2 Model 2: Repeat Sales Analysis

Model 2 takes the form of a Repeat Sales Analysis. This approach allows for a comparison of a given property's value before and after a residential PV system is installed. Here, I restrict the sample to only properties which have sold twice or more, and which at some point will install solar panels. I then construct a cross-section of neighboring transaction pairs and compare the average change in price between the two sales. I then generate a dummy indicator which is 1 when a residential PV system has been installed between two sequential transactions and 0 otherwise. This model estimates the average increase in transaction price between two sales of properties which install PV between sales. Again, if a property within the dataset is sold four times, these four transactions will appear as 3 pairs of neighboring transactions. Transaction one paired to transaction two, transaction two paired with transaction three and transaction three paired with transaction four. This captures the price effect of the installation. The model controls for time and geographic effects, as well as the characteristics of the transacted properties. The coefficient of this installation indicator β provides the estimated capitalization of a residential PV system. Again, year by quarter dummies σ_t and postcode dummies γ_j control for any unobserved factors across time and postcodes which may affect transaction prices of a property between two sales. Crucially, the model also controls for both years of the neighboring transactions, as well as the time between them measured in quarters. The equation of model one is shown below:

$$ln(\frac{Price_{ij(t+T)}}{Price_{ijt}}) = \beta\Delta(Solar_{i(t+T)}) + x'_{i(t+T)}\gamma_i + \sigma_t + +E_{t+T} \tag{5.3}$$

Where:

- $Price_{ij(t+T)}$ and $Price_{ijt}$ are neighboring transaction pairs of property i in postcode j which occur T quarters apart, and the first transaction occurs in quarter t.
- $\Delta(Solar_{i(t+T)})$ is a dummy indicator which is 1 if solar panels have been installed between sales of property i.
- β is the average effect of the average PV installation.
- $x'_{it}\gamma_j$ is a vector of property and geographic characteristics of property i.
- $E_{(t+T)}$ is the error term.

The repeat sales estimate (β) of the capitalization of solar panels in housing prices measures the additional premium of properties with residential PV systems installed between sales above and beyond the price appreciation of like properties with no solar installations. By interpreting β as the effect of PV installation on subsequent transaction price requires the assumption that any idiosyncratic price appreciation of the properties in the study is not associated with the installation of solar panels. This is likely controlled for by the fact that all properties in the study are restricted to those which eventually install residential PV systems as the indicator will be 0 if a PV system was previously installed or has not yet been installed. There is also some potential for treatment effect heterogeneity bias here, though this is partially addressed by controlling for within the neighborhood effects at the postcode level. I discuss this in more depth in 5.6 In the repeat sales sample, 2,541 properties have installed observations have installed solar panels prior to their first sale, 1,785 install between transactions, and 1,801 have not yet installed. These are English and Welsh property transactions cover the period 2008 through 2019, which eventually install residential PV.

This approach has a few advantages over Model 1. As the properties are compared to themselves, this reduces the potential for omitted variable bias, as any unobserved characteristics which are correlated both with PV installation and transaction prices are controlled for. By restricting the analysis to only properties which eventually install PV systems, this analysis addresses the potential of selection bias where PV homes may be of higher quality than non-PV homes. This is simply due to the fact that all properties eventually install solar. I am additionally able to control for other changes in property characteristics, though these are exceptionally rare within the data. For example, between transactions one property gains a room, and therefore its total area also increases, and another property removes a fireplace. Unfortunately, the data does not allow for a full set of characteristics to be tracked over time, but it is assumed that these are also relatively constant.

5.4.3 Model 3: Propensity Score Matching

The final model used to estimate the capitalization of residential solar PV into property transactions is a Propensity Score Matching approach which pairs property transactions where PV systems are installed prior to the transaction to transactions where no PV system is installed before the sale. The underlying goal of this approach is to create a treatment and control group, where one did not exist. In this way, I estimate the counterfactual of what a given property's transaction price would have been if it had not installed a solar panel. As discussed in the previous section, this pairing matches treated and controlled properties based on exact matches on the matching criteria. These criteria are: Floor Area, number of rooms, number of fireplaces, property type, tenure type, the age band of the property, the year and quarter of the transaction, energy efficiency rating, and the postcode where the property is located. A balancing test of these variables is presented in Table 5.3, though this table reports only the core variables as when sale year-quarter and postcode values are included the table will list thousands of variables. The balancing test compares the means of the matching variables between the treated and control groups – there is no statistically significant difference in the means of the two groups. The bias of the matching should remain below $|5\%|$ and this is the case for all variables. I also present a histogram of the propensity scores of the treated and untreated groups in Table 5.6.

By using Propensity Score Matching, the analysis allows for comparisons of the transaction prices of properties who do and do not install residential PV systems, but between properties which are equally likely, or nearly equally likely to install such systems – based on the characteristics discussed above. This model provides the Average Treatment effect for the Treated (ATT) which is the estimated increase in property value for properties which have installed a residential PV system. Applying this approach works in two stages. First, I use a logistic regression model to calculate all properties' propensity for being treated (installing solar PV systems). Following Rosenbaum & Rubin (1983), the propensity score is defined in Equation 5.4 below.

$$p(T) = Pr\{T = 1|S\} = E\{T|S\} \tag{5.4}$$

Where:

- $p(T)$ is the propensity to install a residential PV system.
- T is and indicator that a property has or has not installed a PV system.
- S is the vector of covariates which influence whether a property

has been treated.

The second stage of the analysis uses the estimated propensity scores to match treated and control properties. There are multiple matching techniques, but for this paper I have applied Kernel matching, Standard Nearest-Neighbor matching and Greedy Nearest-Neighbor matching. Kernel matching uses the calculated propensity score to match treated properties to a weighted mean of control properties. The weights assigned to the control properties are based on the distance between the propensity score of the control properties and the score of the treated property which they are being matched to. In this way, all control properties can potentially contribute to the weighted composite mean of the control cases. This improves the estimation power and efficiency of the model (Frisco et al. 2007, Morgan et al. 2017). Nearest-neighbor (NN) matching also uses the calculated propensity score to match a treated property to a control property (Frye & Bartlett 2017). This is achieved by measuring the distance between the propensity scores of each treated and control property and matching between those with the closest scores. Standard NN matching only allows a control property to be matched once, while greedy NN allows the control property to be matched to multiple treatment properties if it has the closest propensity score to that property.

The ATT is the average effect of treatment on the properties which install a residential PV system, whereas the results presented in Model 1 will be the average population-wide effect, Models 2 and 3 estimate the effects for only the group of properties which eventually install the solar systems. I perform the propensity score matching method under 2 alternative matching approaches - nearest neighbor matching and greedy nearest neighbor matching. This essentially acts as a robustness to test whether the results are sensitive to alternative means of generating

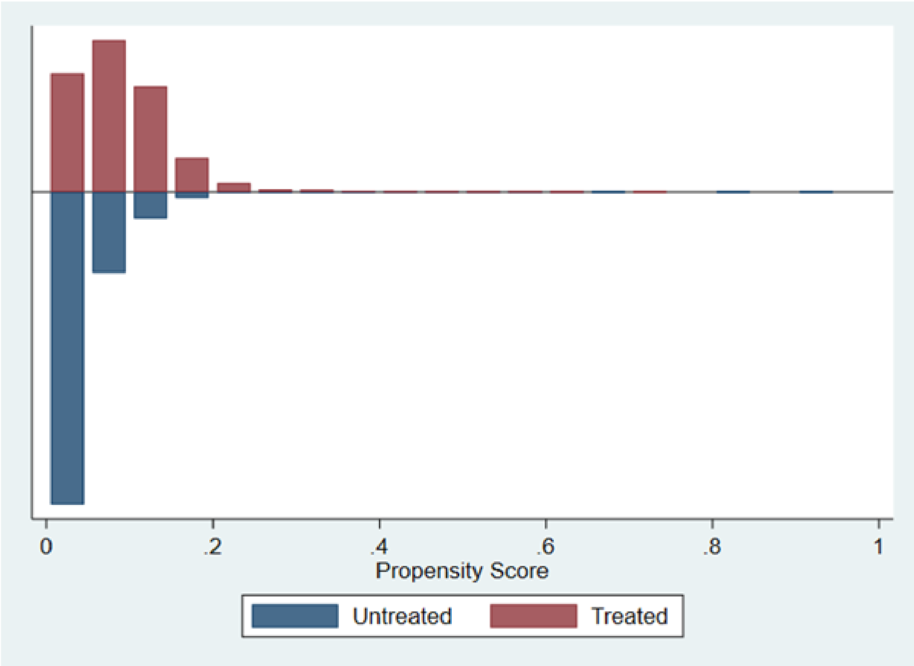

Figure 5.6: Treated and Untreated Propensity Scores

the control group. Nearest neighbor matching is an exact matching approach which allows a control property to be matched to multiple treated property if it is the best match. Under the greedy matching, controls can be paired only once.

5.4.4 Data and Methods

Properties and Transactions

Under each model outlined in the subsequent section, controlling for or matching on property characteristics is key to eliminating biases in the estimations of capitalization of residential PV systems. These property characteristics were obtained from the Energy Performance Certificate (EPC) registry for England and Wales and include detailed property characteristics relating to the energy efficiency of residential properties. All properties must have an EPC lodged with the registry when constructed, sold, let, at assessment for, and after installation of any

Table 5.3: Balancing Tests for Controls

Variable	Mean			t-test	
	Treated	Control	Bias (%)	t	p>t
Floor Area (m2)	99.208	98.741	1.3	0.98	0.328
Number of Rooms	4.9618	4.9693	-0.5	-0.49	0.622
Number of Fireplaces	0.11113	0.10946	0.4	0.37	0.715
	Property Type				
Detached	0.30115	0.30259	-0.3	-0.29	0.77
Semidetached	0.39702	0.39564	0.3	0.26	0.792
Terraced	0.27978	0.27829	0.3	0.31	0.756
Flat/Maisonette	0.02033	0.02154	-0.6	-0.79	0.431
New Build	0.06161	0.05994	0.7	0.65	0.515
	Tenure Type				
Owner-occupied	0.8559	0.84534	3.1	1.37	0.172
Rental (Private)	0.03544	0.03708	-0.7	-0.41	0.685
Rental (Social)	0.00892	0.00915	-0.3	-0.11	0.909
Unknown	0.09974	0.10843	-3.7	-1.31	0.189
	Construction Age Band				
1900-1929	0.12247	0.12173	0.2	0.21	0.831
1930-1949	0.28335	0.28761	-1	-0.88	0.379
1950-1966	0.19572	0.19744	-0.4	-0.41	0.685
1967-1975	0.07986	0.07819	0.6	0.58	0.564
1976-1982	0.06576	0.06593	-0.1	-0.06	0.948
1983-1990	0.04192	0.0418	0.1	0.05	0.957
1991-1995	0.06812	0.06771	0.2	0.15	0.881
1996-2002	0.02781	0.02677	0.6	0.59	0.553
2003-2006	0.01871	0.01589	2.4	2.02	0.044
2007 Onwards	0.0323	0.03432	-1	-1.05	0.295
	Sale Year				
1996	0.04376	0.04226	0.8	0.69	0.492
1997	0.04583	0.04319	1.4	1.2	0.231
1998	0.04837	0.0486	-0.1	-0.1	0.92
1999	0.05436	0.05257	0.8	0.74	0.46
2000	0.05447	0.05413	0.2	0.14	0.887
2001	0.05729	0.05873	-0.7	-0.57	0.566
2002	0.06121	0.06276	-0.7	-0.6	0.548
2003	0.05603	0.05395	0.9	0.85	0.397
2004	0.0558	0.05827	-1.1	-0.99	0.32
2005	0.04301	0.04365	-0.3	-0.29	0.772
2006	0.05522	0.0558	-0.3	-0.23	0.815
2007	0.0562	0.05718	-0.4	-0.39	0.693
2008	0.02332	0.02487	-1.1	-0.94	0.345
2009	0.01894	0.01814	0.5	0.56	0.578
2010	0.01704	0.01848	-0.9	-1.02	0.31
2011	0.0175	0.01768	-0.1	-0.12	0.903
2012	0.01912	0.01848	0.4	0.43	0.664
2013	0.03754	0.03691	0.3	0.31	0.755
2014	0.05257	0.05349	-0.4	-0.38	0.702
2015	0.03288	0.03161	0.6	0.67	0.504
2016	0.0258	0.02637	-0.3	-0.34	0.736
2017	0.02608	0.0262	-0.1	-0.07	0.946
2018	0.02706	0.02591	0.6	0.67	0.504
2019	0.02747	0.02591	0.8	0.9	0.369
2020	0.00685	0.00668	0.2	0.2	0.844

renewable energy system qualifying for the FIT or SEG schemes (EPC 2018). The registry contains duplicate certificates for properties which have met the lodgement requirements multiple times. This dataset was then linked with the transaction data from the Land Registry Price Paid data for England and Wales to create a dataset consisting of detailed property characteristics and transaction prices of these properties.

Linking these two datasets involved a series of complex matching on addresses from each dataset, and not all properties from the EPC registry were linked to a transaction within the Land Registry (LR) data. The headline results presented below represent matched properties which did not require editing the address fields of the two datasets – it is assumed that these are perfect matches and there are no incorrectly linked addresses here. In addition to the linking of the LR and EPC datasets, I then link geocoordinate data from Ordnance Survey's Address Base Plus (ABP) dataset. This again required a series of matching criteria and some properties failed to have corresponding matches across the dataset[1].

To create this dataset, I first restricted the EPC characteristics data to those located in a postcode with at least one installed residential PV system. This included approximately 665,000 properties within solar postcodes. 81,000 properties with installed solar panels and 584,000 properties without, after matching under the strictest criteria described above, roughly 21% of the properties were retained in the final matched dataset. Only 13% of properties with residential PV systems were retained while 22% of non-PV properties were retained. This is shown in Table 5.4.

This reduction is a function of 1) the fact that not all properties which

[1] Instructions for carrying out the multi-stage process of linking these datasets was graciously provided by (Chi et al. 2019).

Table 5.4: Matching

First Stage Perfect Matches

	PV	No PV	Within a PV Postcode
EPC	81,112	584299	665,411
EPC-LR-ABP	10,777	130,159	140,936
Proportion	0.132866	0.222761	0.211802931

appear in the EPC dataset will have been sold during the period and therefore do not appear in the LR dataset and 2) that the addresses reported in the two datasets differ substantially and to increase the matching rate requires the relaxation of several matching restrictions. Such a relaxation will increase the size of the dataset, but also lead to potentially incorrect matches, introducing bias into the results. The analysis of this chapter presents results only form the dataset where perfect linking was achieved. Even with the strict inclusion criteria, my dataset is considerably larger than those within the existing literature. To summarize, the base dataset for this chapter is restricted to 1) properties which appear in the EPC Registry 2) Properties located in a postcode where at some point over the study period a residential PV system is installed; 3) properties which have been sold over the study period and have had transaction(s) recorded in the Land Registry; and 4) Transactions below £10,000 and above £1,000,000 were dropped from the analysis following Sims et al. (2008). I then make further restrictions to this base dataset to suit the repeat sales and PSM approaches, which I discuss in their respective sections.

In Figure 5.7 below, I show the solar PV potential of England and Wales as well as the locations of the properties from the analysis which have installed residential PV systems. Figure 5.8 Is a heat map showing the geographical concentrations of the solar properties in the analysis. A key takeaway from these two figures is that the location of solar properties is less a feature of solar potential as it is a feature of the geography of the English and Welsh populations, coupled with the local stocks of residential properties. This is clearly shown in Figure 5.8 where the highest concentration of solar properties are near Manchester and Leeds rather than in along the southern cost where there is the largest generation potential, or in London where the largest share of residences are located. It is important to reiterate that this is not the distribution of all properties with residential solar installations in England and Wales, but rather the locations of such properties analyzed in the current study. Figure 5.9 below is a histogram of the transaction values of the properties in the current analysis, cutoff at £10,000 and £500,000 to avoid an exceptionally long tail.

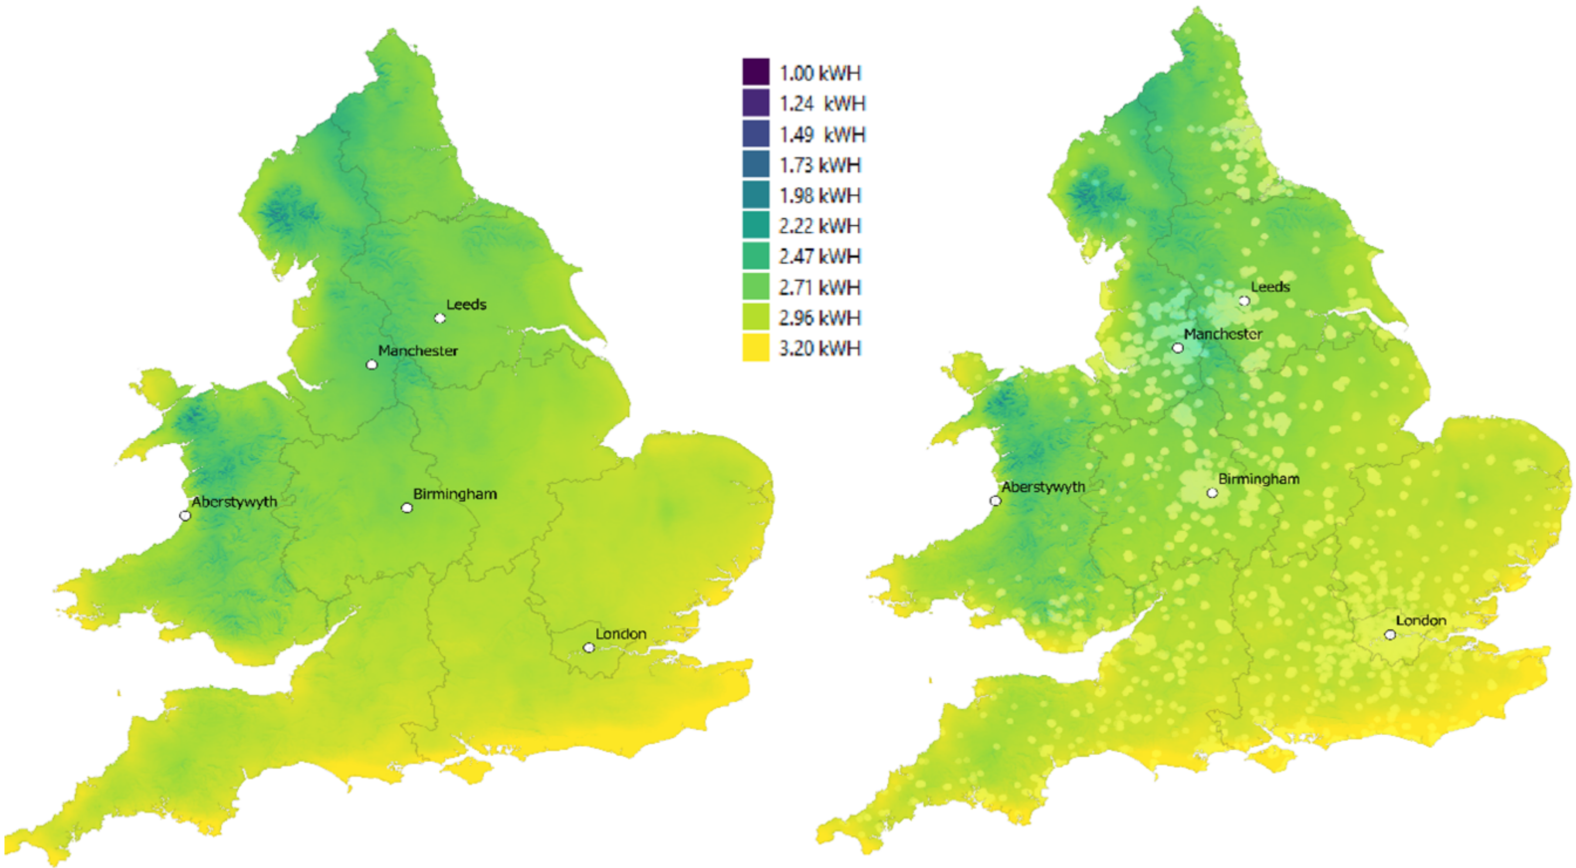

Figure 5.7: UK Solar Potential and Geographic Distribution of Properties in this Analysis

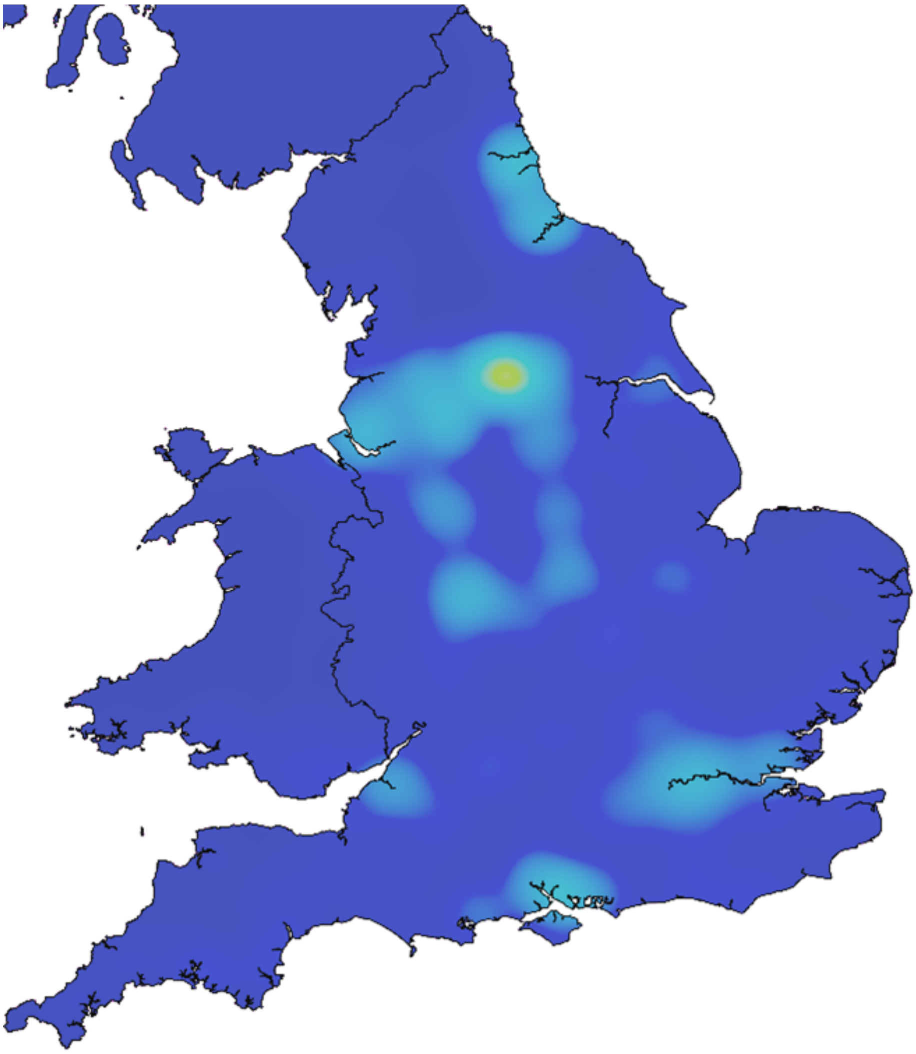

Figure 5.8: Heat Map of Properties with PV Systems

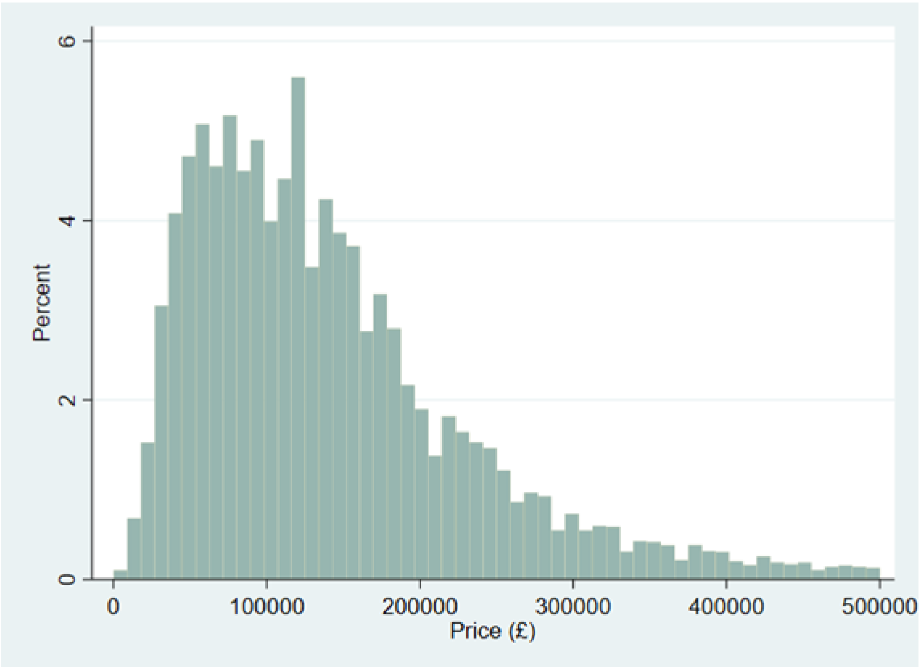

Figure 5.9: Histogram of Transaction Prices within Postcodes with Residential PV

5.5 Results

The following section presents a series of results from the analyses described in the previous sections of this chapter, and are structured as follows: Firstly, in Table 5.5, I present the results from the OLS regression coefficients of the benchmark analysis that form the starting point for the subsequent analyses. Next, in Table 5.6 I present the results from Model 2, a repeat sales analysis applying a pooled cross-section approach which reports the average capitalization of residential solar installations by comparing neighboring transactions of the same properties. Lastly, Table 5.7 reports the results from Model 3, which applied three propensity score matching approaches to determine the extent to which solar panels are capitalized into property transactions.

5.5.1 Hedonic Regression Analysis

Table 5.5 Column (1) reports that in the baseline regression, where the only independent variable is the dummy indicator for solar installations pre-sale, there is a roughly 9.3% increase in price relative to properties which have not installed solar panels prior to being sold. In Column (2), where additional independent variables are included in the analysis, the average price effect of solar installations drops from about 9.3% to 7.8% - still a substantial and statistically significant increase in the value of a property. While both specifications of this model show positive and statistically significant results, the inclusion of the additional characteristics does reduce the estimated effect.

Internal property-specific characteristics which are associated with statistically significant price effects include the total floor area, number of rooms, and the number of fireplaces. For each additional square meter of usable floor area, there is an associated price increase of 0.39%, or for each additional 10 square meters an increase of 3.91%. Each additional room increases price by roughly 3.16%, and each additional fireplace by 12.80%. Additionally, the property type itself is highly correlated with price effects which range from average prices being 12.44% lower for flats to 42.5% higher for detached homes – there is also a 9.4% price premium for newly built homes. Additional controls and associated coefficients are reported in Appendix **??**.

Under this model there is evidence of a solar property premium with an average capitalization of 7.8%. Based on the mean sales price of properties sold without solar panels (£134,531) this translates to an increase of £9,417.18. This result can be interpreted to imply that by installing the average-sized PV system, a property will increase by just under £10,000. This is well above the cost of installing the average sized PV system. Following Wee (2016) I take the 7.8% price increase from solar panels to be an upper bound estimate of the extent that solar panels are capitalized into house prices, and use it as a benchmark to compare the results of the subsequent two models presented in Sections 5.5.2 and 5.5.3.

As discussed in the previous Section (5.4) the main strength of the Hedonic Regression approach is that it makes use of the full dataset available. However, as the list of property characteristics is limited, there may be some omitted covariates which are correlated with both price and the installation of PV. In fact, there is evidence that the inclusion of additional covariates (2) reduces the coefficient relative to the specification with none (1) there is potential that unobserved characteristics may similarly influence the results were they to be included in the model. Additionally, there may be a systemic difference in the distribution of PV and non-pv homes leading to selection bias. I mitigate these potential sources of bias in the subsequent models of the analysis.

Table 5.5: Model 1: Hedonic Regression Analysis

Ln Price	(1)	(2)
Solar	0.09321***	0.07786***
(RSE)	(0.00421)	(0.00816)
Total Floor Area (m2)		0.00391***
(RSE)		(0.00003)
Number Rooms		0.03161***
(RSE)		(0.00072)
Number Fireplaces		0.1280***
(RSE)		(0.00170)
Flat		-0.1244***
(RSE)		(0.00447)
Terraced		0.24408***
(RSE)		(0.03363)
Semidetached		0.28921***
(RSE)		(0.00211)
Detached		0.4252***
(RSE)		(0.00225)
New Build		0.09440***
(RSE)		(0.00379)
Observations	416,508	416,508
Sales with Solar	17,407	17,407
R-squared	.2708	.4405
Property Controls		Yes
Geogrpahic Controls		Yes
Time Controls		Yes

*** p<0.001
**p<0.01
*p<0.05
RSE: Robust Standard Errors Clustered at the Outcode
Geographic Controls: Outcode by Year, Photovoltaic Potential at the Postcode Level.
Time Controls: Sale Year, Sale Month

5.5.2 Cross-Sectional Repeat Sales Analysis

Here I present the capitalization of residential solar installations under the assumptions of Model 2, which applies a Repeat Sales Analysis using a cross-section of sequential transaction pairs, with the key independent variable being a dummy indicator for a property having installed solar panels between two sequential transactions. This analysis is restricted to only properties which eventually install solar panels, and which are sold repeatedly during the study period. The restriction criteria significantly reduce the number of observations in the repeat sales analysis. The benchmark analysis of Model 1 contains over 416,000 observations whilst the repeat sales analysis contains only 6,665 observations. The results of Model 2 are presented in Table 5.6 below and are again split into two columns which present the effects on the change in the natural log of price between two sales where column (1) reports results with no controls and Column (2) reports results with a variety of property-specific, geographic and time controls.

When performing the analysis with no controls I find a rather large effect of 10.13%. However when controlling for the various internal and external property-specific characteristics as well as time and geographic features, the installation of solar panels between two transactions of the same property is estimated to increase the value of a property by an average of 2.7%. This is substantially lower than the estimate presented under the simple regression of Model 1. There is also a statistically significant effect associated with the time between sales, implying that the longer the time between neighboring transactions increases, so does the change in price – this is essentially capturing house price inflation. The coefficients of the internal property characteristics are broadly similar to those reported under Model 1.

Under Model 2 I also find evidence of a solar panel premium with an average capitalization of 2.7%. Based on the mean transaction price for properties before the solar panels are installed (£172,544) this translates to an average solar property premium of £4,651.79. This estimated effect is lower than that found under Model 1. However, the main advantages of the repeat sales analysis relative to Model 1 is the mitigation of the omitted variable bias by arising from comparing properties to themselves. As the repeat sales model functionally takes the form of a property fixed-effects model any unobserved characteristics are inherently controlled for. By restricting the repeat sales analysis to only properties which eventually install PV, the problem of selection bias is also addressed. However, mitigating these potential biases comes at the cost of reducing the dataset substantially, and generalizing the estimated effects outside of the small subset of properties becomes problematic as there may be other features which lead to these properties being sold multiple times.

Table 5.6: Model 2: Cross-Sectional Repeat Sales Analysis

Ln Δ Price	(1)	(2)
Solar	0.10134*	0.02696***
(RSE)	(0.04902)	(0.007987)
Total Floor Area (m2)		0.00334***
(RSE)		(0.00054)
Number Rooms		0.04786***
(RSE)		(0.011751)
Number Fireplaces		0.09016***
(RSE)		(0.03164)
Flat		0.01410***
(RSE)		(0.00336)
Terraced		0.02844***
(RSE)		(0.00366)
Semidetached		0.06071***
(RSE)		(0.00986)
Detached		0.84721*
(RSE)		(0.40976)
New Build at First Sale		-0.04391***
(RSE)		(0.004252)
Years Between Sale		0.12717***
(RSE)		(0.002422)
Observations	6,665	6,665
Sales with Solar	2,089	2,089
R-squared	0.2933	.6007
Property Controls		Yes
Geogrpahic Controls		Yes
Time Controls		Yes
Geography*Time Controls		Yes

*** $p<0.001$
**$p<0.01$
*$p<0.05$
RSE: Robust Standard Errors Clustered at the Outcode
Property Controls: Room count, Useable Floor Area (m^2), Fireplace Count, Newbuild Property at first sale, Property type: Flat, Terraced, Semi-detached, Detached, Construction Age Band.
Geographic Controls: Outcode by Year, Photovoltaic Potential at the Postcode Level
Time Controls: Sale Year, Sale Month, Time Between Sales, Interaction of First Sale Year and Second Sale Year.

5.5.3 Propensity Score Matching

In Table 5.7 I present the results of Model 3. Here, I generated a series of control property transactions (no solar installations before a sale) which are then paired with treated properties (solar installations prior to a sale) to estimate the solar panel premium and the extent to which they are capitalized into housing transactions. Columns (1), (2), and (3) present the average treatment effect on the treated group, or the average effect of being sold with installed solar panels in terms of the natural log of the transaction price. Each column presents results using a different set of matching criteria, with (1) being the least strict – only matching on property characteristics, while (2) applies slightly stricter matching criteria involving time effects, and lastly (3) which requires exact matches between treated and controls on property characteristics, geographic and time controls. I then perform two alternative matching techniques which are listed in Columns (4) and (5) Where I apply a greedy nearest neighbor matching, and standard nearest neighbor matching.

When the matching criteria are at their loosest, I find an average increase in in the treated group of 16.32%. Though this would imply a substantial gain from installing solar panels, the suitability of matching between the treated properties to appropriate controls is questionable. Though it ensures that treated and control properties are structurally the same, it cannot account for general inflation, nor other price effects which may differ over time and space. Under the assumptions of Column (2), where properties are matched exactly on property characteristics as well as the time of the transaction the effect drops to a 10.62% capitalization. Although this estimate is more robust than that presented in (1), it still allows for a treatment and control matching which is geographically unrestricted. Finally in Column (3) where treatment and control matching require exact matches on property characteristics, time and geographic controls I find an estimated capitalization of 3.46% which is statistically significant at the $p<.001$ level. Based on the mean transaction price of the matched control group (£190,335.76) this translates to an average treatment effect of £6,578.00.

Table 5.7: Model 3: Propensity Score Matching Analysis

ATT: Ln Price	(1)	(2)	(3)	(4)	(5)	(6)
Solar Installed Prior to Sale	0.1632***	0.1062***	0.0346***	0.0360***	0.0458***	0.0756***
(SE)	(0.0187)	(0.0131)	(0.0082)	(0.0087)	(0.0095)	(0.0121)
T-Stat	8.73	8.12	4.24	4.12	4.82	6.25
Observations	416,061	387,366	176,916	175,555	196,161	288,437
Treated	4,253	4,253	4,250	4,270	4,270	4,270
Untreated	411,808	391,619	172,666	171,285	191,891	284,167
Property Characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Time Controls		Yes	Yes	Yes	Yes	Yes
Geographic Controls			Yes	Yes	Yes	Yes
Within same postcode			Yes	Yes	Yes	
Within same outcode						Yes

*** $p<0.001$
**$p<0.01$
*$p<0.05$
Bootstrapped Standard Errors

Matching Criteria
Property Characteristics: Room Count, Useable Floor Area (m^2), Fireplace Count, Newbuild Property, Property Type: Flat, Terraced, Semi-detached, Detached, Construction Age Band, Postcode.
Geographic Controls: Parliamentary Constituency, Photovoltaic Potential at the Postcode Level, Elevation, Slope and Aspect.
Time Controls: Sale Year, Sale Quarter

Column (4) shows the estimated ATT when applying a nearest neighbor matching, and (5) the estimated ATT when applying a greedy nearest neighbor matching. The effect remains statistically significant, though under the standard nearest neighbor match there is an increase of about 0.14% above the kernel matching result. When I make use of the greedy nearest neighbor approach, the coefficient increases to about 4.6%. Lastly in Column (7) I relax the matching criteria to allow for matches at the outcode level rather than the postcode level. Outcodes are the first one to three characters in a postcode. For example, with the postcode G1 1XQ, the outcode is G1. Under this relaxation, the ATT increases to about 7.6%.

I have generated results using the Kernel, Standard and Greedy Nearest Neighbor matching techniques to serve as robustness of the model to alternative matching approaches. The Kernel matching approach to PSM essentially allows all control properties to contribute to the estimated effect of the ATT. This has the advantage of not discarding unmatched control properties from the analysis, while the properties used to generate the composite control are weighted by their propensity score, meaning closer matches are weighted higher. Under this matching approach, the estimation power and efficiency of the model are improved through maintaining a larger sample size. The Standard Nearest Neighbor matching simply pairs a treatment to the control with the closest propensity score, and once a control property has been paired it cannot be paired again. The Greedy Nearest Neighbor also pairs a treated property to the control with the closest propensity score, but if a control is the nearest neighbor to multiple treatment properties it is allowed to match them all (Frye & Bartlett 2017).

The PSM approach of Model 3 mitigates the selection bias relative to the benchmark hedonic regression of Model 1, because it creates a set of control properties based on their propensity to install solar PV. This approach ensures comparability between the two groups - PV properties and non-PV properties. Though it addresses selection bias, this model also has potential to suffer from the omitted variable bias due to the incomplete nature of the property characteristics data available, much like the analysis of Model 1. Although the repeat sales analysis of Model 2 addresses both the omitted variable bias and the selection bias, this comes at the cost of a significant reduction in the number of observations within the sample, and the generalizable of the results to properties which do not sell repeatedly. The PSM approach of Model 3 allows for an estimated price effect using a much larger sample than the repeat sales model.

5.6 Discussion

The results presented in the previous section provide evidence that there is a solar property premium within England and Wales, and evidence that this premium is capitalized into property transaction values. All three models report statistically significant and positive increases in home values when solar panels are present. A summary of the headline results of each model are presented in Table 5.8 below. The variation in the estimated capitalization is to be expected as the models apply slightly different statistical approaches to different datasets but all three models point to the same conclusion – that the transaction prices for properties with solar panels is higher than properties without the panels. The results of Models 1 and 3 suggest that the average value of the capitalization exceeds the average cost of a residential solar installations in the UK which ranges from £4,800 for the average installation to £5,600 in 2020 (EST 2020). Model 2 shows that although there is a statistically significant increase in price attributable to the installation of a PV system, this is not quite large enough to fully recover the average cost of installation.

Table 5.8: Solar Capitalization Estimates: All Models

Model	Capitalization (%)	£ Value
1	7.786%	9,417.18
2	2.696%	4,651.79
3	3.456%	6,578.00

The results of all three models show that residential PV systems are capitalized into housing transactions in England and Wales, and PV homes enjoy a price premium over non-PV homes. However, the models do not agree that the installation costs, based on the UK average residential PV system will necessarily be fully recovered through the transaction price of a home. The results of Models 1 and 3 show that the solar premium is in excess of the UK average installation cost. Using the installation cost of the average PV system provided by the Energy Savings Trust, this ranges from £978 to £1,778 for Model 3. For model 1, the premium in excess of the installation cost ranges from £3,817.18 to £4,617. Under the repeat sales analysis of Model 2 the estimated solar premium is not enough to fully recover the costs of installing an average sized PV system. Here, there is a loss of between £949 and £149.

There are four likely contributors to the values of these systems which explain the capitalized value of residential PV and the increase in transaction prices for solar properties. The first is simply the energy savings from the electricity generated by the PV system itself. However, under the highest estimated savings scenario from the Energy Savings Trust (EST 2020), this would amount to a yearly savings of only £240 and would require 20 years of energy savings to break even with the cost of the installation. Of course, the actual value of savings is dependent on the cost of electricity external to the PV system. If the price of electricity imported to the grid increases, the value of the energy savings will also increase - and if the price of electricity declines the opposite is true. The second contributor to the solar premium are subsidies, which may both reduce the cost of installation as well as increase the value of generated electricity.

These subsidies have not been stable across the period of this analysis, and attributing value to any one is beyond the scope of this paper. However, subsidies range from grants which would essentially cover half the installation cost of solar panels, to no interest loans, as well as a guaranteed export rate for unused electricity and a guaranteed payment for each unit of electricity generated. The variety of subsidies which also varied in size and duration may partly explain the differences in the capitalized value of residential solar panels across the three models due to a different set of transactions being analyzed under each which could lead to differences in subsidies incorporated into the valuations. The size and scope of any subsidies may be crucial.

As mentioned previously, Dröes & Koster (2021) find a decrease of 2.6% in the value of properties sited within 1km of commercial-scale solar farms. This is a substantial difference relative to the findings of this chapter and its analysis of residential solar's impacts on house prices. There are two potential avenues which may explain this - firstly The disparity between impacts arising from commercial-scale and residential solar is likely due to the presence of subsidy schemes for residential solar, and the lack of any compensation mechanism for solar farms. Secondly, it may be the difference between a solar home vs a home in a solar neighborhood. If there is a lack of community or recreational amenities associated with commercial scale solar, there would be no compensation mechanism which may explain the positive impacts on house prices from large-scale windfarms reported in the previous two chapters. However, testing this is beyond the scope of this thesis.

Black (2004) found that with the early subsidies provided by the state of California as well as energy savings, the after-tax return on investment for a 5kw solar PV system would outperform the stock market. Lastly, the solar premium will reflect a preference towards solar energy installations by prospective home buyers. It is possible that this preference may not be pecuniary in nature, but rather reflect a preference towards green energy generation. This preference for reducing carbon emissions may be a truly non-market aspect of the solar capitalization found in the current analysis (Dastrup et al. 2012).

5.6.1 Context from the Literature

The results presented in this chapter complement the current literature on the capitalization of a solar premium into property values. All papers in the literature have found that residential solar premiums are either equal to or in excess of the costs of installation, and the results presented here are in agreement and add the United Kingdom to the literature. In addition, the estimated price effects of the analyses presented in the previous section which range from 3.5% to 7.8% capitalizations are well within the ranges reported in the literature which ranges from 2.3% by Ma et al. (2016) to 17% by Qiu et al. (2017).

It is difficult to accurately compare the actual value of these effects in currency terms due to both differing currencies across studies as well as inflation of those currencies based on the periods of their analysis.

A very rough, back of the napkin estimate[2] of the currency value would lead to the capitalization of the results of this analysis to be valued at roughly \$6,311.53 (Model 2), \$8,925 (Model 3) and \$12,777 (Model1)– the lowest, second lowest, and third lowest values within the literature respectively. Table 5.9 below presents the results from this chapter alongside the findings of the literature. Although the estimated effects are smaller than those within the literature regarding the value of the capitalization, as a share of the value of the properties themselves, these results are larger than those found by Dastrup et al. (2012), Hoen et al. (2013) and Ma et al. (2016).

Compared to the estimated treatment effects reported by Dastrup et al. (2012) and Hoen et al. (2013) the 7.8% price increase of Model 1 is quite high, though it should be noted that the property values within California are considerably higher than those of this study, and ultimately the value of the increase is considerably higher in both of these papers. Lan et al. (2020) do not report their results as a percent of property values, but they do find that installation costs are fully recovered or lead to profit in their analysis. The repeat sales analysis of this chapter is finds an effect closest to that of Ma et al. (2016) (2.3%), but it is also quite similar to the effect reported by Hoen et al. (2013) (3.6%) and Dastrup et al. (2012) (3.5%). Though again, the cash value of the solar premium is higher in California than England and Wales. The estimated capitalization of 2.7% is smaller than that of Wee (2016) who found a 5% increase, and a substantially higher solar premium. Under the propensity score matching method of Model 3, the results are again smaller than the comparable analyses within the literature. Qiu et al. (2017) find an increase of 15-17% in property values due to the presence of a

[2]This rough conversion is calculated by multiplying the estimated solar premium in £ of each model and multiplying these by the exchange rate (£1 = \$1.3568) at the time of writing.

residential PV system, which is the largest yet reported within the literature on residential PV capitalization. It should be noted the sample of transactions within their analysis was exceptionally small, covering a period of only 3 months.

The results presented within this chapter show that there is some sensitivity regarding the model applied to the data, as I find a much higher capitalization under the analysis of Model 1, than Models 2 or 3. Each model is also sensitive to the restrictions to the data as well, as the coefficients change substantially depending on the controls included within the models. This has also been the case with the published literature and therefore it is not surprising that similar sensitivities were found here.

Table 5.9: Results of Models 1-3 Compared to Headline Results From the Literature

Author(s)	Hedonic Regression	Repeat Sales	Control-Matching	Solar Premium	Recover Installation Costs	Profit Over Installation Costs
Dastrup et al.	Y	Y	N	3.50%	Y	Y
Hoen et al.	Y	N	N	3.60%	Y	-
Ma et al.	Y	Y	N	2.3-3.2%	Y	-
Wee	Y	Y	N	5%	Y	-
Qiu et al.	N	N	Y	15-17%	Y	Y
Lan et al.	Y	N	Y	$21,403	Y	Y
Model 1	Y	N	N	7.80%	Y	Y
Model 2	N	Y	N	2.70%	N	N
Model 3	N	N	Y	3.50%	Y	Y

5.7 Conclusion

The analyses within this chapter make several contributions to the relatively small, but growing literature in residential PV capitalization. First, this is the first to estimate the solar property premium and the capitalization of residential PV into homes located in England and Wales. Second, this is the first analysis to include Propensity Score Matching techniques, and compare these findings to both a Hedonic Regression and Repeat Sales Model. Third, the analyses of this chapter make use of the largest dataset to date, containing considerably more solar properties than any paper in the literature.

Under all three models, I find evidence that within England and Wales, residential PV systems are capitalized into house prices and there is a solar premium. This is in line with the findings reported in the literature, though there is some evidence under Model 2 that the premium is not large enough to fully recover the cost of installing a residential PV system. The findings of this chapter have mixed implications for policymakers. The findings of Models 1 and 3 show that the premium is large enough to compensate for and profit from the installation of residential PV. These results would suggest that policymakers seeking to encourage adoption of these systems should make an effort to inform homeowners that the installation costs of residential PV can be recouped through the sale of the property. This may encourage additional installation of residential PV.

As there is potential that the residential solar premium arises, at least partly, from government subsidies aimed at increasing adoption, it may be important to preserve the generosity of such subsidies. This is of particular importance within the UK where the FIT scheme has just been replaced with the less generous SEG. In fact, based solely on the findings of Model 2, it may be necessary to increase the size of subsidy available to homeowners who install PV systems. These results imply that at current levels, homeowners may not be able to fully recover the cost of installation and therefore are potentially more hesitant to purchase a PV system.

Of course, the analyses here are not without their limitations. This analysis does not include every property in England and Wales which has installed a residential PV system. There are nearly 800,000 properties which have such a system, but only about 80,000 of these properties appear in the EPC registry, and of those only 10,777 are matched with transactions from the Land Registry Price Paid data. Although this is by far the largest sample of properties to be analyzed within the literature on residential PV capitalization, it does reflect an analysis of only a small sample of the total solar properties. Although the results are statistically significant across the models applied, the use of a dataset containing all transactions of properties with solar panels would lend increased accuracy to the estimation and ensure that there is no selection bias present within the sample. Future research could include a larger dataset which could extend the analysis to cover all or even a majority of these properties.

An additional limitation of the underlying data are the known intrinsic property characteristics. Although this analysis makes use of more detailed characteristics than other studies in the literature, these are limited to those available within the Energy Performance Registry. This data provides invaluable information about the properties in the analysis, but these are mostly related to the energy performance of a given property. Other research makes use of characteristics such as the number of bedrooms, bathrooms, whereas the data available for this work simply lists the number of rooms with no distinction between the room types.

These characteristics have been shown to be important factors in property valuations and therefore a future analysis may be able to apply even more detailed property information within controls or matching variables. In fact, perhaps the most important missing information is the size of a PV system at a given property. Although the EPC registry reports the area of roof covered by solar panels, it does not give an indication of the generation capacity of an installed system – and this is likely to largely drive the value of any capitalization into any transaction price. Larger systems will generate more electricity and may also receive larger subsidies which will effect the value added by the system itself. There is also the matter of the subsidies for residential solar system installation. These were not stable over time, neither in regards to their total value, or what properties might be eligible for installing the panels. Accounting for the impacts of the variety of subsidy programs which were opened, modified, or closed over the period is unfortunately not possible due to data availability constraints.

Lastly, and perhaps the most crucial limitation relates to the timing of an installation. The installation date is taken to be the date an entry is logged in the EPC registry where the presence of a PV system is indicated. For some properties, an earlier record where no PV system is installed is followed by a subsequent record indicating that the property has installed a PV system. However, for other properties there is only an EPC record indicating that a system has been installed with no previous records prior to the installation. The analysis assumes that the EPC inspection date is equivalent to the date of a completed installation

– EPC does not record the presence of solar panels until the installation has been completed and many solar installers either qualify as domestic energy assessors or partner with a qualified energy assessor who may assess the property at the time of installation.

This assumption will have no impact on the results of Model 1 as it simply tests for a price difference between properties which install solar panels eventually and those that never do. However, it could have an impact on the results of Model 2 and which use the time of installation to determine the time of treatment which is a key component of of the model. The EPC registry did not exist before 2008, and any existing PV system which had been installed prior to the FIT could be grandfathered in to the scheme. It is therefore possible that a property had installed a PV system prior to the launch of the FIT may not have an accurate installation date within the analysis - the installation date would be recorded as a date later than the true installation date. This could impact the estimated effects of the Repeat Sales Analysis of Model 2. This is because there is potential that a property had already installed PV, so there is no installation between two sequential transactions. In this way, the estimated effect would be reduced as there would be no change in the presence of solar panels between two sales. The PSM analysis of Model 3 could also be impacted, though to a lesser extent - there would simply be fewer early transactions of PV properties matched to non-pv control properties. However, this issue may be partly mitigated by the fact that the installation data provided by OFGEM, and displayed in 5.3 shows that a majority of these installations occur from 2010, onward - after the EPC registry was created.

This chapter estimated the capitalization of residential solar PV systems into house prices of properties in England and Wales under three different statistical models. The first model performed a basic hedonic regression, the second a repeat sales analysis and the third a propensity score matching analysis. All three models found evidence that there is a solar property premium and that installing a residential PV system and this ranges from a price increase of 2.7% to 7.8%, and under two models this is enough to recover installation costs as well as earn a profit from installing the average sized residential PV system. The analyses performed in this chapter updates the literature to include evidence of a solar property premium even in regions with relatively low PV potential, being the first analysis of its kind in the United Kingdom. In addition, this chapter applies an analysis to the largest dataset to date, and is the first to apply a propensity score matching method to measure the value of the solar premium. Therefore this analysis updates the regionality of the literate as well as the models applied to generate estimates of a solar premium. It provides context to the range of results found within the literature as the estimated effects are sensitive to the models and underlying data restrictions, though it does find consistently positive and significant price effects.

Chapter 6

Conclusions

6.1 Introduction

The growing need to address anthropogenic climate change has led to the United Kingdom making international commitments to reduce greenhouse gas emissions. To achieve these targets, a number of domestic policies have been implemented to support the transition to a low-carbon economy. This thesis has examined the externalities arising within the housing market from the implementation of these policies, specifically those which support the generation of electricity from commercial wind and domestic solar sources. In this final chapter of the thesis, I will summarize the contributions of each chapter. This will include the empirical findings, the contributions to the literature, and future avenues of research which can build upon the work undertaken within this thesis.

6.2 Chapter 1

Chapter 1 provides an overview of the current climate crisis and the international agreements which have been drafted to address it. I then outline the resulting legally binding emissions reductions targets made by the UK government to meet its commitments under the Paris Agreement. The discussion then turns to domestic policies in the electricity sector aimed at supporting the transition away from carbon based electricity. Policies aimed at supporting large-scale commercial low-carbon electricity generation include the now defunct Non-Fossil Fuel Obligation, and its replacement the Renewables Obligation.

The UK government also implemented the Feed-in Tariff Scheme, as well as its replacement the Smart Export Guarantee. The FIT and SEG support the development of small-scale renewable electricity generation such as residential solar PV systems. These policies, while very successful in supporting these technologies, led to externalities in the housing market by altering the characteristics of properties. The presence of an ever-increasing number of large wind turbines and residential PV systems have led to changes in the characteristics of homes located near to windfarms, and homes which have installed PV systems. The chapter then provided an overview of amenities and disamenities in the housing market, as well as the tools generally applied to evaluate them.

6.3 Chapter 2

Building on the context provided in Chapter 1, this chapter provided a detailed review of the literature classifying windfarms and environmental amenity or disamenity. This is the research area in which the empirical analyses of Chapters 2 and 3 are situated. It discussed the current findings presented within the peer reviewed research in this area, as well as the disagreement regarding whether wind energy development will affect house prices. Most papers find no statistically significant effect on house prices arising from wind turbine proximity or visibility, though a large minority find statistically significant negative price effects - and one paper finds positive impacts. The discussion then goes on to describe the data and analytical approaches taken within the literature, as well as the strengths and weaknesses of these approaches. As such, this chapter serves to contextualize the research topic of Chapters 3 and 4, as well as the analytical methods taken.

6.4 Chapters 3 and 4

Chapters 3 and 4 contain complementary analyses, applying Spatial Fixed Effects Difference-in-Differences models to estimate impacts of windfarm proximity and visibility on house prices. Both analyses also include a triple-difference analysis comparing effects of treatment by visible and not-visible windfarms. The empirical work within these chapters serve as two standalone analyses, as well as testing the findings for robustness to the methodological framework applied. Both analyses include only property transactions located near to a windfarm or wind turbine, comparing prices before and after nearby windfarms become operational. The key difference between the two chapters is that Chapter 3 makes use of the average transaction prices at the postcode-quarter level, while Chapter 4 makes use of the transaction prices of properties sold at least twice over the study period. The main advantage of the average price analysis is that it allows for the inclusion of all property transactions in the dataset. The repeat sales analysis has the advantage that individual properties are compared to themselves, ensuring that the observed price impacts are due to treatment by windfarms rather than changes in the properties transacted at a given time. In the following two sections I discuss the key findings and contributions of both chapters, and in the third I discuss the insights for future research into this topic.

6.4.1 Chapter 3

Chapter 3 evaluates the impacts of windfarm proximity and visibility on nearby house prices through the application of an Average Sales Approach Hedonic Pricing analysis. First, this chapter replicated the analysis by Gibbons (2015). I find statistically significant decreases of the average housing transaction price in postcodes located near to visible windfarms. Over the same period, I find statistically significant increases in the average transaction price for postcodes located near windfarms, but for which the windfarms are not visible. The estimated impacts are considerably larger than those reported by Gibbons, particularly for properties located very distant to windfarms. The chapter then goes on to extend the analysis to include an additional 12 years of property transaction data and 6 additional years of windfarm data. In the headline results of the extension, I find that windfarm proximity and visibility are now associated with statistically significant positive price effects. I also find that the analysis is highly sensitive to the study period, and evidence that opinions towards windfarm visibility has changed from an environmental disamenity to and environmental amenity.

This chapter makes several contributions to the literature on windfarms and house prices. First it replicates and updates the literature to include an analysis that extends both 5 years prior to and 6 years after the analysis by Gibbons. This analysis makes use of the largest dataset in the literature to estimate effects of windfarm visibility and proximity.

This extension covers a period of rapid change in the distribution and concentration of wind turbines across England and Wales. It is the first analysis to test for sensitivity to alternative visibility estimates using GIS software. It also presents evidence that the size and direction of the price effect is sensitive to the number of nearby windfarms, which has not previously been tested within the literature. In addition, this is the first analysis to test for differences in price effect arising from the inclusion of windfarms sited in both urban and rural locations. Further, it is the first analysis within the literature to explore the consistency of the price effect across subdivisions of the total study period. By doing so, I find evidence that the price effects of windfarm visibility and proximity are not stable over time - they have gone from largely negative in the earliest periods to largely positive in the latest periods.

6.4.2 Chapter 4

Chapter 4 evaluates the impacts of windfarm proximity and visibility on nearby house prices through the application of a Repeat Sales Approach Hedonic Pricing analysis. In this chapter, I find that wind turbine visibility is not consistently associated with with negative price impacts, and lack of visibility is not consistently associated with positive price impacts, and the impacts while statistically significant are below a 1% change in price. If detailed property characteristics are controlled for, I do find a statistically significant decrease arising from visibility, and a statistically significant increase from lack of visibility at the 0-1km range. However, when the visibility measure is enhanced and takes account of intervening buildings, there are no statistically significant negative impacts from windfarm visibility or lack of visibility on the analyzed properties. I find that the analysis is sensitive to changes in the underlying data regarding the visibility analysis, as well as the inclusion of detailed property characteristics. I also tested these results for robustness to alternative study periods. I find evidence of negative price effects from turbine visibility and positive impacts from non-visible turbines in the earliest periods, but by the final period visibility is associated with a positive price effect. The analysis of this chapter concurs with the findings of Chapter 3, and that under a repeat sales analysis windfarm visibility has transitioned from an environmental disamenity to an amenity in England and Wales.

This chapter makes key contributions to the literature on windfarms and house prices. It updates the literature to include the most recent analysis of windfarm impacts on house prices by applying the first repeat sales analysis examining the house price impacts from windfarm visibility and proximity in England and Wales. The analysis is also the first in England and Wales to incorporate building height data, as well as detailed property characteristics into its analysis to test the robustness of the findings. The analysis finds evidence that there is sensitivity regarding both the inclusion of detailed visibility estimation and property characteristics suggesting that these are key variables that should be included in analyses of this kind. I also show that treatment by multiple wind turbines leads to larger negative impacts from visibility and larger positive impacts from lack of visibility. This is an important finding implying that intensity of treatment is an important factor often overlooked by previous work, as many analyses from the literature define treatment by 'at least' one turbine. The analysis of this chapter finds further evidence that price impacts from windfarm proximity and visibility are not stable over time, which has been untested within the current literature.

6.4.3 Chapters 3 and 4 Together

The empirical analyses of Chapters 3 and 4 make standalone contributions to the research into house price impacts from windfarm siting. The two analyses together also make contributions to the literature. This thesis contains the first application of both an Average Price and a Repeat Sales analysis to the same underlying model to estimate impacts from windfarm siting in England and Wales. In this sense, the two analyses serve as robustness tests for each other, and I find that the estimated impacts are broadly consistent across the two models. This is an important contribution as both approaches have strengths and weaknesses specific to themselves. The differences in the visibility estimation processes allows for an important comparison of how the definition of 'visible' can influence results. The analysis of Chapter 3 found little impact from alternative visibility measures between postcode and windfarm centroids using higher resolution DEMS. However, the work undertaken in Chapter 4 shows that the inclusion of building height data may be a key factor in generating accurate impacts of windfarm visibility on house prices.

The complementary empirical analyses contained in these chapters are in agreement that the house price impacts from wind turbine siting in England and Wales are largely positive. These findings are consistent across both an average price and repeat sales approach in a fixed-effects difference-in-difference setting. Under both analyses, I find that there is sensitivity to the study period analyzed, and that impacts from visibility of nearby windfarms or turbines have transitioned from negative to positive. This implies that not only has the landscape changed significantly since the previous analysis by Gibbons (2015), but so have the opinions of home buyers. As such, the analyses of Chapters 3 and 4 update the analysis by Gibbons by including additional housing transaction data as well as applying a Repeat Sales analysis. They also generate important insights into the findings of the the current literature suggesting that long study periods may hide shifting attitudes towards windfarm proximity and visibility.

6.4.4 Insights and Avenues for Future Research

The literature examining house price impacts from windfarm siting is well-developed, but has considerable room to grow. I have organized the discussion on how new research could build upon the work of Chapters 3 and 4 into the following categories: Data, Analytical Approaches, and Related Research.

Data

In regards to improvements to the data informing the analysis, an obvious improvement would be the use of more detailed property characteristics which would include information on changes to these characteristics over time. The ability to track such changes would ensure that any observed price effects from windfarm siting are not arising from improvements (deterioration) increasing (decreasing) the quality of the homes in the sample independently. A further refinement would be the inclusion of the exact turbine locations of all wind turbines in England and Wales, which was not available at the time of writing. This would allow for a deeper investigation into impacts from different intensities of treatment. It would also expand the analysis to better compare effects from the full set of individual turbines to windfarm centroids in England and Wales.

To better estimate visibility, the use of a Digital Surface Model accounting for land features like trees as well as buildings would be a substantial improvement. Such a model was not available for all of England and Wales, and unfortunately the current available DSM has gaps even within cities which would impose severe restrictions on the analysis. As the DSM coverage increases future research should incorporate this data. Until a national DSM is available, future research could make site visits to compare visibility of windfarms/turbines from postcode centroids and compare this to visibility at the property level. Of course, it may not be feasible to visit each property or postcode in the country, but even visiting a random sample could generate important insights to the accuracy of any GIS modeling used.

It would also be ideal if a UK-wide dataset could be constructed to compare the effects across England, Northern Ireland, Scotland and Wales. This would allow for testing whether all regions experience consistent impacts. They could then test if there is a similar trend within each country where early windfarms cause negative impacts, but these effects decrease or reverse with time - this may provide useful insights to policy makers on how best to compensate affected homeowners. Lastly, a very useful dataset would be location-specific survey data similar to what is provided by the BEIS Public Attitudes Tracker. This survey tracks public opinion on renewable energy developments, and even asks if they would accept a new development near their home. If the responses were broken down into smaller localities rather than the national level, it would be possible to test whether the stated attitudes match the revealed attitudes found through HPM analyses. This would be particularly interesting if the support was stronger in areas unsuited to wind energy developments.

Analytical Approaches

There are a few avenues that future researchers could take regarding the analytical framework applied to estimate house price impacts from windfarm siting. The fixed-effects difference-in-difference model is a strong identification strategy, and was the best option for the analyses of Chapters 3 and 4. However, since these analyses were performed, the approach has been shown to be susceptible to bias if the timing of treatment is staggered (not uniform for all treated units). This is particularly troublesome for long study periods. Future research applying this framework should apply a decomposition as proposed by Goodman-Bacon (2021).

It would also be interesting to apply newer analytical methods to gain new insights into this area. The use of a Synthetic Control approach would be a particularly interesting novel way to test for windfarm siting effects on house prices. This approach has not yet been applied within the literature, but it essentially creates a control group based on pre treatment trends of the control and treated units. There have recently been extensions to the Synthetic Control framework which allow for an analysis of multiple treated units (Kreif et al. 2016), as well as the inclusion of fixed-effects Xu (2017). These recent innovations would allow for analyses to relax the Parallel Trends Assumption required within a DID framework, as the SC approach generates less biased results when the assumption is violated relative to a DID analysis.

Related Research

Future research can also build on the work of this thesis by exploring related questions involving windfarm siting and house prices. First, it would be interesting to determine if there are any house price effects arising from windfarm rejection - i.e. do homes nearby proposed windfarm sites increase in value when planning permission is not granted? In a similar vein, testing for causal relationships between windfarm rejection and local characteristics would also be an interesting research area. This thesis has found that properties near windfarm developments sell at lower prices than those far away, and also that properties with views of turbines sell at lower prices than those without even before treatment. It would be interesting to test for a relationship between rejection and home values. Future research could also explore whether variation in the community investment from windfarms has any impact on home values, or on planning permission being granted.

Just as the landscape has significantly changes since the time of the analysis by Gibbons, the windfarm stock of England and Wales is set to experience further changes in the future. Many early windfarms are now at the end of their expected lifespans, and are beginning to be replaced. Often, the planned replacement turbines are larger than those they are replacing. Future research could test whether this has any impact, or whether the presence of nearby turbines is already baked into house prices regardless of their size. There has also been a substantial increase in the number of commercial solar farms in recent years. Future researchers could use insights from the work of this thesis to explore how other commercial-scale renewable energy technologies impact on nearby house prices.

Summary

The analyses of Chapters 3 and 4 provide several key contributions to the existing literature on the impacts of windfarm siting and house prices. These contributions are made both by each chapter individually as well as the insights gained from comparing the results of both sets of empirical analyses. There are several potential avenues for future researchers to build on this work by deepening our understanding of house price impacts from nearby windfarms. This could be achieved through the use of more detailed datasets or by applying novel analytical frameworks. There is also considerable scope to widen the literature area by researching related questions and issues that have not yet been explored while taking lessons from the research presented within Chapters 3 and 4.

6.5 Chapter 5

Chapter 5 examines whether there is a solar premium for homes in the Uk as well as estimating the capitalization of residential solar PV systems into house prices through the application of a Hedonic Regression Model (Model 1), a Repeat Sales Analysis (Model 2), and Propensity Score Matching (Model 3). I find that there is a solar property premium in England and Wales and that residential PV systems are capitalized into English and Welsh house prices. The capitalization estimates are £9,417 (Model 1), £4,651 (Model 2), and £6,578 (Model 3). The average installation cost for a PV system in the UK is approximately £4,800 - the findings of this chapter show that this cost will either be nearly recouped through the sale of a home, or potentially lead to profit.

The academic literature on solar property premiums and residential PV system capitalization into house prices is in the early stages of development. As such, this chapter makes several key contributions into this research area. It is the first analysis to test for a solar property premium or capitalization of residential PV systems into English and Welsh house prices. This is an important extension of the current literature which has previously examined only locations with considerably larger solar potential and therefore potential for large energy savings or export profits. My findings are consistent with those reported in the literature, though I find a much smaller solar premium and capitalization. The lack of solar potential in England and Wales may partly explain these findings, as well as highlight the important role of government subsidies in generating the returns on installation found. The chapter exploits a large and rich dataset to perform the first application of a Propensity Score Matching analytical framework to estimate the capitalization of residential PV systems into house prices. This is an improvement on the fuzzy matching techniques previously applied within the literature.

The findings of Chapter 5 supports the reported impacts of residential solar within the literature being positively capitalized into house prices. However, the impacts of commercial-scale are negatively capitalized into house prices Dröes & Koster (2021). This implies that the subsidies paid to homeowners, who choose to install residential PV systems or energy savings are likely a large driver of the difference between these two uses of the same technology.

6.5.1 Insights and Avenues for Future Research

The literature examining the capitalization of residential PV systems into house prices, and the solar property premium is small. At the time of writing, there were only six peer reviewed papers exploring this topic area. There is considerable scope for future research exploring this topic, and in this section I will outline how the the work contained in Chapter 5 can be built upon by future researchers, again organized by: Data, Analytical Approaches and Related Research.

Data

The analysis of this chapter was limited by the availability of data on solar properties. Future research should make use of larger datasets which contain the full list of solar properties in England and Wales which is controlled by OFGEM, and which was not accessible for the work of this thesis. This dataset includes information on every property which has entered into the FIT scheme and has details such as: The size of each installation, the total electricity generated by the system, the FIT payments made to the owner and the export payments made to the owner. The OFGEM data also includes information for the entire UK, and would allow for a national analysis. Future research could also make use of more detailed property characteristic information, in particular information on changes to property characteristics over time. This would ensure that any observed price effect is solely due to the installation of a residential PV system and not other changes to homes in the sample.

Analytical Approaches

In regards to alternative analytical approaches, future researchers could consider the application of both a Difference-in-Differences model or synthetic control methods to estimate the capitalization of residential PV into solar properties. The choice of the method, and how it could be applied will likely depend on the data available. If the parallel trends assumption holds, a DID framework would allow for a direct comparison of price trends between solar and non-solar properties. If it does not hold, then the synthetic control approach may be a superior alternative. Rather than using the PSM approach, future researchers could apply Coarsened Exact Matching, or Fuzzy matching techniques to generate a set of control properties. These have been applied within the literature, but not within the context of the UK.

Related Research

Previous research has found that a large portion of the solar premium arises from subsidies, and others have found that in order to recoup the cost of installation these subsidies are required. Future researchers with access to the OFGEM data described above would have information on the subsidies paid to owners of residential PV systems, as well as the total electricity generated. Access to this information would allow future researchers to disaggregate PV capitalization to measure the contributions of government subsidies and energy savings separately. This would be particularly interesting given both the relatively low solar energy endowment of th UK, as well as the generous subsidies provided by the Feed-in Tariff Scheme.

Residential PV systems are not the only technology eligible for the FIT scheme. Future researchers could test whether residential wind-turbines have a similar capitalization to residential PV systems. There are nearly 30,000 residential properties which have installed these small-scale wind turbines within the UK. At the time of writing, this was a completely empty area within the literature. Research into this topic would build on the work of all analytical chapters of this thesis - firstly it could test for residential wind turbine capitalization, and compare this to residential PV capitalization. Second, there would be scope to test whether the presence of a residential turbine has any negative impact on neighboring home values due to either visual effects or noise pollution. This would be particularly interesting given the nature of the FIT which would subsidize the owner, but not any impacted neighbors.

In fact, the FIT should be the subject of future research relating to both the transition to a low-carbon economy as well as unintended policy outcomes. The findings of this chapter show that residential PV systems increase house prices, potentially enough to profit from installing the system. The majority of these currently in operation within the UK participate in the FIT scheme, which provides payment for all electricity generated, even if it is used on site, as well as additional payments for exported electricity. This is paid for by all electricity users in the UK, including those who are unable to take advantage of the subsidy either because they do not own a home or cannot afford to install a PV system. Consider that to benefit from this scheme, an individual would need to both be a home owner, and have the financial ability to make the up front investment in a PV system. Future research should evaluate the welfare effects of this policy in relation to its cost, wealth transfers, and carbon reduction value.

Summary

The analysis of Chapter 5 generates several contributions to the growing literature on residential PV system capitalization into house prices, and the solar property premium. It extends the literature to include locations with relatively low solar PV potential, as well as applying novel analytical frameworks to the the research area. There is substantial potential for future research to build on the work of this chapter by using more detailed data, alternative analytical methods, or exploring related research questions.

6.6 Closing Remarks

This thesis makes several contributions to the literature on the externalities associated with renewable energy developments through the application of Hedonic Pricing Techniques. The thesis builds on and updates the current research into house price impacts from wind turbine visibility. The empirical work generates key insights into how these impacts have evolved over time, and what factors future research should consider to better model these impacts. It extends the research into the capitalization of residential PV systems into house prices to include England and Wales, as well as the first application of a Propensity Score Matching analysis.

Appendices

A1 Implementation of the Renewables Obligation and the Wind Turbine Stock of England and Wales: 1992 - 2011

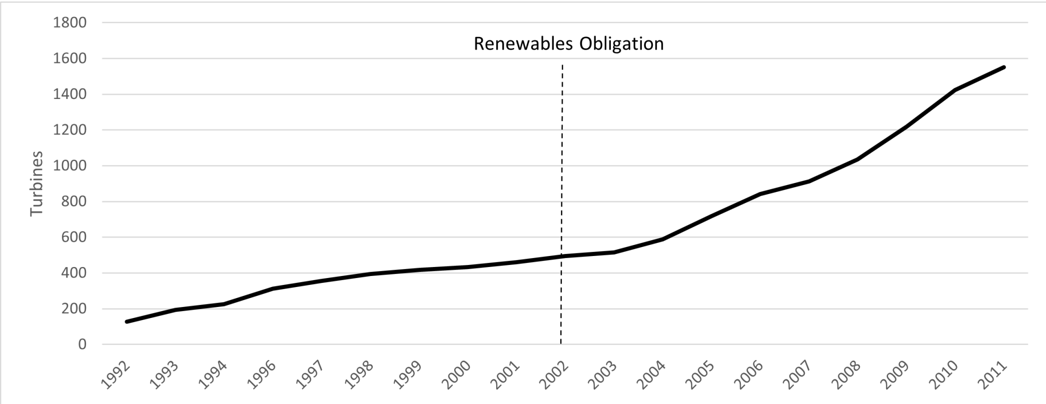

The Renewables Obligation went into effect in 2002, and there is a clear increase in the number of wind turbines in England and Wales following this. The delay between the policy going into effect and the takeoff of the wind turbine stock can be explained by the time required to receive planning permission and complete construction of a windfarm. This is the study period of Gibbons (2015), and one can see there is a substantial increase in the turbine stock - though this increase is relatively small compared to the period after Gibbons' work.

A2 Areas within 14km of English and Welsh Wind Turbines 1992-2017

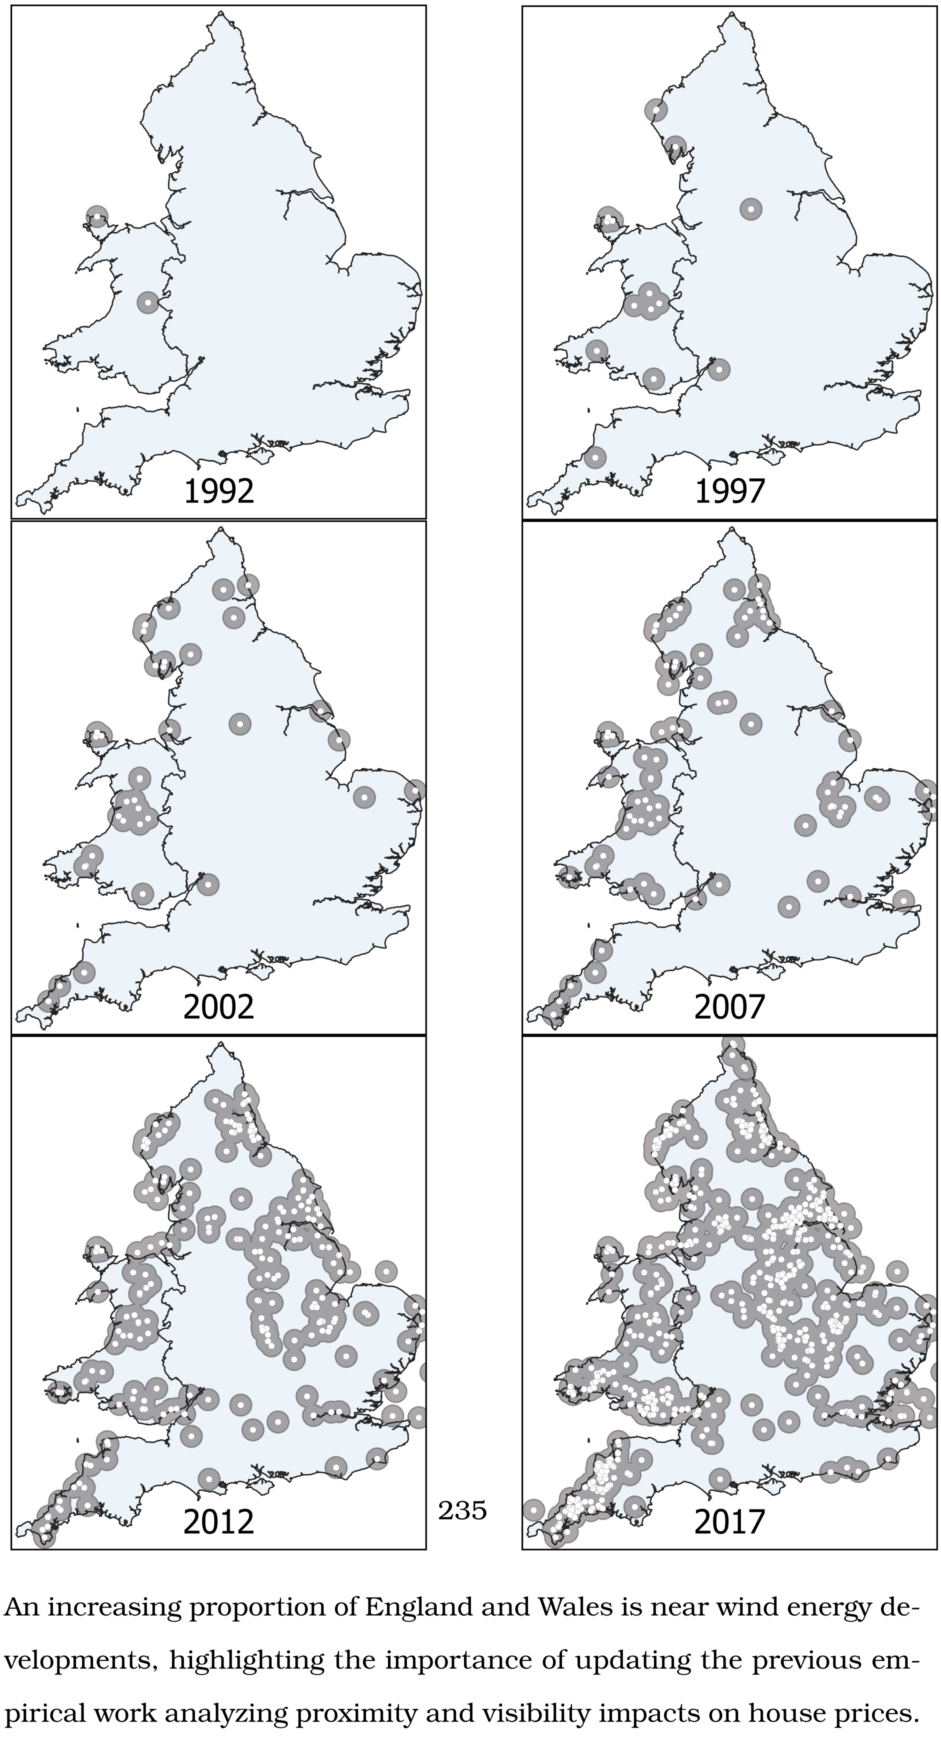

An increasing proportion of England and Wales is near wind energy developments, highlighting the importance of updating the previous empirical work analyzing proximity and visibility impacts on house prices.

A3 Comparison of Gibbons' and Replicated Windfarm Sites

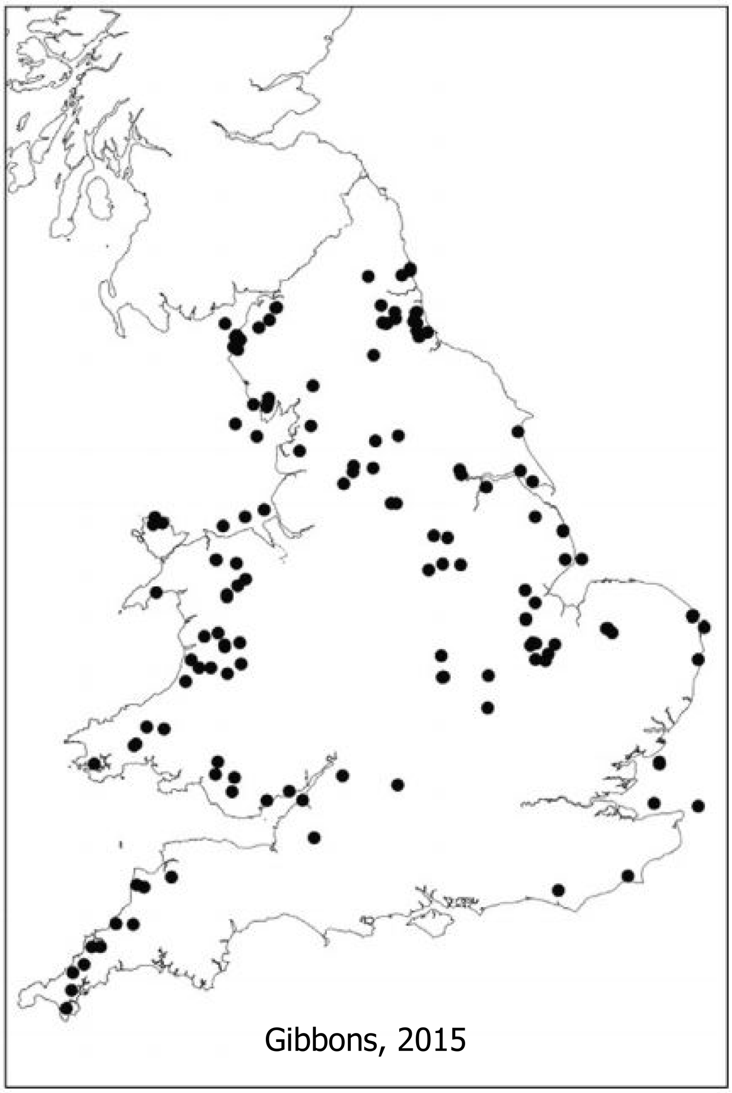

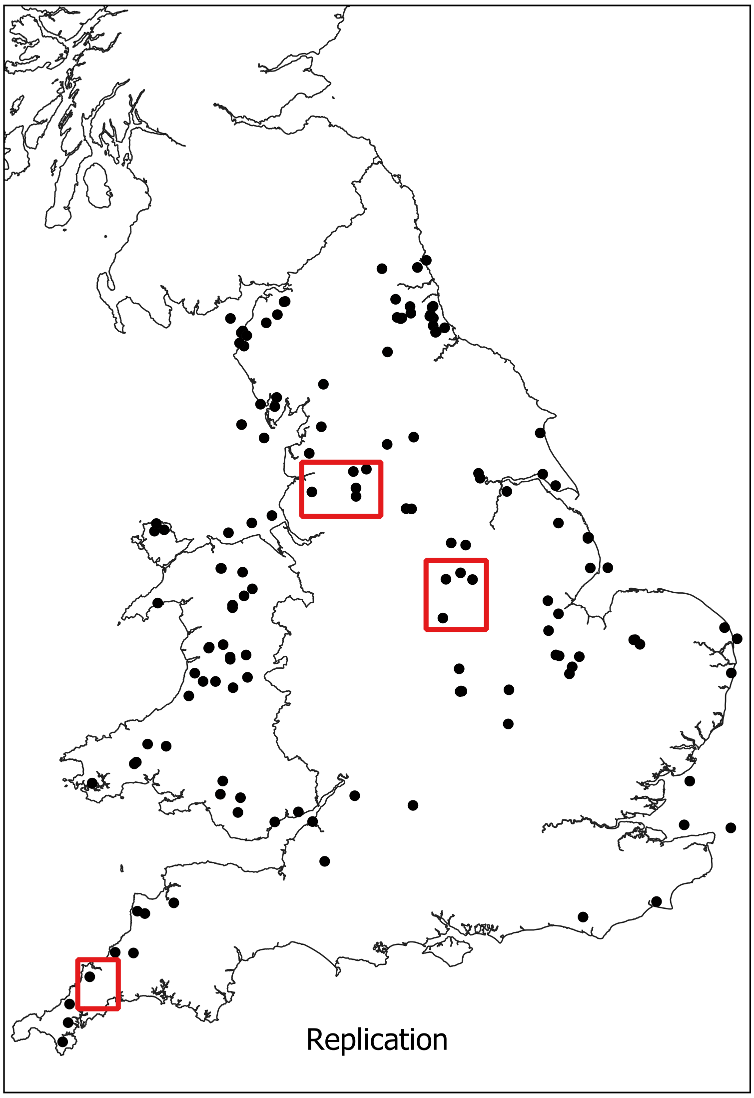

Shown on the left are the locations of the windfarms included in the analysis by Gibbons (2015) as shown within the paper. On the right are the windfarm locations of the replicated dataset, areas of difference are located within the red boxes. The replicated locations were generated by following the inclusion criteria specified by Gibbons. Although the locations do not match exactly to those presented within Gibbon's work, there are 148 windfarms in both.

A4 Example Viewshed

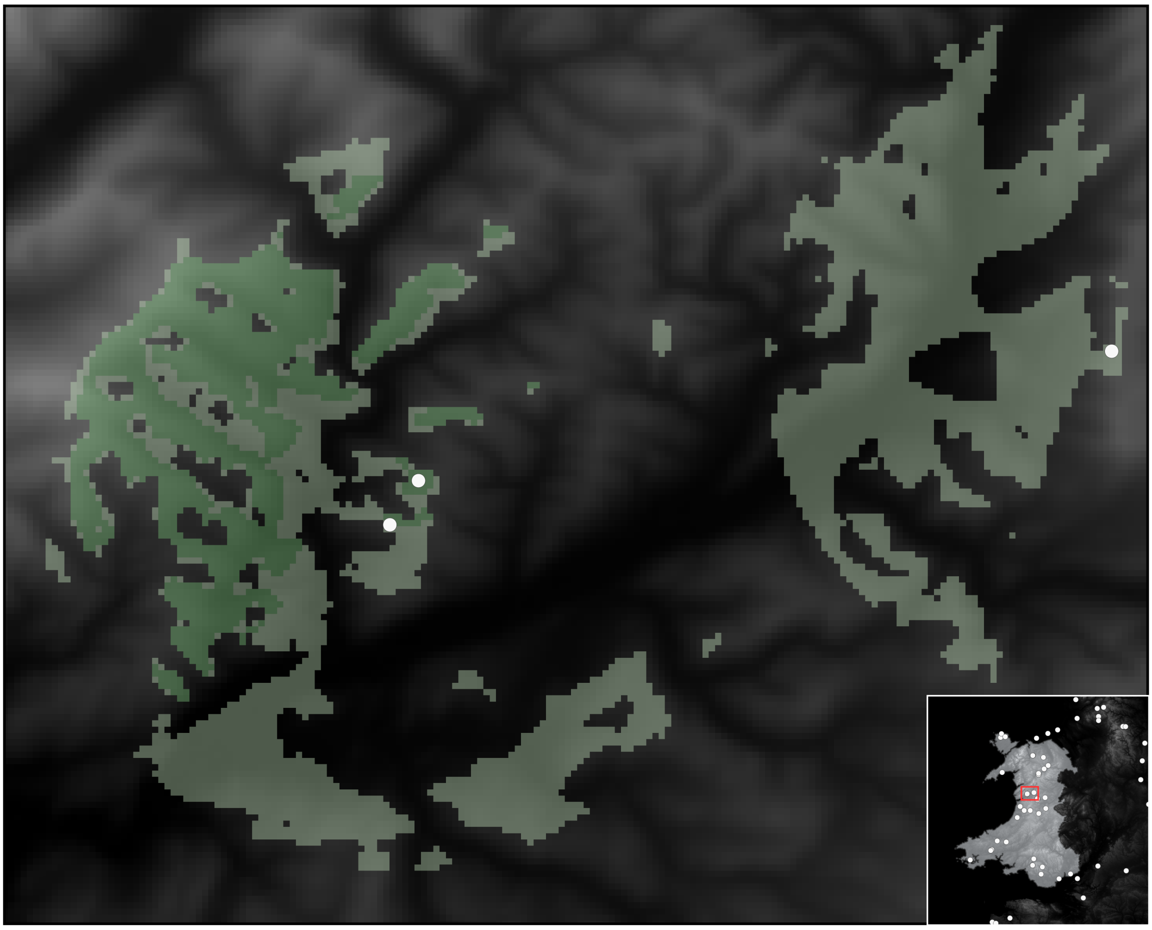

Shown here are the 5km viewsheds of (from left to right) the Bro Dyfi Community Turbine, Dullas Valley Community Turbine, and the Cemmaes windfarms in Wales. These are estimated from the centroids of each windfarm which are shown in white. Areas shaded green are expected to have a view of the windfarm centroids, and areas shaded darker green will have views of more than one centroid. Cemmaes consists of 18 turbines, and this is potentially problematic when calculating viewsheds following the centroid approach of Gibbons.

A5 Testing for Anticipation Effects

Announcement Effect

Distance Band	Visible	Not Visible
0-2km	-0.0062*	-0.0051*
(RSE)	(0.0027)	(0.0026)
2-4km	-0.0191	0.0037
(RSE)	(0.0235)	(0.0381)
4-6km	0.0080	-0.0224
(RSE)	(0.0136)	(0.0351)

When testing for anticipation effects arising from the announcement of a windfarm being approved for construction, I find statistically significant decreases for the average price in postcodes within 2km of windfarms. This decrease of about a half a percent is found for both postcodes which will have views of the future windfarms and postcodes which will not. The anticipation effects are insignificant at distances greater than 2km.

A6 Illustration of Distance Bands

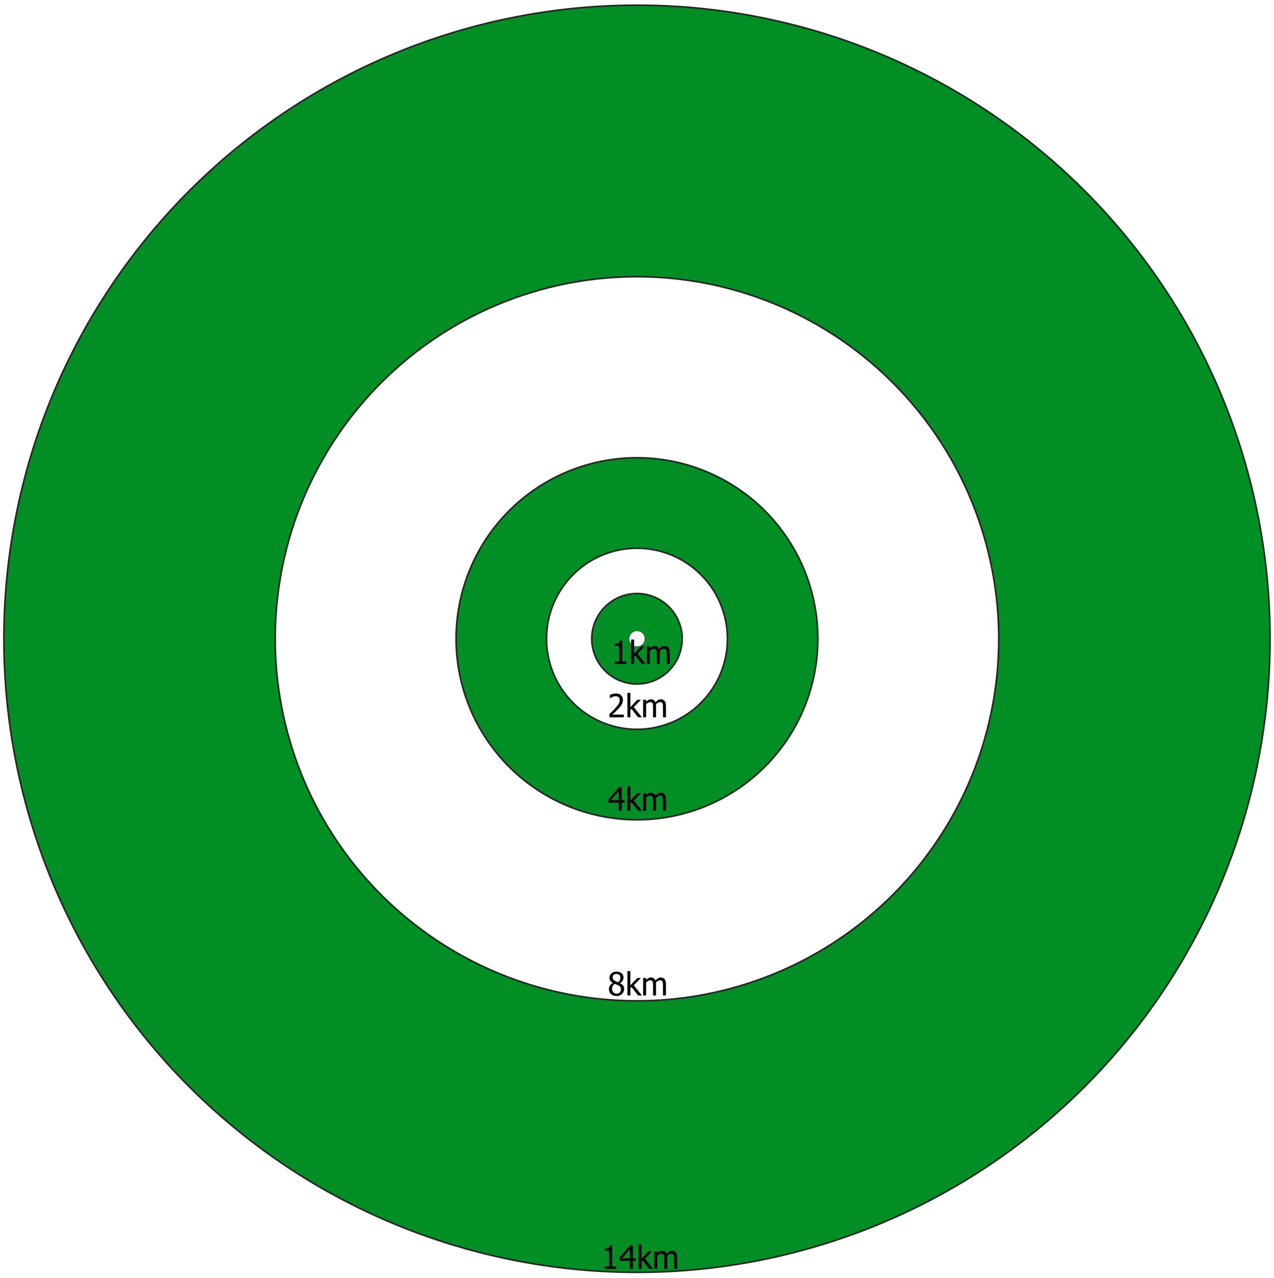

Shown here is an illustration of the distance bands used within the analyses of Chapters 3 and 4.

A7 Transaction Price Trends

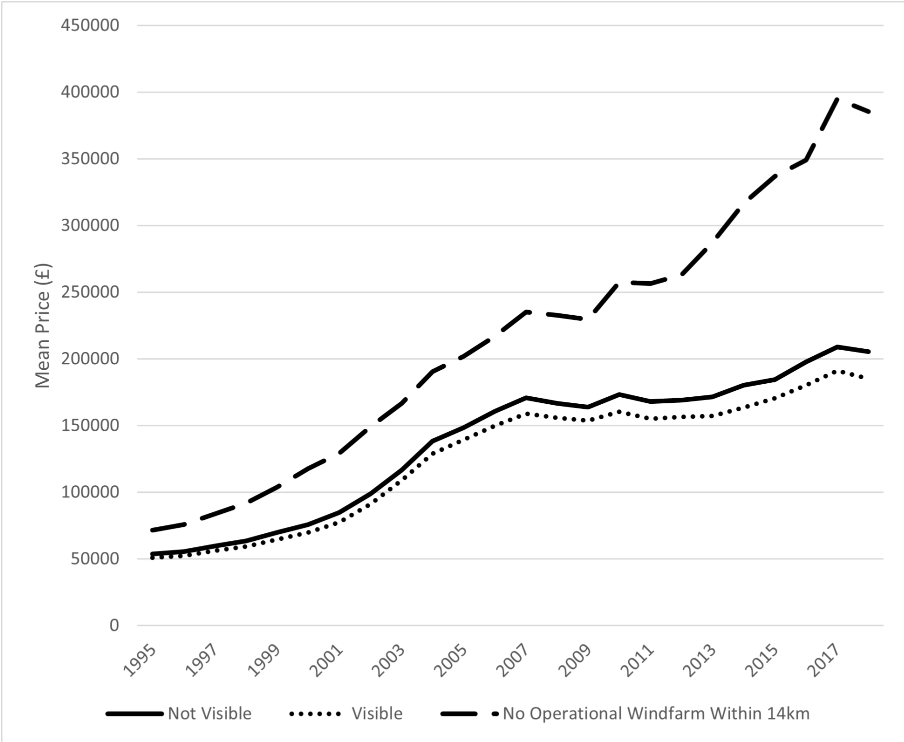

The price trends for properties within 14km of visible and non-visible wind turbines are compared to properties located more than 14km from an operational wind turbine. Much like Figure 3.5 the visible and not visible groups trend together while the set of properties more than 14km from operational windfarms follows a less similar trend. It is also noteworthy that properties within 14km of operational windfarms are considerably lower on average than properties farther away, regardless of windfarm visibility.

A8 Visibility Comparisons: Proportion of properties with visible wind turbines.

Percentage of properties with a visible wind turbine.

Distance Radius	% Visible			
	1	2	3	4
0-1km	97.11	97.04	96.24	85.33
0-2km	94.45	94.30	92.55	80.21
0-4km	84.08	82.47	80.36	62.44
0-8km	68.90	65.02	62.95	53.57
0-14km	52.73	51.80	49.06	39.51
DEM				
200m	X			
90m		X		
5m			X	
5m + BH				X

The first table shows that the proportion of properties with a view of at least one turbine varies substantially with the DEM model used to model visibility. The inclusion of building heights leads to a substantial drop in the proportion of properties with a view. This is to be expected though, as the the inclusion of building height data requires a substantial restriction of the properties included, as well as including additional features which may block views.

A9 Postcode vs property visibility estimates.

Proportion of Properties with visibility estimates different from their postcode centroid.

Distance Radius	% Difference	
	1	2
0-1km	3.22	4.30
0-2km	5.25	7.28
0-4km	9.67	10.94
0-8km	11.39	12.02
0-14km	14.44	16.21
DEM		
200m	X	
90m		X

Here I show that there are substantial differences between neighborhood and property specific windfarm visibility. This highlights the importance of the visibility estimate used to define treatment by visible windfarms.

A10 DEM Resolutions: 200m, 90m, 5m, and 5m with Building Heights

Here I show the alternate DEM resolutions, with the darkest colors representing the lowest elevations and the bright yellow the highest. The locations shown are London, though at the 200m resolution it is only just possible to make out the Thames. When building heights are included, the Thames is clearly visible and London's CBD can be seen clearly due to the concentration of tall buildings.

Shown here is the 200m resolution DEM.

Shown here is the 90m resolution DEM.

Shown here is the 5m resolution DEM.

Shown here is the 5m resolution DEM redrawn to 1m resolution with added building heights.

A11 Solar PV Potential: Hawaii, England and Wales

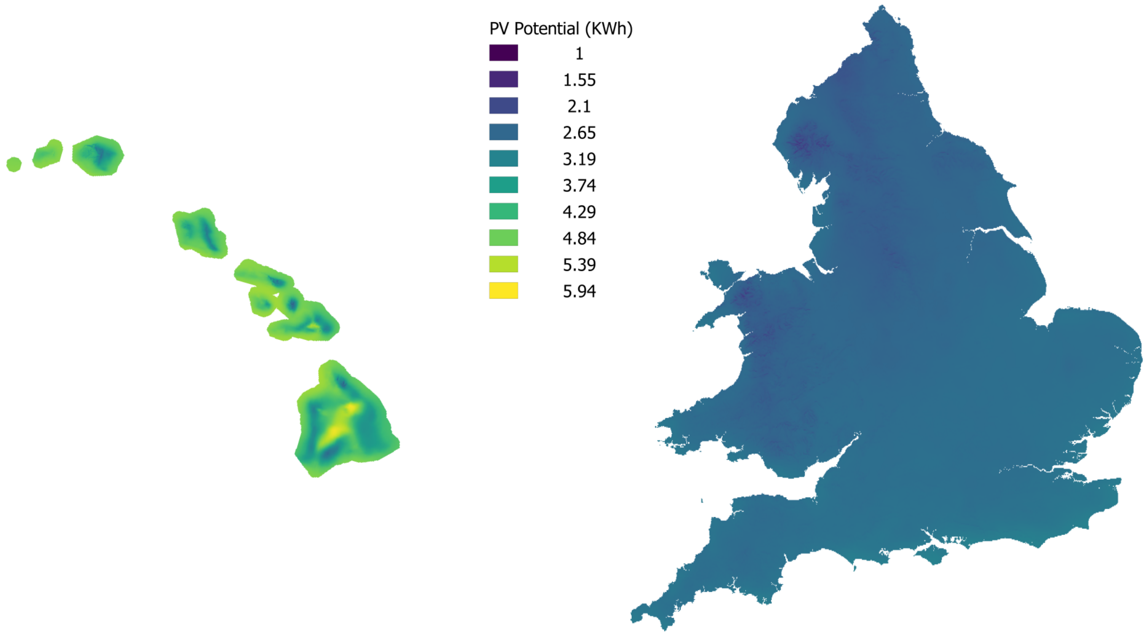

This figure compares the PV potential of Hawaii and England and Wales. Solar PV capitalization into house prices in Hawaii was analyzed by Wee (2016) who found a 5% capitalization, large enough to cover the cost of installing the average PV system in the analyzed dataset.

A12 Solar PV Potential: Arizona, England and Wales

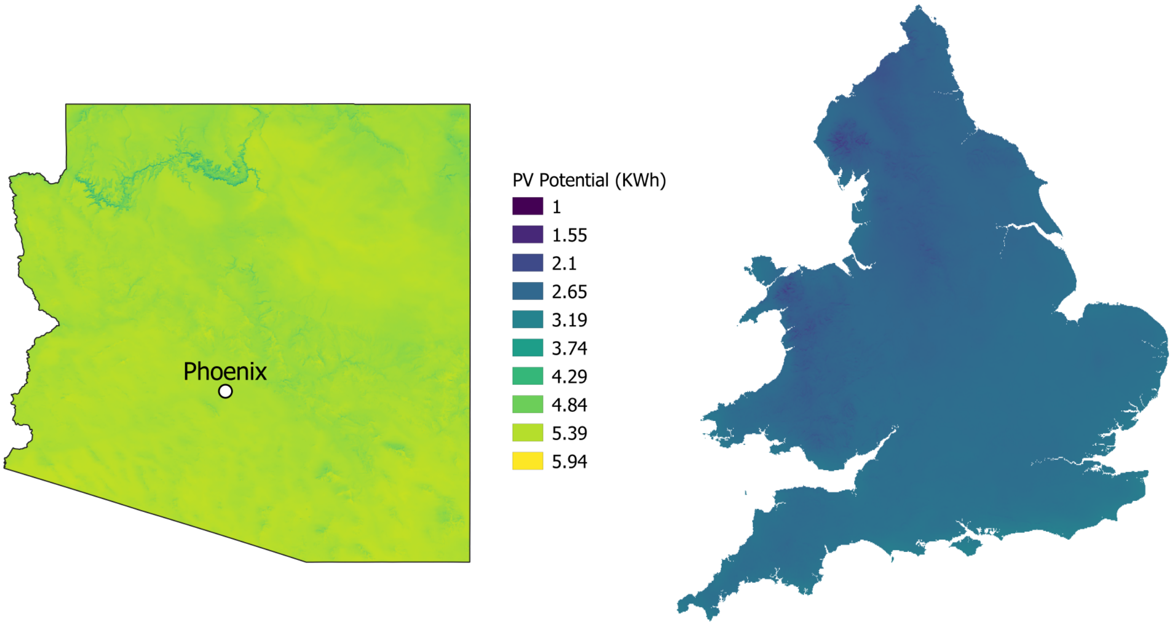

This figure compares the PV potential of Arizona and England and Wales. Solar PV capitalization into house prices in Phoenix Arizona was analyzed by Qiu et al. (2017) who found a 15-17% capitalization, large enough to cover the cost of installing as well as earn a profit over the cost of installing the average PV system in the analyzed dataset.

A13 Solar PV Potential: Western Australia, England and Wales

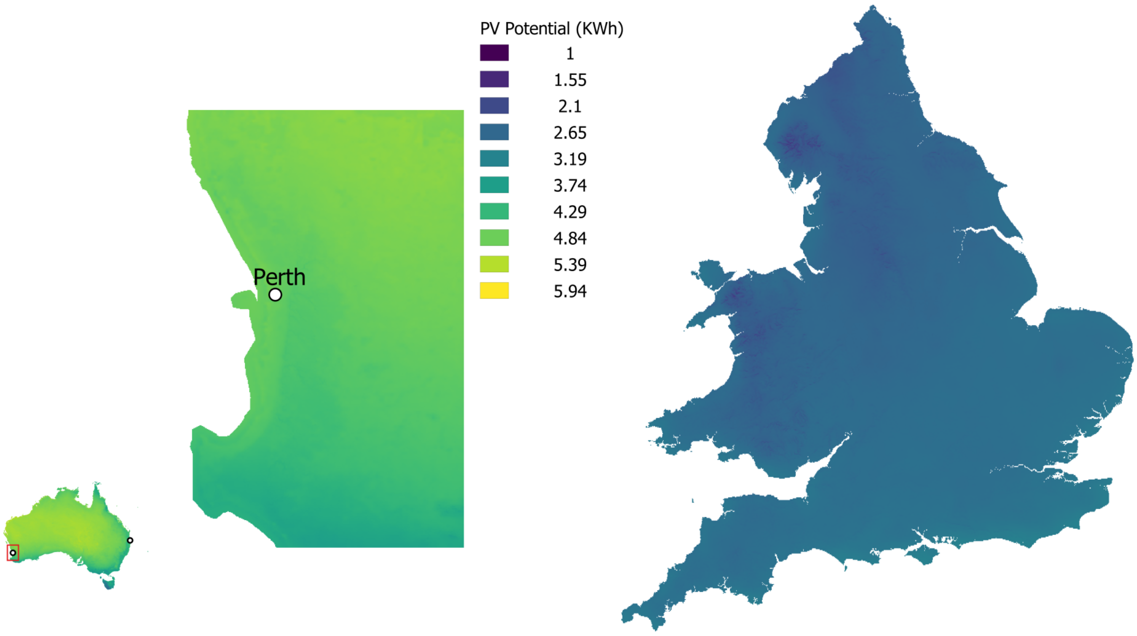

This figure compares the PV potential of Western Australia, England and Wales. Solar PV capitalization into house prices in and around Perth, Western Australia was analyzed by Ma et al. (2016) who found a 2.3-3.2% capitalization, large enough to cover the cost of installing the average PV system in the analyzed dataset.

A14 Solar PV Potential: Southport Australia, England and Wales

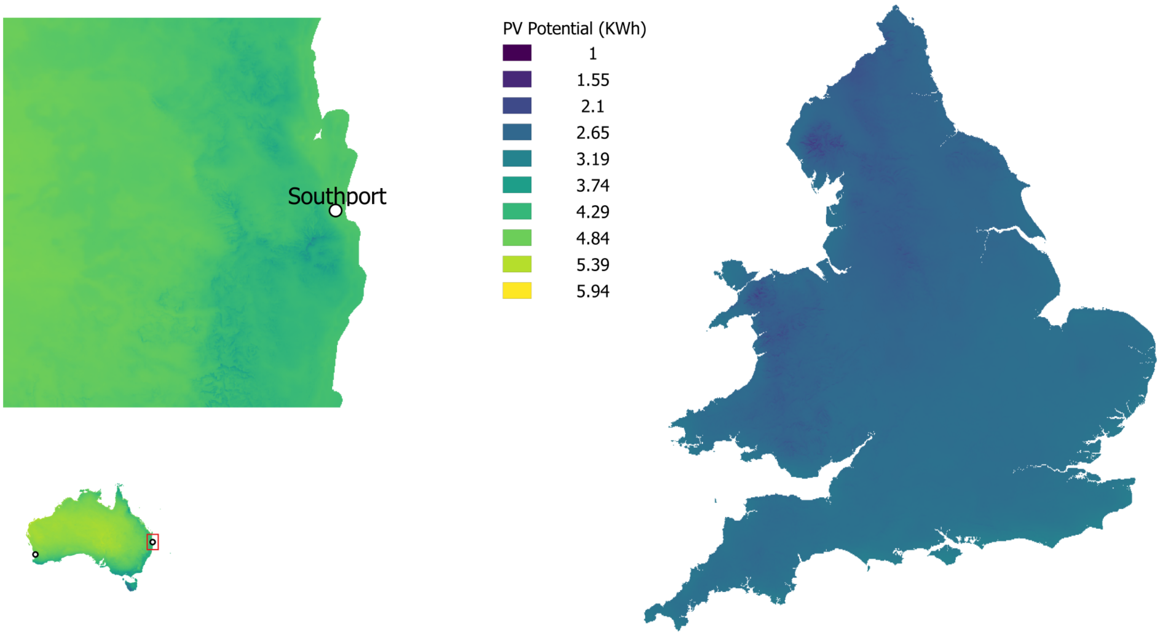

This figure compares the PV potential of Southport, Queensland and England and Wales. Solar PV capitalization into house prices in and around Southport was analyzed by Lan et al. (2020) who found a solar premium of $21,403 AUD, large enough to cover the cost of installing the average PV system, and make a profit in the analyzed dataset.

A15 Additional Control Coefficients

Ln Price	Coefficient	RSE
Repeated Sales	0.0718***	0.0075
Lighting Costs	-0.0051***	0.0001
Heating Costs	-0.0005***	0.0000
Proportion of Windows Multiglazed	0.0029***	0.0003
Number of Extensions	0.0753***	0.0054
Emissions	-0.0298***	0.0045
Energy Rating		
A	0.0506***	0.0209
B	0.0603***	0.0208
C	0.0600***	0.0208
D	0.0524	0.2086
E	-0.0479***	0.0212
F	-0.0674***	0.0228
G	-	-
Year of Sale		
1996	-0.0103	0.0259
1997	0.0895***	0.0255
1998	0.1388***	0.0252
1999	0.2061***	0.0251
2000	0.2560***	0.0248
2001	0.3725***	0.0248
2002	0.5200***	0.0242
2003	0.7315***	0.0247
2004	0.9220***	0.0246
2005	1.0011***	0.0256
2006	1.0863***	0.0248
2007	1.1086***	0.0250
2008	1.1165***	0.0301
2009	1.1147***	0.0303
2010	1.2099***	0.0298
2011	1.1307***	0.0304
2012	1.2298***	0.0289
2013	1.2699***	0.0261
2014	1.2927***	0.0245
2015	1.3128***	0.0268
2016	1.2747***	0.0280
2017	1.2002***	0.0282
2018	1.3222***	0.0277
2019	1.2998***	0.0274
2020	1.3910***	0.0450

Here I report the coefficients of additional property characteristics and their influences on transaction price.

A16 Moran's I Test for Spatial Auto-correlation

	Variable	Moran's I	E(I)	SE(I)	Z(I)	p-Value
Model 1	LnPrice	0.54517	-0.00006	0.00935	58.30666	0.00000
Model 2	Ln Price	0.30110	-0.00006	0.00188	160.35806	0.00000

Here I present the Moran's I tests for Spatial Auto-correlation within the datasets of Models 1 and 2 of Chapter 5. The null hypothesis is rejected and therefore there is spatial auto-correlation within the dataset. This is to be expected, as it simply implies that properties located near to each other sell at similar prices. This auto-correlation is absorbed in the model through the inclusion of location-specific and time specific controls.

Data and Copyright Acknowledgments

Property Transactions

Property Addresses and Geo-coordinates

Property Characteristics

Postcode Information

Postcode geo-coordinates were obtained via the Office for National Statistics Postcode Directory and made available through the Open Government License.

Windfarms and Wind Turbines

Windfarm centroid locations, windfarm characteristics and windfarm status were provided by renewableUK through an academic license.

Wind turbine locations were provided by Simon Mallet at RenewablesMapUK on an academic license.

Digital Maps

This Thesis includes data obtained through Digimaps, an online mapping data service managed by the University of Edinburgh.
Landcover Map 2010: LCM2010 © NERC (CEH) 2011. Contains Ordnance Survey data © Crown Copyright 2007, Licence number 100017572.
Landcover Map 2000: LCM2000 © NERC (CEH).
90m Digital Elevation Model: Pope, Addy. (2017). GB SRTM Digital Elevation Model (DEM) 90m, [Dataset]. EDINA. https://doi.org/10.7488/ds/1928.
200m Digital Elevation Model: Pope, Addy. (2017). GB SRTM Digital Elevation Model (DEM) 90m, [Dataset]. EDINA. https://doi.org/10.7488/ds/1928.
5m Digital Elevation Model: Landmap (2014): Landmap: Elevation Earth Observation Collection. NERC Earth Observation Data Centre, 2018. http://catalogue.ceda.ac.uk/uuid/f5d48fb0372be25b737f92976e386f53.
Building Height Data: Landmap; The GeoInformation Group (2014): UK building heights. NERC Earth Observation Data Centre, 2018. https://catalogue.ceda.ac.uk/uuid/e1449eac35108b49d3a0af26bb6d2060

Bibliography

Anderson, S. T. & West, S. E. (2006), 'Open space, residential property values, and spatial context', *Regional science and urban economics* **36**(6), 773–789. 26

ASP 15 (2019), 'Climate change (emissions reduction targets) (scotland) act 2019'. 2

Athey, S. & Imbens, G. W. (2018), Design-based analysis in difference-in-differences settings with staggered adoption, Technical report, National Bureau of Economic Research. 50, 69

Athey, S. & Imbens, G. W. (2021), 'Design-based analysis in difference-in-differences settings with staggered adoption', *Journal of Econometrics* . 50

Austin, P. C. (2011), 'An introduction to propensity score methods for reducing the effects of confounding in observational studies', *Multivariate behavioral research* **46**(3), 399–424. 167, 169, 170, 171

BEIS (2020*a*), 'Beis public attitudes tracker'. 56, 63, 116, 157

BEIS (2020*b*), 'The uk's nationally determined contribution communication to the unfccc'. 2

Black, A. J. (2004), 'Financial payback on california residential solar electric systems', *Solar energy* **77**(4), 381–388. 159, 207

Black, D. A. & Kniesner, T. J. (2003), 'On the measurement of job risk in hedonic wage models', *Journal of risk and uncertainty* **27**(3), 205–220. 24

Blumenschein, K., Blomquist, G. C., Johannesson, M., Horn, N. & Freeman, P. (2008), 'Eliciting willingness to pay without bias: evidence from a field experiment', *The Economic Journal* **118**(525), 114–137. 64, 119, 163

Bourassa, S. C., Hoesli, M. & Sun, J. (2004), 'What's in a view?', *Environment and Planning A* **36**(8), 1427–1450. 107, 151

Brander, L. M. & Koetse, M. J. (2011), 'The value of urban open space: Meta-analyses of contingent valuation and hedonic pricing results', *Journal of environmental management* **92**(10), 2763–2773. 26

Brinkley, C. & Leach, A. (2019), 'Energy next door: A meta-analysis of energy infrastructure impact on housing value', *Energy Research & Social Science* **50**, 51–65. 27

Brown, J. N. & Rosen, H. S. (1982), 'On the estimation of structural hedonic price models', *Econometrica: Journal of the Econometric Society* pp. 765–768. 25

Brown, J. P., Pender, J., Wiser, R., Lantz, E. & Hoen, B. (2012), 'Ex post analysis of economic impacts from wind power development in us counties', *Energy Economics* **34**(6), 1743–1754. 3, 5

Callaway, B., Sant'Anna, P. H. et al. (2018), 'Difference-in-differences with multiple time periods and an application on the minimum wage and employment', *arXiv preprint arXiv:1803.09015* pp. 1–47. 50

Case, B. & Quigley, J. M. (1991), 'The dynamics of real estate prices', *The Review of Economics and Statistics* pp. 50–58. 167

Cebula, R. J. (2009), 'The hedonic pricing model applied to the housing market of the city of savannah and its savannah historic landmark district', *Review of Regional Studies* **39**(1), 9–22. 27

CEDA (2020), 'Uk building heights'. 47

Chay, K. Y. & Greenstone, M. (2005), 'Does air quality matter? evidence from the housing market', *Journal of political Economy* **113**(2), 376–424. 24

Chen, W. Y. & JIM, C. Y. (2010), 'Amenities and disamenities: A hedonic analysis of the heterogeneous urban landscape in shenzhen (china)', *Geographical Journal* **176**(3), 227–240. 26

Chen, W. Y., Li, X. & Hua, J. (2019), 'Environmental amenities of urban rivers and residential property values: A global meta-analysis', *Science of the Total Environment* **693**, 133628. 59

Chi, B., Dennett, A., Oléron-Evans, T. & Morphet, R. (2019), Creating a new dataset to analyse house prices in england, Technical report, CASA Working Paper 213. Available at: https://www. ucl. ac. uk/bartlett/casa 188

Cohen, M. A. (1990), 'A note on the cost of crime to victims', *Urban Studies* **27**(1), 139–146. 20

Coley, M. C. (2005), House and landscape value: An Application of hedonic pricing technique investigating effects of lawn area on house selling price, PhD thesis, Citeseer. 27

Dalla Longa, F., Kober, T., Badger, J., Volker, P., Hoyer-Klick, C., Hidalgo Gonzalez, I., Medarac, H., Nijs, W., Politis, S., Tarvydas, D. et al. (2018), 'Wind potentials for eu and neighbouring countries'. 16

Dastrup, S. R., Zivin, J. G., Costa, D. L. & Kahn, M. E. (2012), 'Understanding the solar home price premium: Electricity generation and “green” social status’, *European Economic Review* **56**(5), 961–973. 3, 10, 159, 163, 165, 166, 172, 173, 207, 208

Denant-Boemont, L. & Hammiche, S. (2019), Economic measurement of environmental costs for transportation activity, *in* ‘Sustainable Transportation and Smart Logistics’, Elsevier, pp. 153–167. 20

Dröes, M. I. & Koster, H. R. (2016), ‘Renewable energy and negative externalities: The effect of wind turbines on house prices’, *Journal of Urban Economics* **96**, 121–141. 3, 6, 33, 34, 36, 38, 39, 40, 43, 46, 50, 52, 121, 122, 123, 125, 126, 150

Dröes, M. I. & Koster, H. R. (2021), ‘Wind turbines, solar farms, and house prices’, *Energy Policy* **155**, 112327. 158, 206, 228

EMBER (2020), ‘European electricity’. viii, 176

EPC (2018).
URL: *https://epc.opendatacommunities.org/docs/guidance* 188

EST (2020), ‘A comprehensive guide to solar panels’.
URL: *https://energysavingtrust.org.uk/advice/solar-panels/* 204, 205

Frisco, M. L., Muller, C. & Frank, K. (2007), ‘Family structure change and adolescents’ school performance: A propensity score approach’, *Journal of Marriage and Family* **69**(3), 721–741. 185

Frye, B. E. & Bartlett, J. E. (2017), ‘Performing the psm analysis: An applied example’, *Journal of Applied Research in the Community College* **24**(1), 43–58. 185, 203

Gibbons, S. (2015), ‘Gone with the wind: Valuing the visual impacts of wind turbines through house prices’, *Journal of Environmental Economics and Management* **72**, 177–196. 3, 5, 6, 29, 33, 34, 36, 38, 40, 41, 43, 46, 50, 52, 53, 54, 56, 57, 59, 61, 64, 66, 67, 68, 71, 74, 78, 79, 82, 83, 85, 88, 89, 98, 99, 100, 103, 104, 111, 115, 125, 126, 132, 136, 155, 157, 164, 219, 222, 226, 234, 236

Gibbons, S., Mourato, S. & Resende, G. M. (2014), 'The amenity value of english nature: a hedonic price approach', *Environmental and Resource Economics* **57**(2), 175–196. 26

Goodman-Bacon, A. (2021), 'Difference-in-differences with variation in treatment timing', *Journal of Econometrics* . 225

Greenstone, M. (2017), 'The continuing impact of sherwin rosen's "hedonic prices and implicit markets: product differentiation in pure competition"', *Journal of Political Economy* **125**(6), 1891–1902. 21, 22, 23, 24, 108

Heblich, S., Olner, D., Pryce, G. & Timmins, C. (2016), 'Impact of wind turbines on house prices in scotland'. 3, 5, 6, 32, 34, 36, 37, 38, 41, 43, 46, 47, 50, 52, 53, 56, 63, 66, 100, 103, 111, 122, 123, 124, 125, 126, 150, 152, 153, 157

Heintzelman, M. D. & Tuttle, C. M. (2012), 'Values in the wind: a hedonic analysis of wind power facilities', *Land Economics* **88**(3), 571–588. 3, 4, 33, 34, 36, 37, 38, 42, 44, 49, 121, 126, 150

Hoen, B. & Atkinson-Palombo, C. (2016), 'Wind turbines, amenities and disamenitites: Astudy of home value impacts in densely populated massachusetts', *Journal of Real Estate Research* **38**(4), 473–504. 3, 5, 29, 33, 34, 36, 38, 43, 44

Hoen, B., Brown, J. P., Jackson, T., Thayer, M. A., Wiser, R. & Cappers, P. (2015), 'Spatial hedonic analysis of the effects of us wind energy facilities on surrounding property values', *The Journal of Real Estate Finance and Economics* **51**(1), 22–51. 3, 5, 33, 34, 36, 38, 42, 43, 44, 49

Hoen, B., Wiser, R., Cappers, P., Thayer, M. & Sethi, G. (2011), 'Wind energy facilities and residential properties: the effect of proximity and view on sales prices', *Journal of Real Estate Research* **33**(3), 279–316. 3, 5, 33, 34, 36, 38, 39, 43, 44, 45, 50, 51, 100, 124, 126, 150

Hoen, B., Wiser, R., Thayer, M. & Cappers, P. (2013), 'Residential photovoltaic energy systems in california: The effect on home sales prices', *Contemporary Economic Policy* **31**(4), 708–718. 3, 10, 159, 163, 164, 165, 172, 173, 208

Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., Rasmussen, D., Muir-Wood, R., Wilson, P., Oppenheimer, M. et al. (2017), 'Estimating economic damage from climate change in the united states', *Science* **356**(6345), 1362–1369. 13

IEA (2020*a*), 'Projected costs of generating electricity report 2020'. **URL:** *https://www.iea.org/reports/projected-costs-of-generating-electricity-2020* 19, 56, 64

IEA (2020*b*), 'Renewables 2020: Analysis and forecast to 2025'. 2

Imbens, G. W. (2004), 'Nonparametric estimation of average treatment effects under exogeneity: A review', *Review of Economics and statistics* **86**(1), 4–29. 171

IPCC (2014), 'Climate change 2014: Synthesis report. contribution of working groups i, ii and iii to the fifth assessment report of the intergovernmental panel on climate change', *United Nations Intergovernmental Panel on Climate Change* . 13

Jensen, C. U., Panduro, T. E. & Lundhede, T. H. (2014), 'The vindication of don quixote: The impact of noise and visual pollution from wind turbines', *Land Economics* **90**(4), 668–682. 3, 5, 33, 34, 36, 38, 42, 46, 49, 103

Jensen, C. U., Panduro, T. E., Lundhede, T. H., Nielsen, A. S. E., Dalsgaard, M. & Thorsen, B. J. (2018), 'The impact of on-shore and off-shore wind turbine farms on property prices', *Energy Policy* **116**, 50–59. 3, 5, 33, 34, 36, 37, 38, 42, 44, 49

Kreif, N., Grieve, R., Hangartner, D., Turner, A. J., Nikolova, S. & Sutton, M. (2016), 'Examination of the synthetic control method for evaluating health policies with multiple treated units', *Health economics* **25**(12), 1514–1528. 225

Kyoto Protocol (1997).
URL: *https://unfccc.int/resource/docs/convkp/kpeng.pdf* 14

Lan, H., Gou, Z. & Yang, L. (2020), 'House price premium associated with residential solar photovoltaics and the effect from feed-in tariffs: A case study of southport in queensland, australia', *Renewable Energy* **161**, 907–916. 3, 10, 11, 126, 160, 163, 167, 168, 173, 208, 249

Lang, C., Opaluch, J. J. & Sfinarolakis, G. (2014), 'The windy city: Property value impacts of wind turbines in an urban setting', *Energy Economics* **44**, 413–421. 3, 5, 33, 34, 36, 38, 43, 44, 45, 50, 51, 65, 100, 119, 122, 124, 125, 126, 150

Liao, W.-C. & Wang, X. (2012), 'Hedonic house prices and spatial quantile regression', *Journal of Housing Economics* **21**(1), 16–27. 27

Littlecott, C. (2016), 'Uk coal phase out: the international context', *E3G Briefing Paper* **9**. 157

Ma, C., Polyakov, M. & Pandit, R. (2016), 'Capitalisation of residential solar photovoltaic systems in western australia', *Australian Journal of Agricultural and Resource Economics* **60**(3), 366–385. 3, 10, 160, 163, 165, 166, 173, 207, 208, 248

McCarthy, I. & Balli, H. O. (2014), 'Windfarms and residential property values', *International Journal of Strategic Property Management* **18**(2), 116–124. 3, 5, 33, 34, 36, 37, 38, 41, 44, 45, 46, 49, 124, 150, 154

Monson, M. (2009), 'Valuation using hedonic pricing models'. 26, 65, 119, 163

Morgan, P. L., Frisco, M. L., Farkas2, G. & Hibel, J. (2017), 'Republication of "a propensity score matching analysis of the effects of special education services"', *The Journal of Special Education* **50**(4), 197–214. 185

OECD (2012), 'Linking renewable energy to rural development'. 62, 63, 64

OFGEM (2019), 'About the smart export guarantee (seg)'.
URL: *https://www.ofgem.gov.uk/environmental-programmes/smart-export-guarantee-seg/about-smart-export-guarantee-seg* 19, 177

OFGEM (2020*a*), 'About the fit scheme'.
URL: *https://www.ofgem.gov.uk/environmental-programmes/fit/about-fit-scheme* 18, 177

OFGEM (2020*b*), 'Electricity generation mix by quarter and fuel source (gb)'. 19, 56, 114, 157, 160

OFGEM (2020*c*), 'Feed-in tariff installation report'. 17, 175

OFGEM (2020*d*), 'What is the nffo?'.
URL: *https://www.ofgem.gov.uk/environmental-programmes/nffo* 16

OFGEM (2020*e*), 'What is the ro?'.
URL: *https://www.ofgem.gov.uk/environmental-programmes/ro/energy-suppliers* 16

ONS (n.d.), 'The decoupling of economic growth from carbon emissions: Uk evidence kernel description'. 64, 111, 113

Oreskes, N. (2004), 'The scientific consensus on climate change', *Science* **306**(5702), 1686–1686. 13

Paris Agreement (2015).
URL: *https://unfccc.int/sites/default/files/english_paris_agreement.pdf* 14

Parmeter, C. F. & Pope, J. C. (2013), Quasi-experiments and hedonic property value methods, *in* 'Handbook on experimental economics and the environment', Edward Elgar Publishing. 100, 153

Plott, C. R. & Zeiler, K. (2005), 'The willingness to pay-willingness to accept gap, the" endowment effect," subject misconceptions, and experimental procedures for eliciting valuations', *American Economic Review* **95**(3), 530–545. 20

PMO (2021), 'Uk enshrines new target in law to slash emissions by 78% by 2035'. 2, 14, 15, 18, 56

Powell, J. (2017), 'Scientists reach 100% consensus on anthropogenic global warming', *Bulletin of Science, Technology & Society* **37**(4), 183–184. 13

Pride, D., Little, J. & Mueller-Stoffels, M. (2018), 'The value of residential energy efficiency in interior alaska: A hedonic pricing analysis', *Energy policy* **123**, 450–460. 27

Qiu, Y., Wang, Y. D. & Wang, J. (2017), 'Soak up the sun: Impact of solar energy systems on residential home values in arizona', *Energy Economics* **66**, 328–336. 3, 10, 11, 160, 163, 168, 173, 207, 208, 247

Randolph, J. J. & Falbe, K. (2014), 'A step-by-step guide to propensity score matching in r.', *Practical Assessment, Research & Evaluation* **19**. 170

Richter, M. K. (1966), 'Revealed preference theory', *Econometrica: Journal of the Econometric Society* pp. 635–645. 20

Rosen, A. M. (2015), 'The wrong solution at the right time: The failure of the kyoto protocol on climate change', *Politics & Policy* **43**(1), 30–58. 14

Rosen, S. (1974), 'Hedonic prices and implicit markets: product differentiation in pure competition', *Journal of political economy* **82**(1), 34–55. 20, 21, 23, 107, 163

Rosenbaum, P. R. & Rubin, D. B. (1983), 'The central role of the propensity score in observational studies for causal effects', *Biometrika* **70**(1), 41–55. 170, 171, 184

Rubin, D. B. (1978), 'Bayesian inference for causal effects: The role of randomization', *The Annals of statistics* pp. 34–58. 170

Schaeffer, Y. & Dissart, J.-C. (2018), 'Natural and environmental amenities: a review of definitions, measures and issues', *Ecological Economics* **146**, 475–496. 27, 118

Schläpfer, F., Waltert, F., Segura, L. & Kienast, F. (2015), 'Valuation of landscape amenities: A hedonic pricing analysis of housing rents in urban, suburban and periurban switzerland', *Landscape and Urban Planning* **141**, 24–40. 26

Sims, S. & Dent, P. (2007), 'Property stigma: wind farms are just the latest fashion', *Journal of Property Investment & Finance* . 28, 33, 34, 36, 37, 42, 44, 56, 104

Sims, S., Dent, P. & Oskrochi, G. R. (2008), 'Modelling the impact of wind farms on house prices in the uk', *International Journal of Strate-*

gic Property Management **12**(4), 251–269. 33, 34, 36, 37, 38, 42, 44, 45, 49, 121, 189

SolarGIS (2020), 'Solar resource maps and gis data'. viii, 172, 174

Song, H., Dwyer, L., Li, G. & Cao, Z. (2012), 'Tourism economics research: A review and assessment', *Annals of tourism research* **39**(3), 1653–1682. 20

Stetler, K. M., Venn, T. J. & Calkin, D. E. (2010), 'The effects of wildfire and environmental amenities on property values in northwest montana, usa', *Ecological Economics* **69**(11), 2233–2243. 26

Sunak, Y. & Madlener, R. (2016), 'The impact of wind farm visibility on property values: A spatial difference-in-differences analysis', *Energy Economics* **55**, 79–91. 3, 5, 33, 34, 36, 37, 38, 43, 46, 47, 50, 52, 61, 100, 102, 103, 122, 123, 124, 125, 126, 150, 152

Sunak, Y. & Madlener, R. (2017), 'The impact of wind farms on property values: A locally weighted hedonic pricing model', *Papers in Regional Science* **96**(2), 423–444. 3, 5, 33, 34, 36, 38, 43, 46, 50, 52, 59, 100, 102

Tapsuwan, S., Polyakov, M., Bark, R. & Nolan, M. (2015), 'Valuing the barmah–millewa forest and in stream river flows: A spatial heteroskedasticity and autocorrelation consistent (shac) approach', *Ecological Economics* **110**, 98–105. 26

United Nations Framework Convention on Climate Change (1992).
URL: *https://unfccc.int/files/essential_background/background_publications_htmlpdf/applicati*
14

Vyn, R. J. & McCullough, R. M. (2014), 'The effects of wind turbines on property values in ontario: does public perception match empirical evidence?', *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie* **62**(3), 365–392. 3, 5, 33, 34, 36, 37, 38, 43, 44, 45, 49, 122, 123, 125

Watkiss P, Downing T, D. J. (2010), 'Adaptcost project: Analysis of the economic costs of climate change adaptation in africa'. 13

Wee, S. (2016), 'The effect of residential solar photovoltaic systems on home value: A case study of hawai 'i', *Renewable Energy* **91**, 282–292. 3, 10, 160, 163, 165, 166, 173, 195, 208, 246

WGCCC (2021), 'Wales commits to net zero by 2050, but sets out ambitions to get there sooner'. 2

Xu, X.-l., Chen, H. H., Feng, Y. & Tang, J. (2018), 'The production efficiency of renewable energy generation and its influencing factors: Evidence from 20 countries', *Journal of Renewable and Sustainable Energy* **10**(2), 025901. 175

Xu, Y. (2017), 'Generalized synthetic control method: Causal inference with interactive fixed effects models', *Political Analysis* **25**(1), 57–76. 225