Reykjavík University

In collaboration with ETH Zürich

Analysis of Wind Farms Impact on Property Prices Using Difference-in-Differences

A Case Study of Southern Great Britain

Daniel Z. El Deen K. Al Hennaw
Franz Ísak Ingólfsson

Project Supervisor
Dr. Christoph Lohrmann

Project Co-supervisor
Dr. Alena Lohrmann

Project Examiner
Dr. Gylfi Þór Guðmundsson

July 1, 2026

Abstract

This thesis investigates whether onshore wind turbines are associated with changes in residential property prices in England and Wales, with a particular focus on proximity, visual exposure, and shadow flicker. Using geocoded property transactions from 2011–2019, wind turbine commissioning data, terrain-aware viewshed rasters, and simulated shadow-flicker exposure, the study applies Difference-in-Differences models to estimate how property prices vary with distance to turbines and with direct physical exposure to turbine-related effects. The results suggest that properties located closest to turbines experience the largest price discounts, with the strongest proximity effect observed within 1 km. Visibility also appears to matter independently of distance, especially when many turbines are visible from a property. Shadow flicker provides an additional exposure channel, helping distinguish general proximity effects from more specific physical and visual impacts. However, event-study results show clear pre-trend violations, meaning that the dynamic estimates should be interpreted cautiously. Overall, the thesis contributes evidence that both distance and actual visibility are relevant when assessing the local property-market effects of wind farm development.

In collaboration with:

Table of Contents

1 Collaborators **1**

2 Research Overview **2**
2.1 Background . . . 2
2.2 Objectives . . . 2
2.3 Research Questions . . . 3
2.4 Contributions and Novelty . . . 3
2.5 Stakeholders . . . 4
2.6 Tech Stack . . . 4

3 Literature Review **5**
3.1 Visual Intrusion and Viewshed Analysis . . . 5
3.2 Noise and Shadow Flicker Exposure . . . 6
3.3 Planning Decisions and NIMBYism . . . 6

4 The Data **8**
4.1 Property Transactions . . . 8
4.2 Wind Turbine Installations . . . 9
4.3 Viewshed Rasters . . . 11
4.4 Digital Surface Model . . . 12

5 System Design **13**
5.1 Preprocessing and Stage Sequencing . . . 13
5.2 Three Analysis Branches . . . 13

6 Data Preprocessing **15**
6.1 Turbine Cleaning and Wind Farm Construction . . . 15
6.2 Transaction Cleaning and Geocoding Merging . . . 15
6.3 Final Model-Ready Samples . . . 16

7 Methodology **17**
7.1 Difference-in-Differences . . . 17
7.2 Two-Way Fixed Effects . . . 18
7.3 Estimating Equation . . . 19
7.4 Event Study Design . . . 20
7.5 Staggered Adoption and Dynamic Effects . . . 20
7.6 Inference and Standard Errors . . . 21

8 Models & Results **22**
8.1 Proximity Models . . . 22
8.1.1 ProxGrad — Proximity Gradient (Baseline Model) . . . 22
8.1.2 Proximity Event Study . . . 26
8.1.3 Results Summary — Proximity Models . . . 28
8.2 Visibility Models . . . 29
8.2.1 Visibility DiD . . . 29
8.2.2 Disentangling Visibility from Distance . . . 32
8.2.3 Visibility Event Study . . . 38
8.2.4 Results Summary — Visibility Models . . . 41
8.3 Shadow-Flicker Models . . . 42
8.3.1 Shadow-Flicker DiD . . . 42

9 Conclusion **46**
9.1 Summary of Main Findings . . . 46
9.2 Answering the Research Questions . . . 47
9.3 Limitations . . . 48
9.4 Final Remarks . . . 48
9.5 Future Work . . . 49

10 Bibliography **51**

CHAPTER 1

Collaborators

This project is conducted in collaboration with ETH Zürich, Switzerland, one of the world's leading technical universities.

The collaboration emerged through Dr. Christoph Lohrmann, an assistant professor and head of the MSc in Data Science at Reykjavík University, who supervises this project, and Dr. Alena Lohrmann, a postdoctoral researcher at ETH Zürich who serves as co-supervisor.

This work aligns with the goals of the WIMBY project (Wind In My Backyard)[1], which is an EU-funded initiative that supports the adoption and acceptance of wind power across Europe. The project's main objective is to mitigate the "Not In My Backyard" (NIMBY) effect, meaning the local opposition that often threatens wind farm deployment despite broad public support for renewable energy. WIMBY aims to develop innovative models to assess wind power impacts, including visual disruption, noise, and shadow flicker, and translates these into practical information that stakeholders and citizens can use for informed decision-making. By facilitating societal engagement and providing tools for comprehensible impact assessments, the project aims to contribute to the EU's decarbonization strategy.

Our work builds on and extends previous research conducted at ETH Zürich, which examined the relationship between wind turbine visibility and residential property values using regression models. We have been provided with spatial data on property prices in England and Wales, wind turbine location data from the WIMBY project, and pre-computed visibility metrics. These resources serve as a foundation and reference point for our investigation.

This cross-institutional collaboration enables us to contribute to a broader European research effort while applying rigorous methods to the specific context of property markets in the south of Great Britain.

CHAPTER 2

Research Overview

2.1 Background

Over the past decade, transitioning away from fossil fuels has been one of the defining challenges of our time, and many efforts have focused on the transition to renewable energy sources. Among these, wind energy emerged as one of the most promising renewable alternatives. Although most people support the usage of wind turbines as a renewable energy source, support can become more mixed when wind farms are proposed close to residential areas[2], where residents may support renewable energy in principle but oppose nearby developments because of concerns such as visual impact, noise, shadow flicker, and possible effects on property values[3].

These factors can create real problems for energy planners. Community resistance may result in project delays, increased costs, or difficulty in securing approval for new wind farm developments[4]. At the same time, meeting climate targets requires a substantial expansion of wind energy capacity. This makes it important to understand what the actual impacts are, whether residents' concerns are reflected in property prices, and if so, to what extent.

2.2 Objectives

The primary objective of this study is to estimate the causal effect of wind farms/turbines on property prices in England and Wales using DiD algorithm. This geographic scope was chosen because of the available property transaction data, wind turbine location data, and pre-computed visibility metrics were available for this region. England and Wales also provide a suitable case study because they contain a large residential property market and a substantial number of onshore wind turbines, making it possible to study whether turbine exposure is reflected in nearby property prices.

The project considers three main channels through which wind turbines may affect residential property values: proximity, visual intrusion, and shadow flicker.

Proximity refers to the effect of living near a wind turbine or wind farm, regardless of whether the turbine is visible from the property. This channel may capture physical concerns or local disruption, but also perception-based effects, where buyers value a property less simply because they know it is close to a turbine.

Visual intrusion refers to the effect of turbines being visible from a property. This is different from proximity because two properties may be located at similar distances from a turbine but have very different levels of visual exposure depending on terrain, landscape structure, and line-of-sight conditions. This channel is central to the project because the analysis uses visibility measures to test whether actually seeing turbines has an effect on property prices beyond distance alone.

Shadow flicker occurs when sunlight passes through rotating turbine blades and creates moving shadows on nearby buildings or land. Its intensity depends on the position of the turbine, the position of the sun, the time of year, and whether a property lies within the affected shadow path.

2.3 Research Questions

The main research question is:

> *How do onshore wind farms affect residential property prices in England and Wales, and how much of that effect is attributable to visibility rather than proximity alone?*

The report addresses four supporting questions:

- Does wind turbine *visibility* have a statistically significant effect on residential property prices in England and Wales?
- Is the effect on property prices driven by the *actual visibility* of wind turbines, or by *proximity* alone, that is, does living near a turbine affect prices even when the turbine is not visible?
- Does the *shadow flicker* from wind turbines have an additional effect on property prices beyond visual impact?
- Do property prices recover over time after wind turbine installation?

2.4 Contributions and Novelty

This project contributes to the ongoing ETH Zürich and WIMBY research by extending the analysis of wind farm impacts on residential property prices in England and Wales. Its main novelty is that it considers proximity, visibility, and shadow flicker as separate channels through which turbines may affect property values, while applying a DiD framework to estimate changes before and after turbine installation. This allows the study to move beyond simple distance-based comparisons and examine whether observed price effects are more closely related to actual visual exposure or other forms of local disturbance.

2.5 Stakeholders

The primary audience for this work is academic. The project is conducted as part of a BSc thesis at Reykjavík University and contributes to the broader research agenda of at ETH Zürich. The results, methodology, and codebase are intended to be used and extended by researchers at both institutions.

Beyond academia, the findings may be of practical interest for energy planning authorities, wind farm developers, and local governments who must weigh the economic externalities of turbine installations against renewable energy targets. Whether and how the model is adopted for applied use will be determined primarily by the ETH Zürich research team, who oversee the wider initiative and engagement with industry and policy stakeholders.

2.6 Tech Stack

The project was implemented in Python, using a small set of core libraries for data processing, geospatial analysis, statistical modelling, and visualisation. The main libraries used were:

- **Data processing and numerical analysis:** Pandas, NumPy.
- **Geospatial analysis:** GeoPandas, Rasterio, Shapely, PyProj.
- **Visualisation:** Matplotlib.
- **Statistical modelling:** Statsmodels, with additional use of scikit-learn for spatial clustering and nearest-neighbour calculations.

Version control was managed using Git and GitHub. Development and exploratory analysis were carried out using Jupyter Notebooks and Visual Studio Code.

About AI usage:
Large language models, mainly ChatGPT and Claude, were used as supporting tools during the project. They assisted with debugging, code review, LaTeX formatting, and reviewing written explanations, but all code, results, interpretations, and final decisions were manually made, checked and verified by the authors.

CHAPTER 3

Literature Review

The economic impact of wind energy infrastructure on residential property markets is a well-established field of study, typically framed within the context of local externalities. Early research often relied on simple Euclidean distance from property to turbine as a proxy for disturbance. However, recent advancements in spatial econometrics and computational modeling have shifted the focus toward quantifying specific physical disturbances: visual intrusion, noise pollution, and shadow flicker.

3.1 Visual Intrusion and Viewshed Analysis

Visual impact is frequently cited as the primary driver of local opposition and subsequent property value depreciation. In a seminal study of England and Wales [5], utilized a quasi-experimental design to isolate the visual component of wind farm externalities. By comparing properties with a direct line-of-sight to those where turbines were obscured by terrain, Gibbons found price reductions of 5–6% for properties with a visible wind farm within 2 km, declining to under 2% between 2–4 km and near zero by 8–14 km. Interestingly, locations close to wind farms but with lower turbine visibility showed smaller price reductions, suggesting the effect is driven by visual exposure rather than proximity alone. This supports using viewshed modeling over proximity-based distance thresholds.

More recently, there was a large-scale analysis in the United States showing that visibility leads to statistically significant negative impacts on home values in close proximity (under 8 km) [6]. However, their findings suggest that the price penalty does not stay constant. As wind turbines become a more familiar part of the landscape, the discount appears to shrink over time [6]. This could reflect people simply getting used to having turbines nearby or a broader shift in how communities feel about renewable energy infrastructure.

3.2 Noise and Shadow Flicker Exposure

While visual intrusion covers the esthetic cost, noise and shadow flicker represent distinct physical disturbances that occur at much closer ranges. Chen[7] addressed this issue in the Netherlands by moving beyond binary visibility measures and using continuous exposure metrics. Using simulation models for wind turbine noise (WTN) and shadow flicker (SF), Chen observed that properties exposed to severe impacts, defined as noise levels $\geq$ 45 dB or more than 30 hours of shadow flicker per year experienced the sharpest price declines of approximately 3.24%. Notably, this study found evidence of an anticipation effect, where prices began to decline 6–8 years before a project was even commissioned.

The specific impact of shadow flicker is further investigated by the Danish market. Using a large dataset of 2.4 million transactions and 6,878 wind turbines, they distinguished between general proximity and the shadow flicker effect caused by rotor blades. Their results indicate that while large modern modern turbines cause a baseline reduction of up to 12% for nearby homes, properties specifically situated in the shadow flicker path experience an additional 8.1% loss in value.[8]

3.3 Planning Decisions and NIMBYism

The spatial distribution of these externalities is often a byproduct of the planning process itself. Jarvis [9] explores the "NIMBY" (Not In My Backyard) phenomenon in the UK, evaluating how local opposition influences where turbines are eventually built. Jarvis notes that while wind projects impose real local costs, revealed through hedonic price drops, the planning system's tendency to reject projects in higher-value areas can lead to a misallocation of renewable resources. By using a DiD approach across three decades of UK data, Jarvis demonstrates that the local costs are significant enough to sway planning outcomes.

Summary

The literature shows that wind turbines can affect property values through several channels. Earlier studies often measured exposure using simple distance to nearby turbines, while more recent work uses more detailed spatial measures such as visibility, noise, and shadow flicker. The reviewed studies suggest that visual exposure is an important channel, but also that physical disturbances such as noise and shadow flicker may matter for properties located very close to turbines.

For this thesis, the main implication is that distance alone is an incomplete measure of wind-farm exposure. Two properties at the same distance from a turbine may experience different impacts depending on whether the turbines are visible from the property. This motivates the focus of this project on separating proximity effects from visibility effects when estimating the relationship between wind farms and residential property prices.

Paper	Description	Priority	Methodology
Gibbons (2015) [5]	The definitive UK study using Land Registry data and a spatial DiD framework.	High	Quasi-experimental
Andersen & Hener (2025) [8]	Detailed shadow flicker modelling with 2.4 million transactions in Denmark.	High	Shadow-mapping, HPM
Guo et al. (2024) [6]	Found visual effect on property values is small and diminishing in space and time.	High	Geospatial DiD
Dröes & Koster (2016) [10]	Renewable energy and negative externalities: the effect of wind turbines on house prices in the Netherlands.	High	—
Chen (2025) [7]	MSc thesis examining wind turbine noise and shadow flicker impacts on 3.6 million Dutch property transactions, finding a 3.24% decline for severe exposure and recovery after ∼10 years.	High	DiD, Fixed Effects
Schütt (2024) [11]	Meta-regression analysis synthesising 25 peer-reviewed studies, comparing DiD vs. HPM approaches.	Moderate	DiD vs. HPM
Brunner et al. (2024) [12]	US study on anticipatory price drops (pre-construction), finding an 11% decline after announcement that recovers over 9 years.	Moderate	—
Ebersole (2021/2025) [9]	Identifies "anticipation stigma" — price drops during the planning stage.	Moderate	—
Sims, Dent & Oskrochi (2010) [13]	Early UK study modelling the impact of wind farms on house prices.	Moderate	Hedonic Pricing Model
Jensen, Panduro & Lundhede (2018) [14]	The impact of on-shore and off-shore wind turbine farms on property prices.	Low	—

CHAPTER 4

The Data

This study combines four spatial datasets: *residential property transactions, wind turbine installations, cumulative turbine viewshed rasters*, and *a digital surface model*. The property data define the housing market observations (the outcome variable), the turbine data define the location and timing of wind farm commission (treatment event), the viewshed rasters provide a terrain-aware measure of whether turbines are visible from each property location, and the digital surface model provides ground elevation data used to compute turbine shadow projections. All spatial data are in EPSG:3035.[1]

4.1 Property Transactions

The property transaction dataset contains approximately ~ 5.7 million residential sales in England and Wales from 2011 to 2019, sourced from University College London[15]. Each record represents one property sale and includes the attributes shown in Table 4.1. These variables are used both to define the outcome variable and to control for basic differences between properties.

Attribute	**Description**	**Type**
`transactionid`	Unique transaction identifier	String
`postcode`	Property postcode	String
`price`	Sale price (£)	Integer
`dateoftransfer`	Date of sale	Date
`propertytype`	Detached / semi-detached / terraced / flat	Categorical
`oldnew`	New build or existing stock	Categorical
`tfarea`	Total floor area (m^2)	Float
`numberrooms`	Number of rooms	Integer
`priceper`	Price per m^2 (raw CSV only)†	Float

†`priceper` is present in the raw CSV but is not retained after preprocessing; the outcome variable is *log*(`price`) instead.

Table 4.1: Attributes in the property transaction dataset.

[1] A European projected coordinate system that represents locations in metres and allows distances and areas to be measured consistently.

4.2 Wind Turbine Installations

The wind turbine dataset contains onshore and offshore turbine locations across Europe from the WIMBY project[1], based on data from DTU Wind Energy[16]), The raw dataset covers 38 countries and 84,488 turbines. For each turbine, the dataset includes the attributes shown in Table 4.2. The commissioning date is especially important because the empirical strategy depends on comparing property prices before and after nearby turbines are installed.

Attribute	**Description**	**Type**
Longitude	Longitude (decimal degrees east)	Float
Latitude	Latitude (decimal degrees north)	Float
ISO_code	Country identifier	String
Offshore	Onshore (0) or offshore (1)	Integer
Hub_height_m	Hub height (m)	Float
P_rated_kW	Rated power (kW)	Float
Diameter_m	Rotor diameter (m)	Float
Commissioning_date	Commissioning date (YYYY-MM-DD)	Date

Table 4.2: Attributes in the Europe turbine dataset.

The turbine data is further summarized in Table 4.3, which reports the main physical characteristics used to describe the turbines in the study area.

Characteristic	**Min**	**Median**	**Max**
Hub height (m)	24.0	67.0	125.0
Rotor diameter (m)	21.0	82.0	167.0
Rated power (kW)	100	2,050	7,000
Tip height (m)	34.5	110.0	189.5

Table 4.3: Summary statistics of turbine physical characteristics ($total = 5{,}441$).

A visualization of the spatial distribution of property transactions and onshore turbines is shown in Figure 4.1. Figure 4.2 then shows the distribution of onshore turbines by year of commission. This is important for the empirical strategy because there is a notable increase in turbine commissioning during the study period, between 2011 and 2019, which provides useful variation in treatment timing for the DiD models.

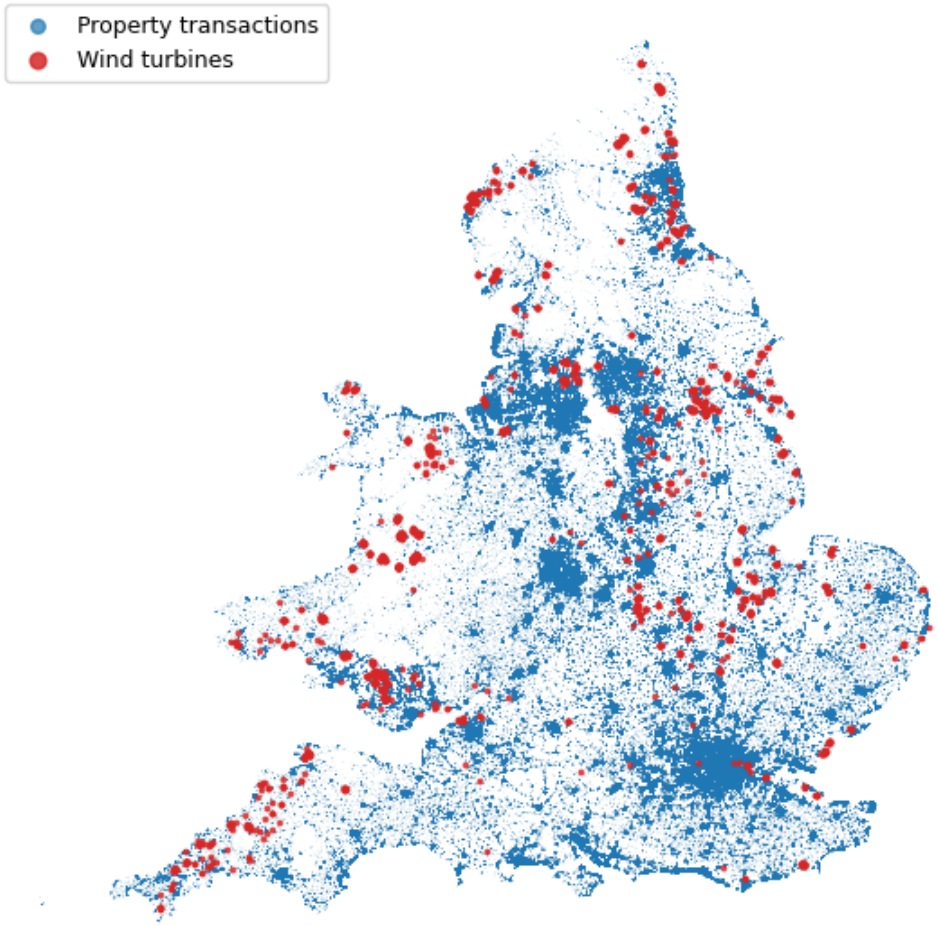

Figure 4.1: Spatial distribution of property transactions and onshore turbines across England and Wales.

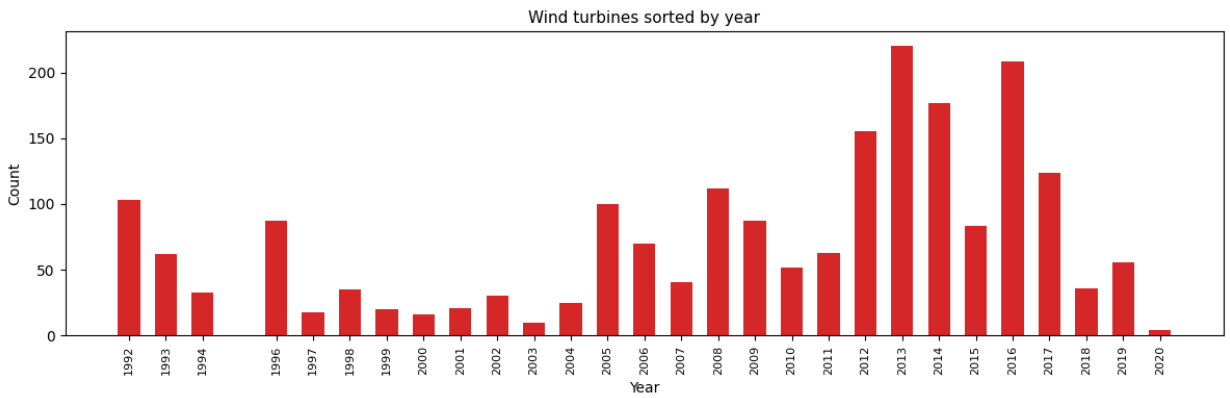

Figure 4.2: Onshore turbines by commissioning year in Great Britain ($total = 5{,}441$).

4.3 Viewshed Rasters

The visibility data consist of 28 pre-computed cumulative annual viewshed rasters ETH Zürich, covering the years 1992–2020. Each raster is projected to EPSG:303, where each pixel represents a $100\,m \times 100\,m$ resolution, and stores the total number of turbines visible from that location at that year.

Sampling the raster at each property's location is used to assign each transaction a turbine visibility count for the year in which the sale occurred, from which both a binary treatment flag and a continuous dose are derived directly.

Table 4.4 summarizes how turbine visibility coverage changes across England and Wales during the transaction period and the following year.

Year	Pixels with visibility	% of total	Max visible turbines	Mean where visible
2011	2,710,722	3.49	174	7.8
2012	3,220,543	4.15	174	7.9
2013	3,889,772	5.01	174	8.3
2014	4,480,127	5.77	174	8.3
2015	4,749,875	6.11	174	8.3
2016	4,996,590	6.43	174	8.9
2017	5,221,237	6.72	174	9.2
2018	5,275,368	6.79	174	9.2
2019	5,352,301	6.89	174	9.3
2020	5,360,232	6.90	174	9.3

Table 4.4: Visibility coverage across England and Wales, 2011–2020.

The spatial expansion of turbine visibility between the beginning and end of the study period is shown in Figure 4.3. The same growth pattern is summarized over time in Figure 4.4, which shows the increasing share of pixels with turbine visibility.

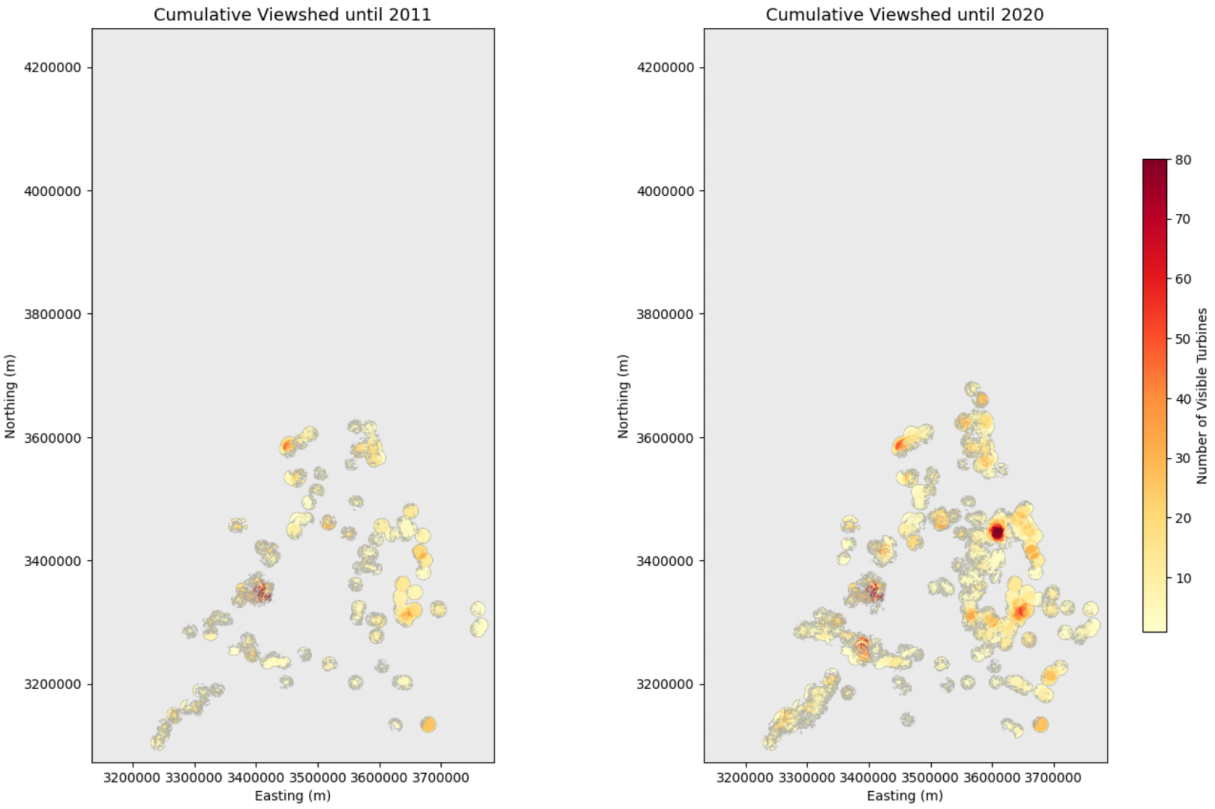

Figure 4.3: Expansion of turbine visibility between 2011 and 2020.

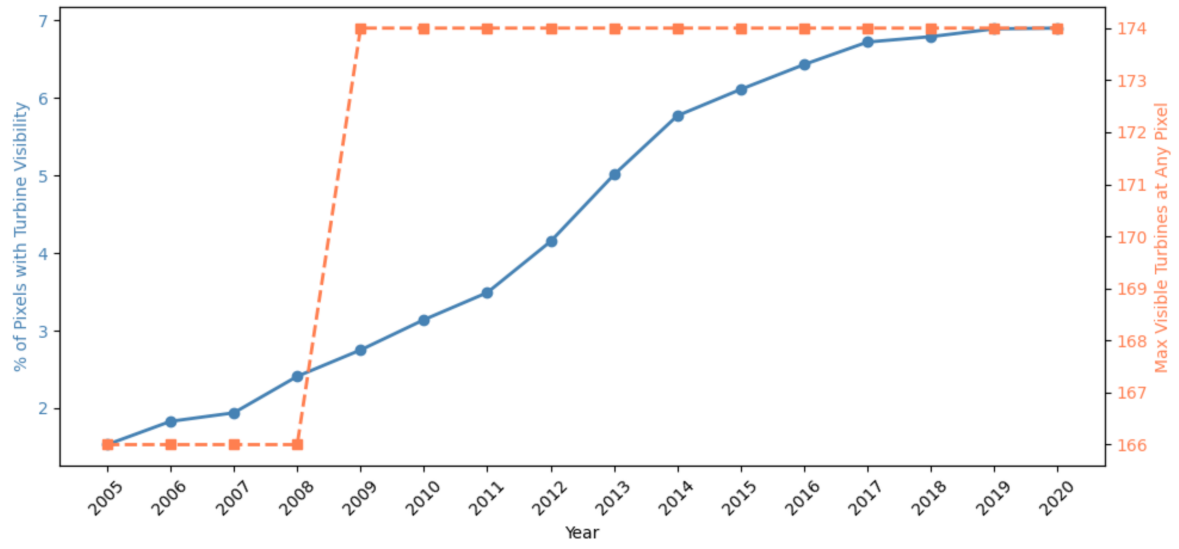

Figure 4.4: Growth of turbine visibility coverage over time.

4.4 Digital Surface Model

The Digital Surface Model (DSM) is a gridded raster covering Great Britain at 100 m × 100 m resolution in EPSG:3035, provided by the project supervisor. Each pixel encodes the ground surface elevation in metres. In this study, the DSM serves as a terrain input to the shadow flicker simulation, which requires surface elevation data to project turbine blade shadows realistically onto the surrounding landscape.

CHAPTER 5

System Design

The analysis is implemented as a step-based pipeline where each stage specifies its input files, output files, and the source files it depends on. The pipeline spans four stages (preprocessing, proximity, visibility, and shadow flicker), each isolated so that no stage depends on another completing successfully. A single configuration module centralises all study parameters and file paths, so a change to any constant automatically propagates and only the affected steps rerun rather than the full pipeline. Figure 5.1 gives a high-level overview of the data flow.

5.1 Preprocessing and Stage Sequencing

The preprocessing stage ingests the four raw data sources described in Chapter 4 and produces two model-ready artefacts: a cleaned turbine and wind-farm table and a cleaned transaction table. Both outputs feed the proximity and visibility branches directly. The shadow flicker branch runs last because it additionally draws on the visibility property dataset produced by the visibility stage.

5.2 Three Analysis Branches

The three treatment channels (proximity, visibility, and shadow flicker) are introduced conceptually in Chapter 2. At the implementation level each channel is a separate pipeline branch consuming the shared clean data layer. Shadow flicker differs from the other two in that it calls `wimby-sf` [17] to build per-farm annual shadow-flicker rasters from the DEM before running its models.

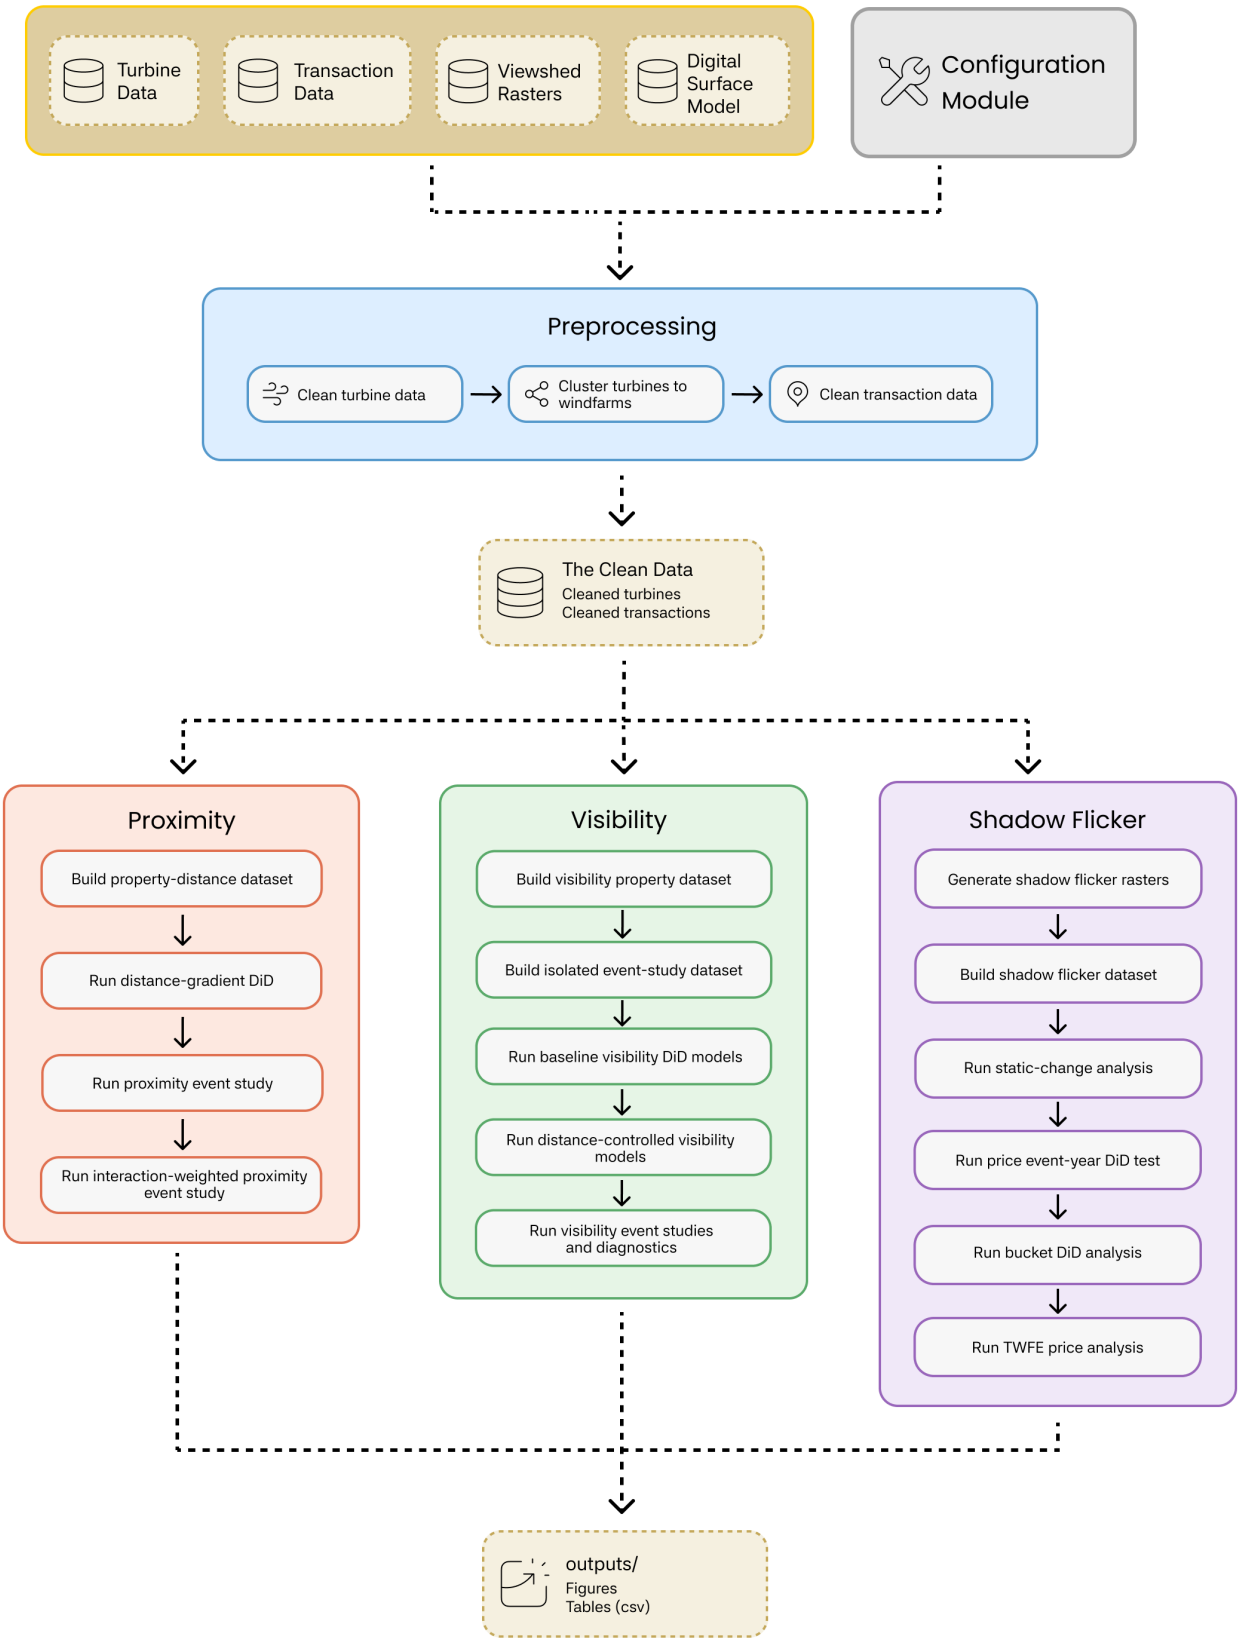

Figure 5.1: High-level data flow of the analysis pipeline. Raw inputs are normalised once in preprocessing and then consumed by three analysis branches, proximity, visibility, and shadow flicker.

CHAPTER 6

Data Preprocessing

The preprocessing pipeline cleans converts the raw property, and turbine datasets into model-ready samples. The main goal is to keep the data consistent across the proximity and visibility analyses while removing observations that cannot be reliably linked to turbine exposure.

6.1 Turbine Cleaning and Wind Farm Construction

1. The turbine data are first restricted to relevant GB onshore turbines.
2. Turbines are then grouped into estimated wind farms using DBSCAN clustering with a 1 km radius. This creates 835 estimated wind farms in the cleaned turbine sample.
3. Missing commissioning dates are handled using information from nearby turbines in the same estimated wind farm. In total, 15 missing dates are filled using the modal commissioning date within the farm. A further 13 turbines cannot be matched to a dated wind farm and remain undated. These undated turbines are not used as treatment units, but they are retained to define exclusion zones so that nearby transactions are not included in the study.

6.2 Transaction Cleaning and Geocoding Merging

1. Property transaction records are linked to their geocoded spatial locations using the transaction identifier (`transactionid`). Records without a matching geocoded location are dropped because they cannot be assigned distance or visibility measures.
2. The property sample is then cleaned by removing implausible values: prices below £1,000 or above £10,000,000, floor areas outside 0–1,000 m^2, and properties with fewer than one room.
3. The maintained analysis uses a study horizon to keep exclusively the transactions around dated turbines. The horizon is **11 km for proximity** models and **13 km for visibility-distance** models (these decisions will be explained later in the section of visibility models with distance control 8.2.2). While transactions within 13 km of an undated turbine are removed to avoid uncertain treatment timing.

4. GB turbines just outside the England and Wales administrative boundary are retained when they are within range of properties in northern England, since they can still affect those properties.

5. **For the proximity analysis**, each transaction is assigned its distance to the nearest dated turbine and the commissioning year of that turbine.

6. **For the visibility analysis**, each transaction is matched to the viewshed raster for its sale year. The raster value at the property's 100 m pixel gives the number of visible turbines, stored as `n_visible`. From this, the pipeline derives a binary visibility indicator, the first year in which the property becomes visible to any turbine, and visibility bands such as 1–5, 6–10, 11-15, and 16+ visible turbines.

7. Sale prices are strongly right-skewed, with a small number of very expensive transactions as shown in Figure 6.1. To reduce the influence of these extreme values and make the outcome more suitable for regression modelling, the final outcome variable is defined as log(price). This also makes the estimated coefficients easier to interpret, since they can be read approximately as percentage changes in property prices, and that will be our **outcome variable**.

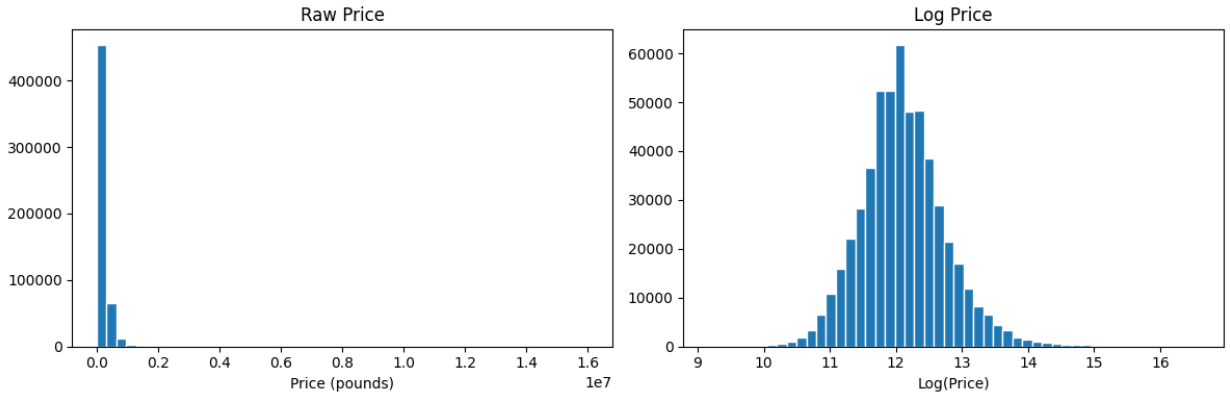

Figure 6.1: Raw price and log-price distributions for the transaction data.

6.3 Final Model-Ready Samples

The final model-ready datasets are summarized in Table 6.1. These maintained samples form the basis for the proximity and visibility analyses used in the rest of the report.

Table 6.1: Preprocessing outputs used by the maintained models.

Output	Rows / units
Clean transactions	2,949,650
Proximity analysis rows	2,478,283
Visibility analysis rows	2,949,650
Clean turbines	5,441
Estimated wind farms in UK	835

CHAPTER 7

Methodology

This chapter describes the general empirical strategy used throughout the analysis. The purpose is not to present the individual models or their results, which are discussed in the next chapter, but to introduce the common methodological framework behind them.

The analysis studies whether exposure to wind turbines is associated with changes in residential property prices. Since wind turbines are commissioned in different places and at different points in time, the project uses a difference-in-differences framework. This allows property prices before and after turbine exposure to be compared between treated and control properties, while accounting for general price trends and fixed differences between geographic areas.

All models share the same basic logic. They use residential transaction prices as the outcome, include hedonic property controls, account for geographic and year fixed effects, and estimate the relationship between turbine exposure and property prices. The exact definition of turbine exposure differs across the model families and is introduced in the Models & Results (Chapter 8).

7.1 Difference-in-Differences

The main identification strategy is difference-in-differences. The central idea is to compare how prices change over time for properties that become exposed to wind turbines with how prices change over the same period for properties that do not become exposed.

A simple before-and-after comparison would not be enough, because property prices may change over time for many reasons unrelated to wind turbines. For example, national housing market conditions, interest rates, inflation, and local economic conditions may all affect prices. Similarly, a simple comparison between exposed and unexposed properties would also be problematic, because areas near turbines may differ systematically from areas without turbines.

Difference-in-differences addresses this by combining both comparisons. It compares the before-after change for treated properties against the before-after change for control properties. In doing so, it removes time-invariant differences between treated and control areas and also removes common shocks that affect all properties in the same year.

In this study, treatment does not occur at one single point in time for all properties. Wind turbines are commissioned across multiple years, meaning that different properties become exposed at different times. This setting is known as staggered adoption. Each cohort [1] contributes its own before-after comparison, and these comparisons are combined in the regression framework. This staggered treatment structure is illustrated in Figure 7.1, where each cohort becomes

[1] The set of properties that got their first turbine treatment in the same year.

treated in a different calendar year while the never-treated group remains available as a control group.

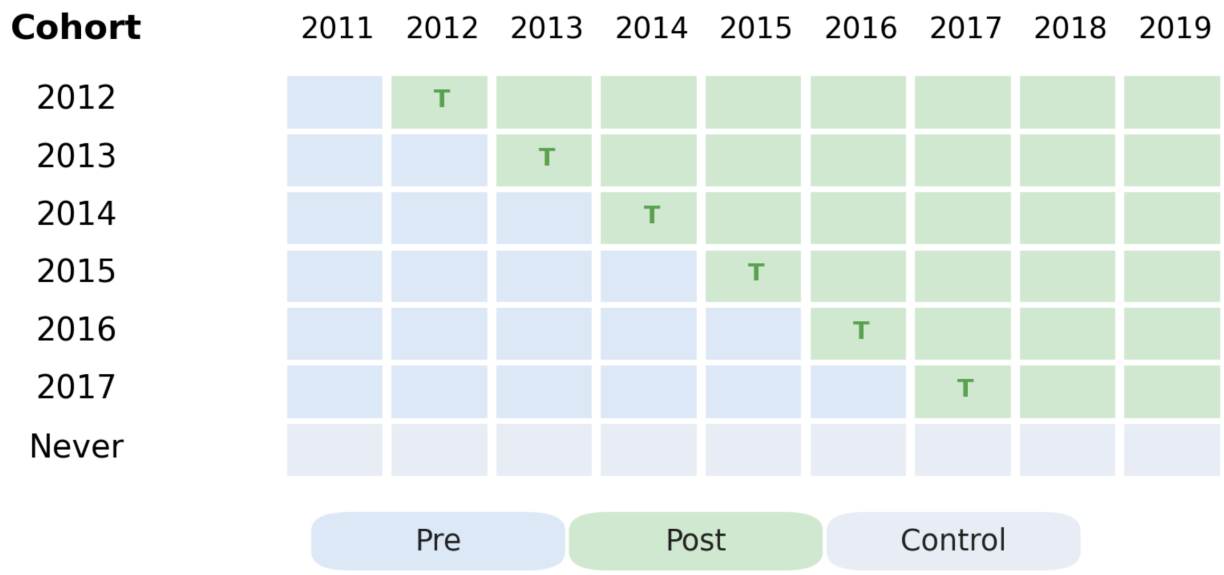

Figure 7.1: Staggered adoption structure used in the difference-in-differences design. Each row represents a treatment cohort, defined by the year in which properties first become exposed to turbines. Blue cells show pre-treatment observations, green cells show post-treatment observations, and the never-treated group provides the comparison (control) group.

7.2 Two-Way Fixed Effects

The static[2] DiD models are estimated using a two-way fixed effects specification. The two dimensions are geographic fixed effects and year fixed effects. Geographic fixed effects account for persistent differences between areas, while year fixed effects account for common changes over time.

This specification estimates one average post-treatment effect. In other words, it summarizes the association between turbine exposure and property prices into a single coefficient. This is useful for estimating the overall relationship, but it does not show whether the effect appears before treatment, immediately after treatment, or gradually over time.

[2]Models that estimate one overall treatment effect, without showing how the effect changes over time before and after turbine installation.

7.3 Estimating Equation

The models estimate the relationship between property prices and turbine exposure using the regression framework summarized in Figure 7.2:

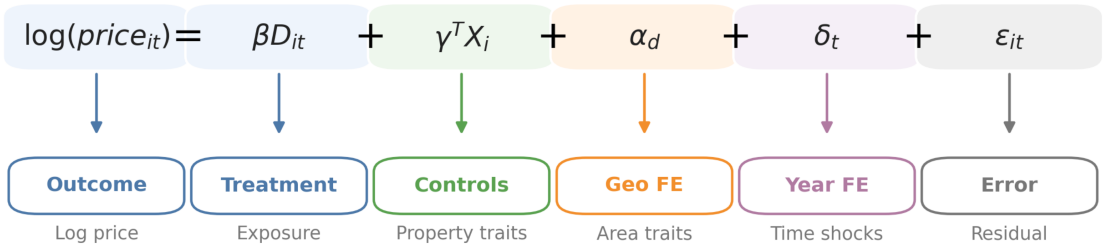

Figure 7.2: Main components of the regression framework used across the models

where i is a property transaction, t is the calendar year, and d is the postcode geographic unit for fixed effects.

Outcome: $log(price_{it})$
As discussed in Section 6.2, we use the *log* transformation of a transaction price.

Treatment variable: D_{it}
The treatment variable differs across models. For proximity models it captures distance to the nearest turbine. For visibility models it captures the number of turbines visible from the property in year t, derived from cumulative viewshed rasters. For shadow-flicker models it captures simulated annual shadow-flicker exposure at the property. Each model section defines D_{it} precisely.

Hedonic controls: X_i
All models include floor area (m^2), number of rooms, property type dummies (semi-detached, terraced, flat, with detached as reference), and a new-build indicator. These control for property quality differences within groups, which helps ensure that the estimated treatment effect is not simply capturing differences in property quality.

Geographic fixed effects: α_d
This is important because properties near turbines may be located in systematically different areas from properties farther away. By including geographic fixed effects, the model compares properties within the same broad geographic area over time rather than comparing all properties directly against each other. This reduces bias from fixed local characteristics that do not change during the study period.

The geographic granularity differs by model family. Proximity models use **postcode area** fixed effects (approximately 100 groups): the treatment variable varies across a wide distance range and a finer grouping would absorb most of that variation. Visibility models use **postcode district** fixed effects (approximately 2,200 groups), since visibility contrasts are sharper and more localised and a finer grouping is appropriate. Shadow-flicker models use **postcode area** fixed effects, since exposure is sparse and confined to within 3 km and a finer grouping would leave too few exposed properties per cell.

Year fixed effects: δ_t
The year fixed effect absorbs and control for shocks that affect the housing market as a whole in a given year. For example, if property prices rise nationally in one year, the year fixed effect absorbs that common increase. This means that the treatment effect is identified from differences between treated and control properties within the same time period.

The error term: ε_{it}
It captures variation in property prices that is not explained by the treatment variable, property controls, geographic fixed effects, or year fixed effects.

7.4 Event Study Design

The event study design extends the difference-in-differences framework by estimating effects separately for different years before and after treatment. Instead of estimating one average post-treatment effect, it traces how property prices evolve relative to the treatment year.

Event time k is defined as:

$$k = \text{transaction year} - \text{treatment year}$$

where the treatment year is the commissioning year (proximity) or the year turbines first became visible from the property (visibility). A transaction at $k = 0$ happened in the same year as treatment. The reference period is $k = -1$, so all coefficients are measured relative to the year immediately before treatment.

The event study replaces the single treatment indicator with a set of event-time indicators. This allows the model to estimate how prices differ in each period before and after treatment. The post-treatment coefficients describe the estimated price path after treatment, while the pre-treatment coefficients are used to examine whether treated and control properties were already following different trends before treatment occurred.

This is important because difference-in-differences relies on the parallel trends assumption. If treated and control properties have similar trends before treatment, the assumption becomes more plausible. If they already differ before treatment, the estimated treatment effects must be interpreted more cautiously.

7.5 Staggered Adoption and Dynamic Effects

Because turbine exposure occurs in different years for different properties, the event study must account for staggered adoption. In a staggered setting, properties treated earlier may already be exposed by the time later properties become treated. This creates a problem for standard two-way fixed effects event studies, because already-treated observations may be used as part of the comparison group for later-treated observations.

Before estimating the event studies, transactions with ambiguous treatment timing are removed. This is needed because some locations may be exposed to turbines from more than one treatment cohort. In that case, the same transaction could be before treatment relative to one cohort but already after treatment relative to another, which makes the event time difficult to

interpret. The analysis therefore keeps only observations with a clear treatment timing for the relevant event study.

To avoid this issue, the event study models use an interaction-weighted (IW)[18] estimator. This estimator first estimates event-time effects separately by treatment cohort and then aggregates them across cohorts. This produces a pooled event-time path while avoiding comparisons that use already-treated properties as inappropriate controls. The logic of this estimator is shown in Figure 7.3: cohort-specific event-time paths are estimated first, then combined into a pooled dynamic treatment path.

The event study estimates are reported as $ATT(k)$, where ATT means the average treatment effect on the treated. In this context, $ATT(k)$ describes the average price difference for treated properties at event time k, measured relative to the reference period $k = -1$. Before aggregation, the IW estimator estimates the same quantity separately for each treatment cohort, denoted $ATT(g, k)$, where g identifies the cohort of properties first treated in the same year. These cohort-specific estimates are then combined into the pooled $ATT(k)$ path shown in Figure 7.3.

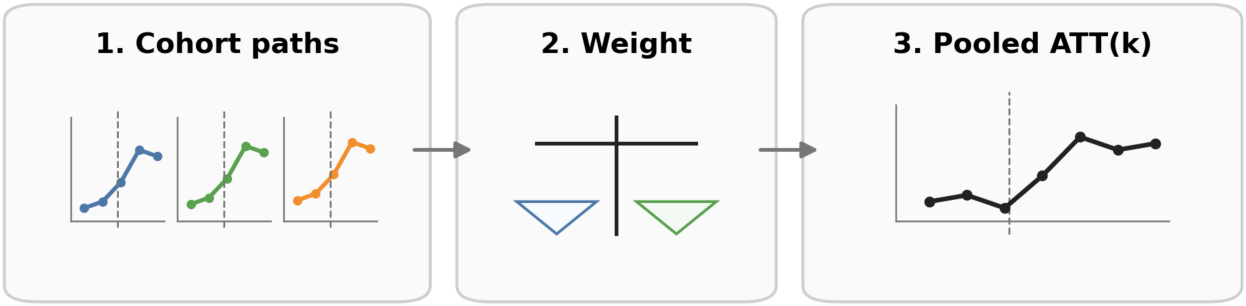

Figure 7.3: Interaction-weighted event study estimation under staggered adoption. Event-time effects are first estimated separately for each treatment cohort and then aggregated into a pooled event-time path. This avoids using already-treated properties as controls for later-treated cohorts.

7.6 Inference and Standard Errors

The models use robust standard errors to account for uncertainty in the estimated coefficients. Property price data are naturally heteroskedastic, meaning that the variance of prices is not constant across observations. Robust standard errors reduce the risk that statistical significance is overstated because of unequal error variance.

For static models, heteroskedasticity-robust standard errors (HC1) are used. For event study models, cluster-robust standard errors are used to account for correlation between observations within geographic areas. This is important because nearby properties may be affected by similar local housing market conditions, planning decisions, or unobserved neighbourhood characteristics.

CHAPTER 8

Models & Results

This chapter presents the empirical models and their results. Building on the common DiD framework introduced in the previous chapter, the analysis first examines proximity-based exposure, then visibility-based exposure, and finally shadow-flicker exposure. For each model, the design choices are introduced before the results are interpreted, with particular attention to whether the estimated effects are stable across model specifications and distance or exposure levels.

8.1 Proximity Models

8.1.1 ProxGrad — Proximity Gradient (Baseline Model)

Design

This model asks: *how does the price effect change as a property gets closer to a turbine?*

Rather than a single binary treatment (near vs. far), we split the area around each turbine into 1 km rings and estimate a separate DiD coefficient for each ring. This provides a more detailed view of how the estimated effect varies with distance, rather than assuming it is uniform up to some cutoff.

Key design choices

Figure 8.1 shows a visualization of the main design choices that was taken for proximity models, which are:

- **9 treatment rings:** 0–1, 1–2, ..., 8–9 km. Each ring is interacted with the post-commissioning indicator, giving one DiD estimate per ring.
- **Omitted buffer ring at 9–10 km:** this gap separates the treated rings from the reference group, reducing boundary contamination between near-treated and control observations.
- **Control ring at 10–11 km:** all effects are measured relative to properties in this ring. We chose 10–11 km on two grounds. First, the hedonic pricing literature finds proximity effects on property prices become negligible well before this distance. For England and Wales specifically, Gibbons (2015)[5] finds that only the largest wind farms (20+ turbines) produce small residual effects beyond 8 km. Secondly, properties at 10–11 km are still in the same rural, wind-farm geography as the treated rings.

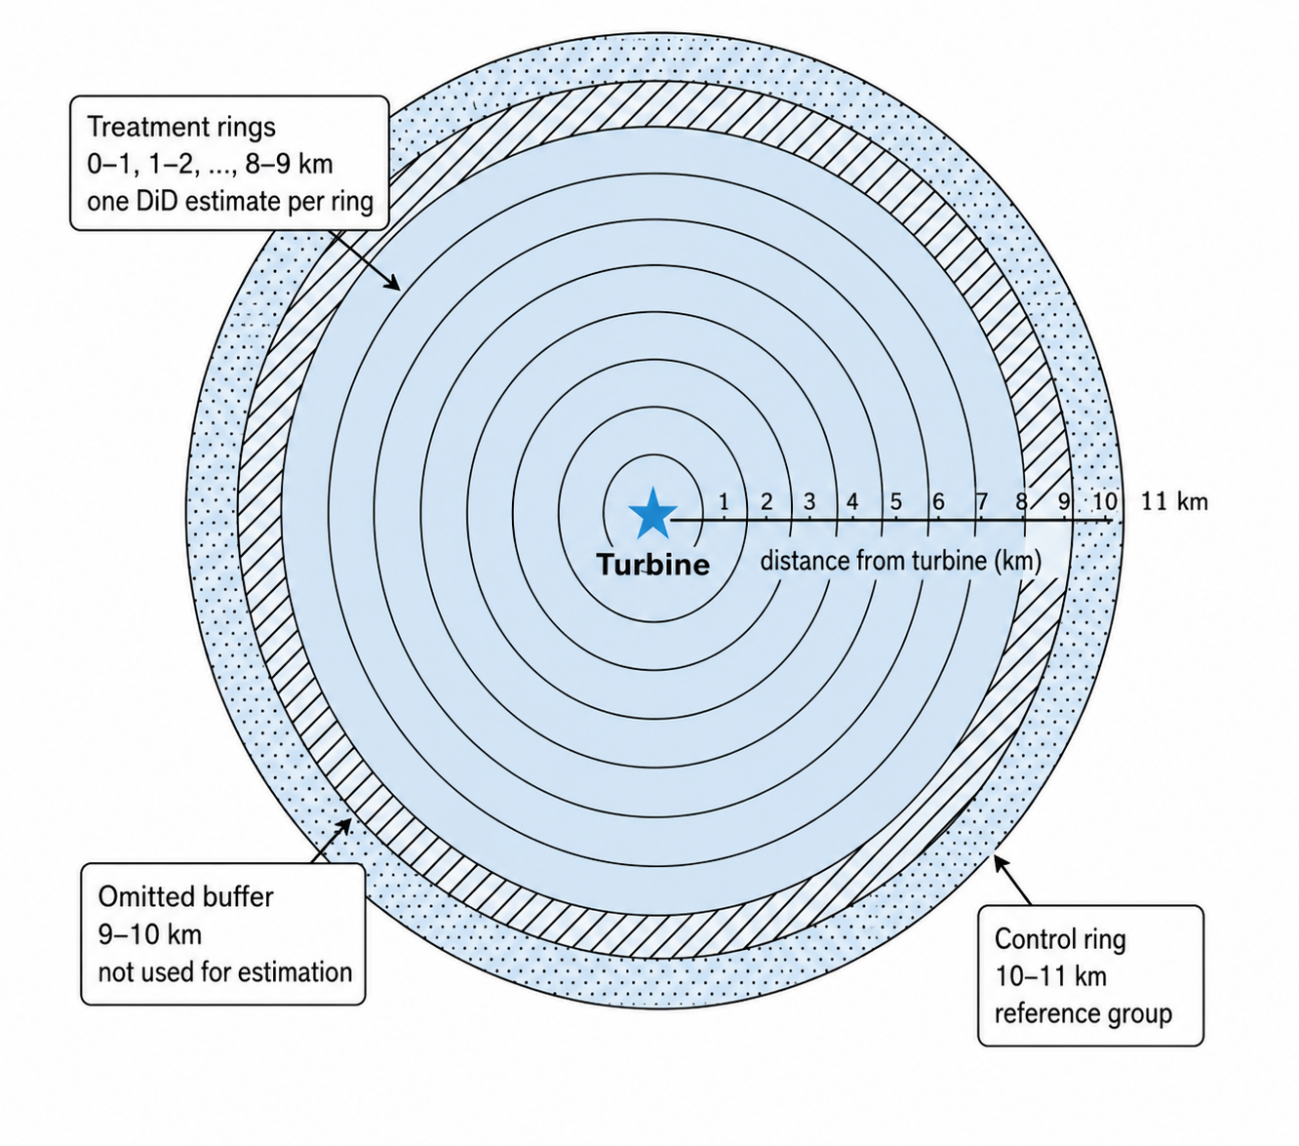

Figure 8.1: Visualization of each of treatment and control rings around a turbine, with the omitted buffer from 9–10 km away.

- **Postcode area FE** (≈100 groups) rather than district. Within a postcode district (∼1 km across), distance to a turbine varies very little, meaning that most properties are roughly the same distance away. Using postcode district fixed effects would remove most of the distance-band variation we are trying to estimate. Postcode area is broad enough that multiple distance rings are represented within each group, so identification is preserved.

Results

The estimated distance-band effects are reported in Table 8.1. While Figure 8.2 visualizes the same distance gradient and shows the transaction count in each distance band. The figure makes the main pattern clearer: the estimated discount is largest close to turbines and gradually fades as distance increases. All coefficients are measured relative to the 10–11 km control ring.

Table 8.1: Distance-band DiD results (ProxGrad).
$R^2 = 0.7562$

Distance band	Effect	95% CI	p-value	N
0–1 km	**−4.29%**	[−5.22, −3.34]	< 0.001	31,391
1–2 km	**−2.63%**	[−3.20, −2.04]	< 0.001	118,214
2–3 km	**−3.30%**	[−3.77, −2.82]	< 0.001	194,212
3–4 km	**−1.48%**	[−1.93, −1.02]	< 0.001	260,277
4–5 km	**−2.23%**	[−2.67, −1.79]	< 0.001	274,037
5–6 km	**−1.72%**	[−2.15, −1.28]	< 0.001	272,789
6–7 km	**−0.21%**	[−0.66, +0.24]	0.350	270,526
7–8 km	**+0.44%**	[−0.01, +0.89]	0.054	269,407
8–9 km	**+1.70%**	[+1.25, +2.15]	< 0.001	265,592
10–11 km	0.00% (control ring)	—	—	256,897

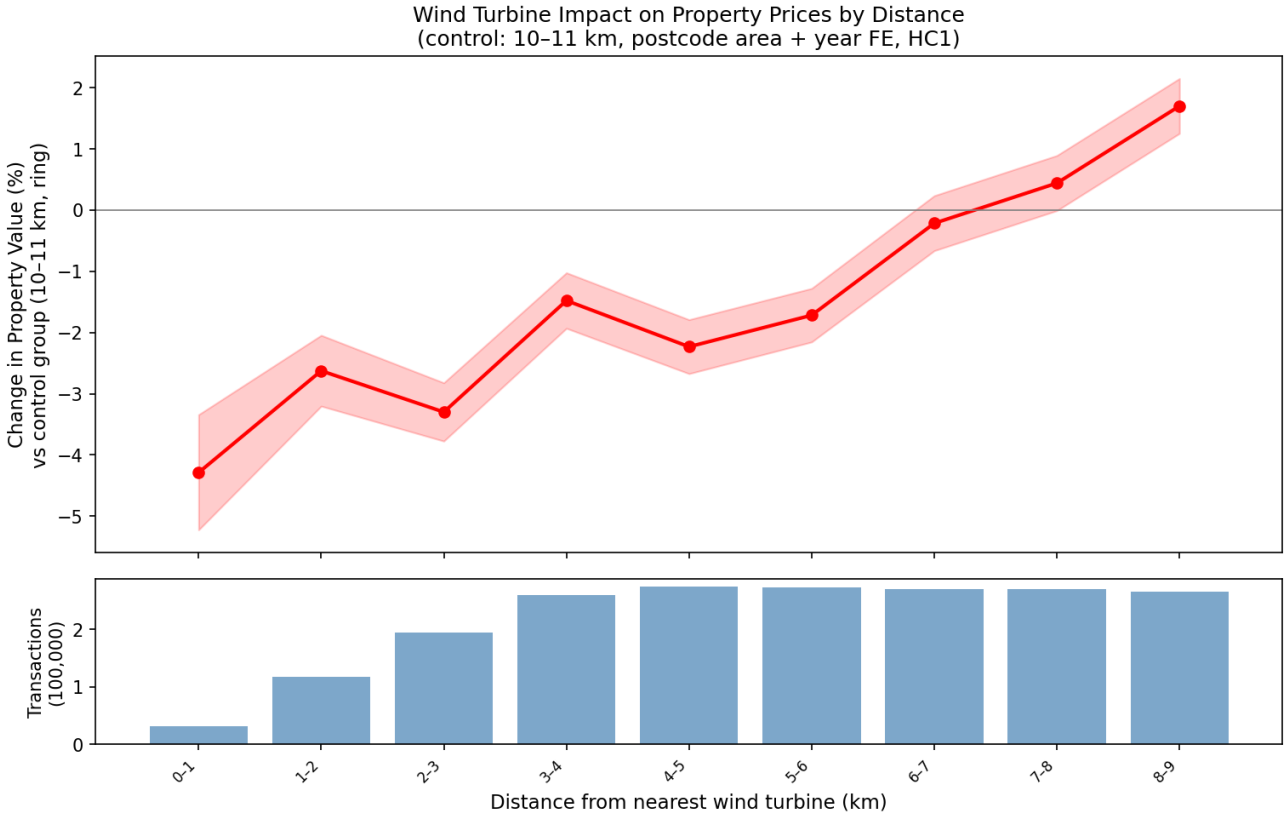

Figure 8.2: Distance gradient from the ProxGrad model. Price effect relative to the 10–11 km control ring, with 95% confidence bands. Bottom panel: transaction counts per band.

Interpretation

The estimates show the expected distance gradient, a strong negative effect close to the turbine that fades as distance increases. The estimated discount is largest within 1 km, at −4.29%, and becomes statistically indistinguishable from zero at 6–7 km ($p = 0.350$, CI spanning zero on both sides). Two observations are worth noting:

- The gradient is not perfectly monotonic in the 1–6 km range, there are small fluctuations between adjacent bands. These fluctuations are likely due to sampling variation or local composition differences between adjacent distance bands. All bands through 5–6 km are significantly negative and the overall pattern is clearly downward.
- The 8–9 km band shows a small but significant positive effect (+1.70%), just inside the omitted buffer and control ring. This boundary result should be interpreted cautiously, since it may reflect residual near-turbine exposure or local composition differences close to the reference ring rather than a genuine price premium.

8.1.2 Proximity Event Study

Design

This model asks: *when does the proximity effect appear, and does it persist after commissioning?*

Where ProxGrad estimates how the effect varies by distance, ProxES estimates how the proximity effect changes over time before and after commissioning. The model keeps the same treatment and control geography as ProxGrad, but replaces the static distance-band effects with event-time estimates.

Event time k and the reference period $k = -1$ follow the convention introduced in the Methodology chapter.

Key design choices:

- **Event window** $k \in [-6, +4]$**:** the window is asymmetric because the 2011–2019 transaction panel provides more pre-commissioning coverage for late cohorts than post-commissioning coverage. Strict window enforcement is applied, so only transactions that fall within this range are included.
- **Cohorts 2013–2017:** properties are grouped by the commissioning year of their nearest turbine or farm. Five cohorts enter the analysis.
- **Cohort isolation:** Each postcode is checked against all study period turbine cohorts, and transactions are kept only if their postcode is exposed to at most one commissioning cohort. Postcodes exposed to more than one cohort are removed because their event time is ambiguous.
- **Treated group:** isolated properties within 9 km of the nearest commissioned turbine.
- **Omitted buffer at 9–10 km:** excluded from both treatment and control to reduce boundary contamination.
- **Control ring at 10–11 km:** effects are measured relative to properties in this anchor ring, the same reference geography as ProxGrad.
- **Flexible distance controls:** the model controls for distance to the nearest turbine using breakpoints at 2, 5, and 8 km. This helps ensure that the event-time estimates reflect changes around commissioning rather than only the distance gradient already shown in ProxGrad.
- **Postcode area and year fixed effects:** postcode area fixed effects account for persistent differences between broader local housing markets, while year fixed effects absorb common housing-market shocks.
- **Support thresholds:** a cohort-event-time cell must contain at least 100 treated and 100 control observations to enter the estimation. The pooled path requires at least 3 for each event-time bin.

Estimators

ProxES-IW is the headline dynamic estimator. It estimates cohort-specific event-time effects and then combines them into a pooled $ATT(k)$ path. The cohort-specific paths are also reported as diagnostics, since they show whether the pooled result is driven by one cohort or whether the pattern appears across several commissioning years. A separate cohort-level ProxES specification is used as a sensitivity check.

Results

The results are presented in two steps: first the pooled ProxES-IW path, and then the cohort-specific paths used to assess heterogeneity across commissioning years.

ProxES-IW — pooled dynamic path

The pre-trend joint F-tests yield $p = 0.221$ for 2013 cohort and $p < 0.001$ for the 2014, 2015, 2016, and 2017 cohorts. These tests examine whether the pre-treatment event-time coefficients are jointly equal to zero. Since this is rejected for four of the five cohorts, treated and control properties were already following different price paths before commissioning. **The parallel trends assumption is therefore not supported**, and the post-commissioning estimates should be interpreted as descriptive timing patterns rather than clean causal effects.

The pooled IW path (shown in Figure 8.3) declines from roughly +6.50% at $k = -4$ to +1.42% at $k = -2$ before reaching the reference period at $k = -1$. After commissioning, the estimate becomes negative, $\approx -2.01\%$ at $k = 0$, and $\approx -3.13\%$ at $k = 1$. It then remains negative, at around -1.94% by $k = 4$. This pattern is consistent with a modest post-commissioning discount, but because the pre-treatment path is already declining, it should not be read as a clean causal timing effect.

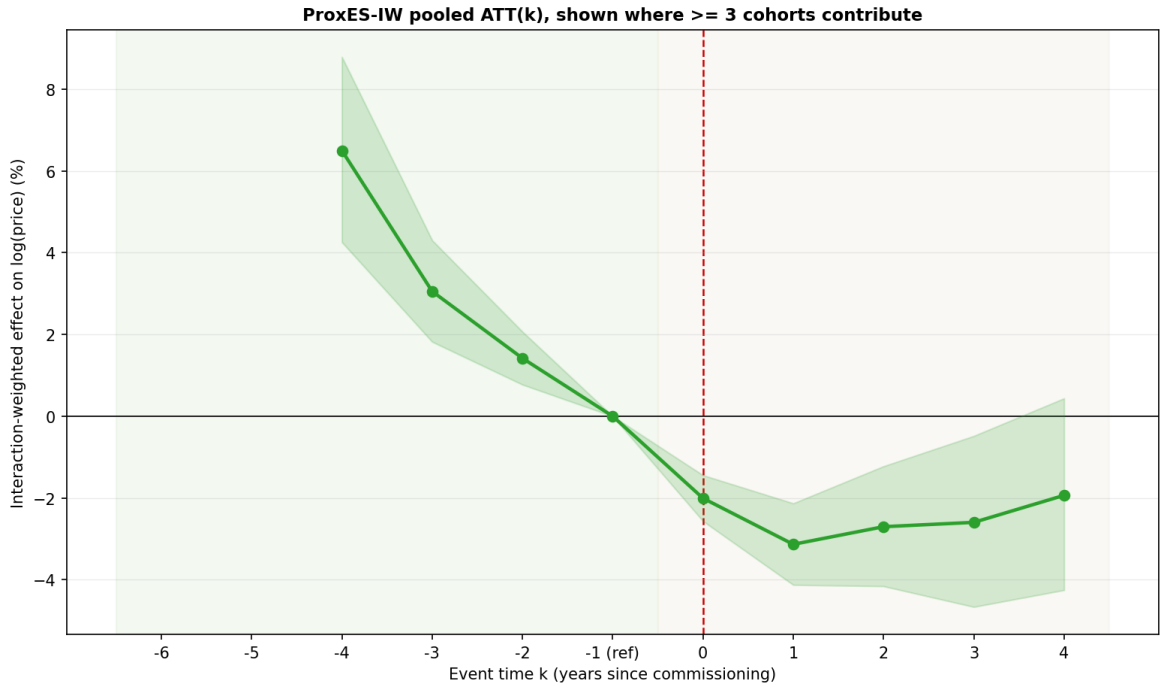

Figure 8.3: Pooled ProxES-IW dynamic $ATT(k)$ path. Points report the interaction-weighted proximity effect on log price at each event time relative to the reference year $k = -1$. The shaded band shows the 95% confidence interval and the red dashed line marks the commissioning year.

ProxES — cohort-by-cohort sensitivity

Figure 8.4 reports the cohort-specific ProxES-IW paths and shows that the pooled result hides substantial variation across cohorts. The 2013 cohort has the most persistent negative post-commissioning estimates, around -5.11% to -5.57% for $k = 1$ to 3. The 2014 and 2015 cohorts show smaller and shorter-lived negative effects, while the 2016 and 2017 cohorts are closer to zero after commissioning. However, the later cohorts have less post-treatment support, with only $k \leq 3$ observed for 2016 and $k \leq 2$ for 2017. Overall, the pooled path is not driven by one single cohort, but the cohort-level variation reinforces the need to interpret the dynamic results cautiously.

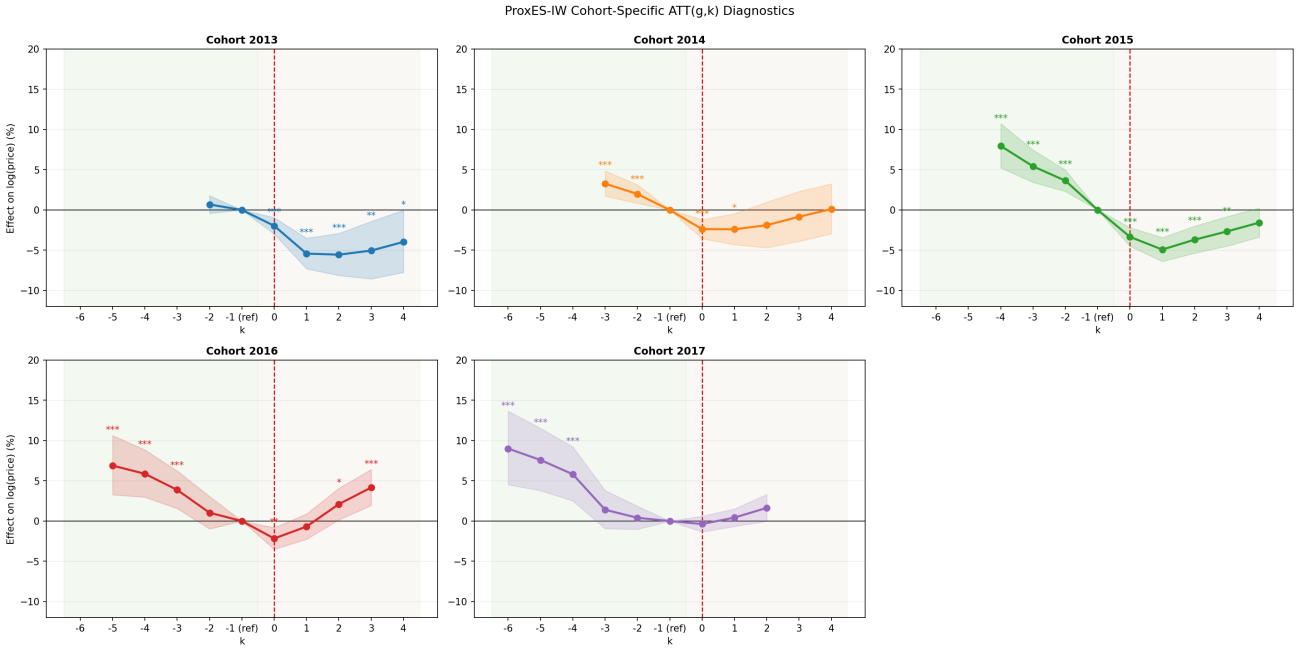

Figure 8.4: Cohort-specific ProxES-IW event-study paths. Each panel reports $ATT(g, k)$ for one commissioning cohort, with effects measured relative to $k = -1$. Stars mark statistically significant event-time coefficients, and the red dashed line marks the commissioning year.

8.1.3 Results Summary — Proximity Models

ProxES supports the static ProxGrad result by showing negative estimates after commissioning. However, because several cohorts fail the pre-trend tests, the dynamic estimates are best treated as descriptive evidence on timing rather than as strong causal estimates.

8.2 Visibility Models

The proximity models treat distance to the nearest turbine as a proxy for exposure. This proxy is imperfect: two properties at the same distance can have entirely different experiences depending on the terrain between them and the turbines. The visibility models replace the distance proxy with a direct measure of exposure drawn from the annual cumulative viewshed rasters.

Section 8.2.1 presents the three core visibility DiD specifications. Section 8.2.2 tests whether visibility operates independently of proximity by including both channels in the same regression. Section 8.2.3 traces how the visibility effect evolves over time using an event study design.

8.2.1 Visibility DiD

Design

The visibility models use the model-ready visibility sample described in the Data Preprocessing chapter 6. In this part of the analysis, exposure is measured directly from the annual viewshed rasters rather than inferred from distance to the nearest turbine. The key variables are the number of turbines visible from each property in the transaction year and the first year in which the property location gains any turbine visibility.

Using these variables, we estimate three increasingly flexible visibility specifications, each adding a layer of detail:

Vis-Binary — *Does gaining* any *turbine visibility reduce prices?*
It uses a single coefficient `treat_post`, which equals 1 when a property location has gained visibility of at least one turbine by the transaction year. It estimates the average price difference associated with moving from no turbine visibility to any turbine visibility.

Treatment variable: `treat_post` = `ever_treated` × `post`.

Vis-Dose — *Does each additional visible turbine add to the penalty?*
This specification uses the raw number of visible turbines in the transaction year. It imposes linearity, meaning that each additional visible turbine is assumed to have the same marginal association with price.

Treatment variable: `n_visible`

Vis-Bands — *Is the relationship non-linear?*
This specification groups visibility into four bands (1–5, 6–10, 11–15, 16+), using 0 visible turbines as the control group. This allows the estimated price effect to differ across levels of visual exposure.

Results

Vis-Binary
Table 8.2 shows that properties that gain any turbine visibility are associated with sale prices approximately -0.83% lower than comparable properties that have not gained visibility in the

same postcode district and year. The confidence interval of [-0.97%, -0.69%] is narrow, and the p-value ≈ 0 suggests that the estimate is unlikely to be explained by sampling variation alone.

Vis-Dose

Table 8.2 shows that each additional visible turbine is associated with an additional -0.25% lower sale price. The Vis-Bands model below tests whether this linearity assumption holds.

Table 8.2: Vis-Binary and Vis-Dose results.
$R^2 = 0.839$

Model	Variable	Effect	p-value
Vis-Binary	`treat_post`	$\mathbf{-0.83\%}$	≈ 0
Vis-Dose	`n_visible` (per turbine)	$\mathbf{-0.25\%}$	≈ 0

Vis-Bands

The non-linear visibility estimates are shown visually in Figure 4.1 and reported numerically in Table 4.2. Together, they show clearly how the estimated price discount changes as the number of visible turbines increases.

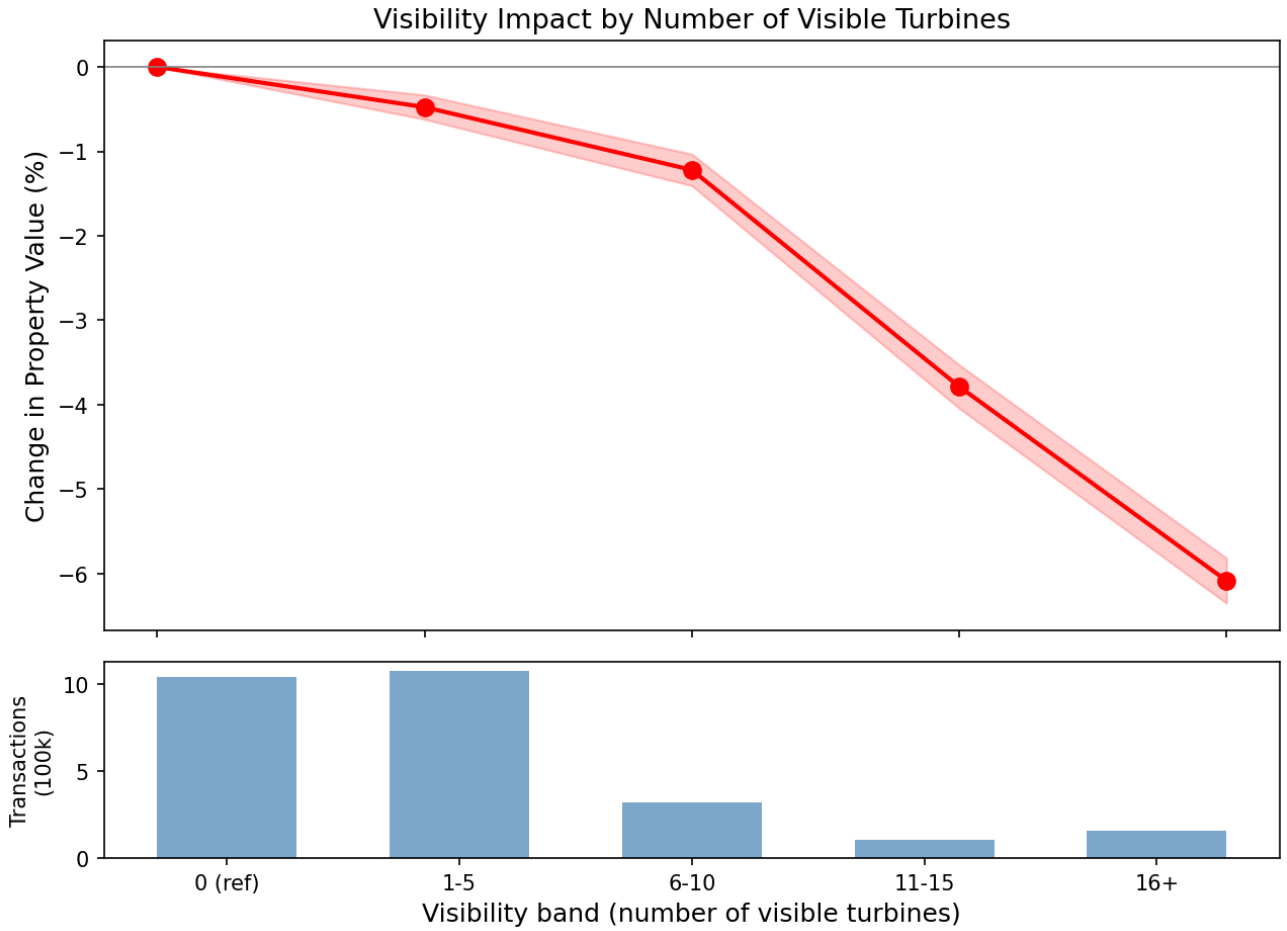

Figure 8.5: Visibility dose-response curve (Vis-Bands). Price effect by number of visible turbines, with 95% confidence bands. Bottom panel: transaction counts per band.

Table 8.3: Vis-Bands results.
$R^2 = 0.840$

Visible turbines	Effect	95% CI	p-value	N
0 (ref)	0.00%	—	—	1,040,900
1–5	**−0.48**%	[-0.62%, -0.33%]	< 0.001	1,074,728
6–10	**−1.22**%	[-1.41%, -1.03%]	≈ 0	321,935
11–15	**−3.79**%	[-4.04%, -3.53%]	≈ 0	104,802
16+	**−6.08**%	[-6.35%, -5.82%]	≈ 0	155,889

Three patterns stand out:

- **The effect is strongly non-linear.**
  Seeing 1–5 turbines reduces prices by -0.48%. Seeing 16 or more reduces prices by -6.08%. The Vis-Dose linearity assumption is clearly violated.
- **The largest escalation occurs between the 6–10 and 11–15 bands.**
  The penalty increases from -1.22% to -3.79%. This suggests a threshold effect: a horizon dominated by turbines is experienced qualitatively differently from one in which a handful are visible in the periphery.
- **The high penalty for properties with 16 or more visible turbines**
  These locations face a landscape in which turbines are a prominent and unavoidable visual feature. The penalty of -6.08% [-6.35%, -5.82%] is large and highly significant.

8.2.2 Disentangling Visibility from Distance

The visibility models above show that turbine visibility is associated with lower property prices. However, visibility and distance are naturally correlated, because properties close to turbines are more likely to have a clear line of sight to them. The key question is therefore whether the visibility estimate captures actual visual exposure, or whether it simply reflects the fact that visible properties also tend to be closer to turbines.

If visibility were only a proxy for proximity, the visibility coefficient should shrink substantially or disappear once distance is added to the model. This section tests that directly by including both channels in the same regression.

Design

We restrict to the maintained transactions sample described in the Data Preprocessing chapter 6, and estimate four specifications, each pairing a visibility treatment with a distance control (these design choices are explained later in this section):

VisBin-DistLog: combines binary visibility with the log of distance to the nearest turbine.

VisBin-DistBands: combines binary visibility with one-kilometre distance-band controls. The 11–12 km ring is omitted as a buffer, and the 12-–13 km ring is used as the distance control group.

VisDose-DistLog: combines the continuous visible-turbine count with the log of distance.

VisBands-DistBands: combines visibility bands and distance bands in the same regression. This is the most flexible specification because it allows both the visibility channel and the distance channel to be non-linear.

Visibility-Distance treatment and control groups

This part is dedicated to answer the question of: *Why do we use 12–13 km as the distance control group?*

Figure 8.6 shows why the combined design needs both distance and visibility. The two measures are related, with a negative rank correlation of $\rho = -0.4616$, meaning that properties farther from turbines are generally less likely to have high turbine visibility. The figure itself plots all transactions within 14 km for visual context. But the relationship is not deterministic. In the near rings, some properties see no turbines, while farther rings still contain properties with positive visibility, which shows that distance alone is an imperfect treatment definition for visual exposure.

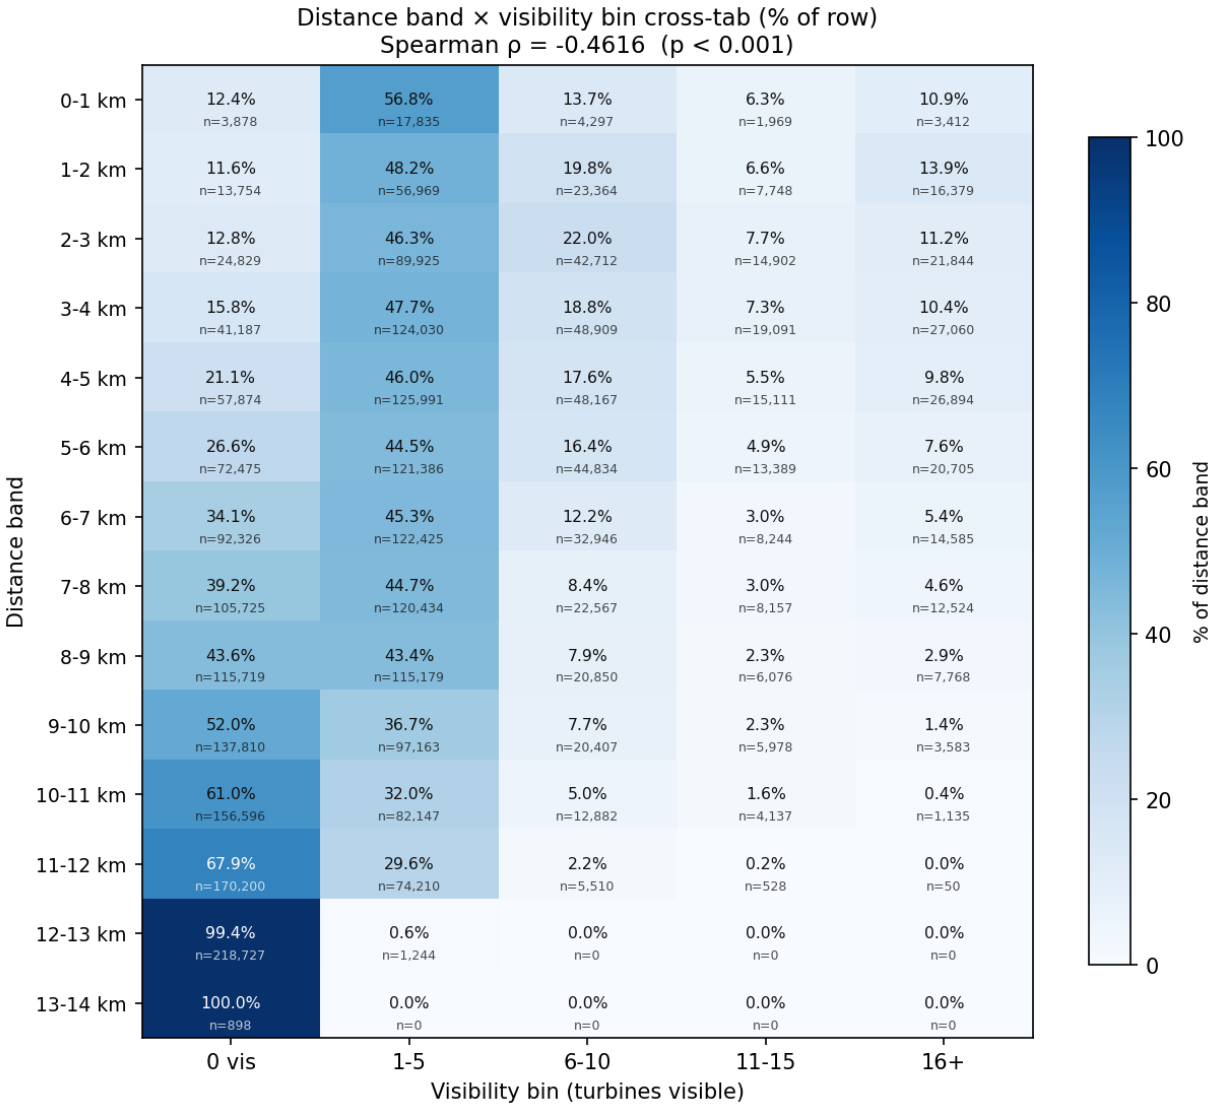

Figure 8.6: Distance and visibility support. Cells report the share of transactions in each distance band falling into each visibility bin. Visibility generally declines with distance, but the within-band variation shows that distance and visibility are related exposure measures rather than substitutes.

The distance structure used here differs from the proximity-only models. **In the proximity models**, the treatment rings end at 9 km, with a 9–10 km omitted buffer and a 10–11 km control ring. This choice produced the most stable proximity gradient and is consistent with the literature reviewed earlier, where proximity effects are generally not expected beyond this range.

For the visibility-distance models, the distance support is extended because turbine visibility can still occur at larger distances. Some properties still have turbine visibility beyond 11 km, even if they are outside the main proximity-effect range. For this reason, the combined models use 0-–11 km as the main distance-treatment rings, omit the 11-–12 km ring as a buffer, and use the 12—13 km ring as the distance-control ring. This gives a nearby control group that remains within wind-farm geography but has very limited turbine visibility.

In the combined visibility-distance model, the baseline comparison is therefore a property in the 12-–13 km reference ring with zero visible turbines.

Why $log(distance)$ rather than distance in levels?

The log-distance specifications are included because turbine-related proximity effects are unlikely to change linearly with distance. Moving from 1 km to 2 km is a much larger change in exposure than moving from 9 km to 10 km. A log transformation captures this diminishing sensitivity, because equal changes in log(distance) correspond to equal proportional changes in distance, so the model naturally weights short-range variation more heavily. This is consistent with ProxGrad, where the gradient is steepest close to the turbine and then flattens at larger distances.

Results

Does visibility survive distance controls?
Table 8.4 reports the visibility estimates before and after adding distance controls. The binary visibility estimate changes only slightly, from -0.83% without distance controls to -0.78% after distance is controlled. The continuous dose estimate also changes only modestly, from -0.25% to -0.23% per additional visible turbine. Both effects remain highly significant.

Table 8.4: Visibility estimates before and after adding distance controls.

Model	Visibility variable	Effect	p-value
Vis-Binary (no distance control)	`treat_post`	$\mathbf{-0.83}\%$	≈ 0
VisBin-DistLog	`treat_post`	$\mathbf{-0.78}\%$	≈ 0
VisBin-DistBands	`treat_post`	$\mathbf{-0.78}\%$	≈ 0
Vis-Dose (no distance control)	`n_visible`	$\mathbf{-0.25}\%$/turbine	≈ 0
VisDose-DistLog	`n_visible`	$\mathbf{-0.23}\%$/turbine	≈ 0

This shows that the **visibility estimate is not simply capturing proximity**. Adding distance controls explains only a small part of the visibility effect, meaning that **visibility remains an independent channel associated with property prices.**

The log-distance models also show that distance continues to matter after visibility is controlled. **In VisBin-DistLog, the coefficient on log(distance) is +3.92%**. Since doubling distance increases log distance by $log(2)$, which is ≈ 0.693, doubling the distance to the nearest turbine is associated with approximately $+3.92\% \times 0.693 \approx 2.8\%$ higher sale prices, holding visibility and the other controls constant. This remaining distance gradient represents the proximity component that is not explained by visibility.

VisBands-DistBands

The VisBands-DistBands model is the most flexible specification, because it allows both channels to be non-linear simultaneously, by including visibility bands and distance bands in the same regression. Its purpose is to separate the two channels directly: it compares properties with *different visibility* levels within the same distance band, and properties at *different distances* within the same visibility level.

Figure 8.7 summarizes the combined predicted effect from both channels using a property in the 12—13 km control ring with zero visible turbines as the baseline. The largest predicted discounts occur where close distance overlaps with high visibility.More importantly, within the same distance band, higher visibility is associated with lower prices. This directly answers the *research question of whether the price effect is driven by actual visibility or by proximity alone*: **visibility is not only a proxy for proximity, because it remains associated with lower prices even after distance is held constant.**

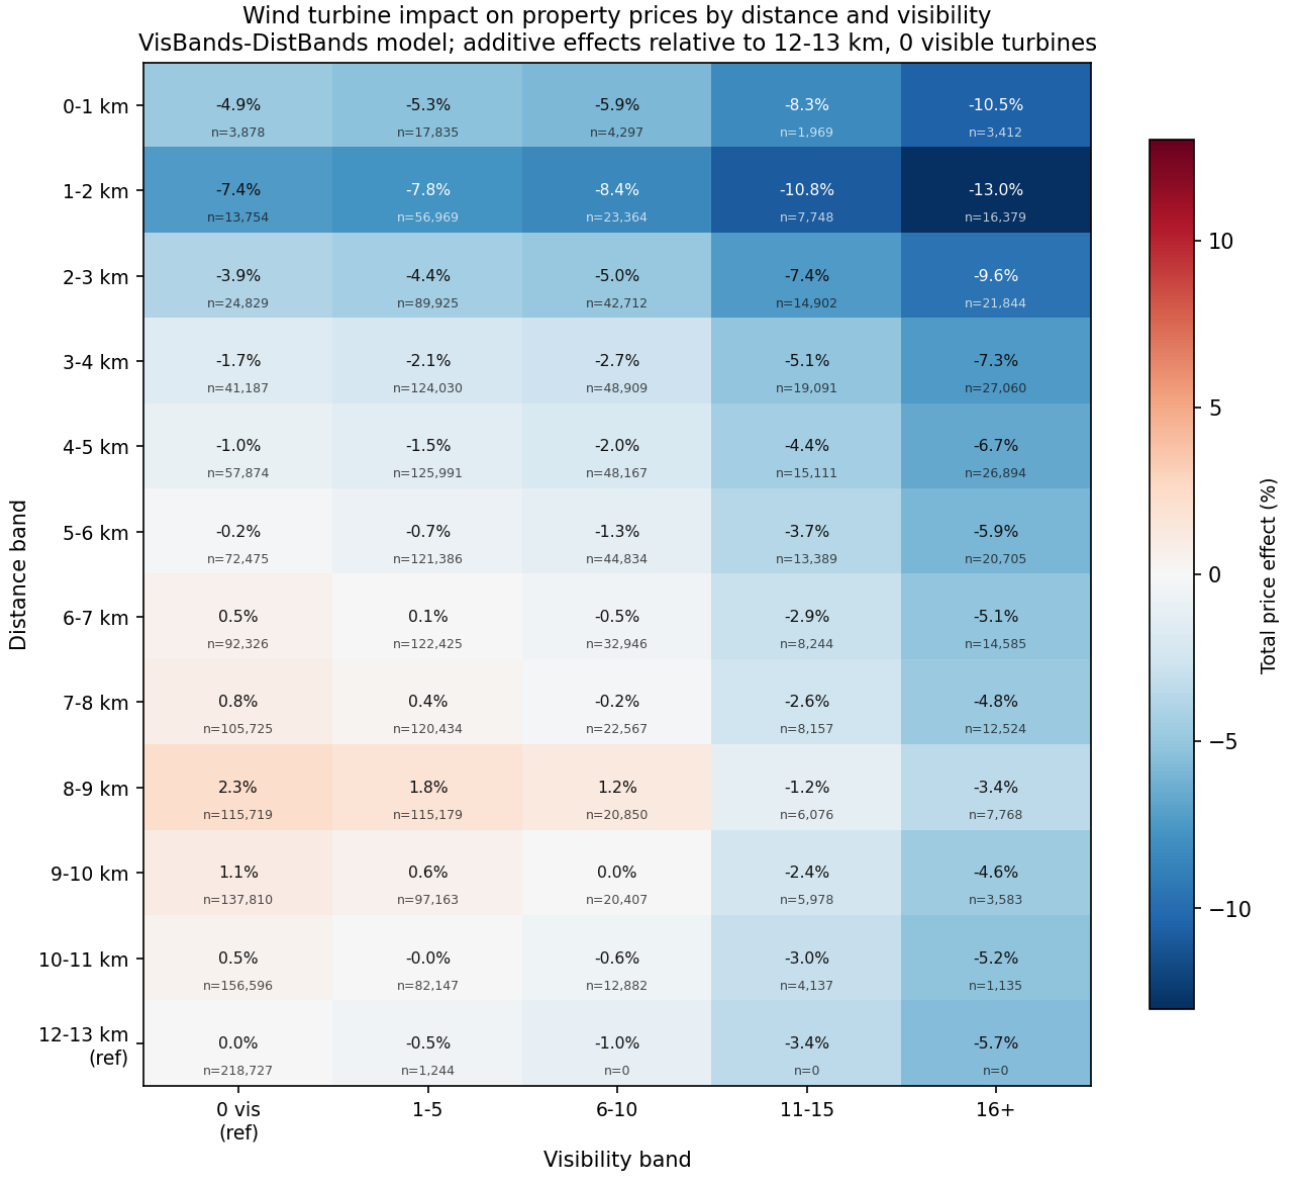

Figure 8.7: Combined distance and visibility price effects from VisBands-DistBands. Each cell adds the estimated distance-band effect and visibility-band effect.

Visibility channel:

Table 8.5 and Figure 8.8 compare visibility-band estimates before and after adding distance-band controls. In Vis-Bands (no distance), the control group is any property with no visible turbines regardless of how far away it is. While in VisBands-DistBands, the control group is properties with no visible turbines *in the same distance band*, so distance is held constant, and what remains is the pure visual channel.

Table 8.5: Visibility band estimates before (Vis-Bands) and after (VisBands-DistBands) controlling for distance.

Visible turbines	Vis-Bands (no dist.)	VisBands-DistBands	Change
1–5	−0.48%	−0.47%	+0.01 pp
6–10	−1.22%	−1.03%	+0.18 pp
11–15	−3.79%	−3.44%	+0.35 pp
16+	−6.08%	−5.67%	+0.41 pp

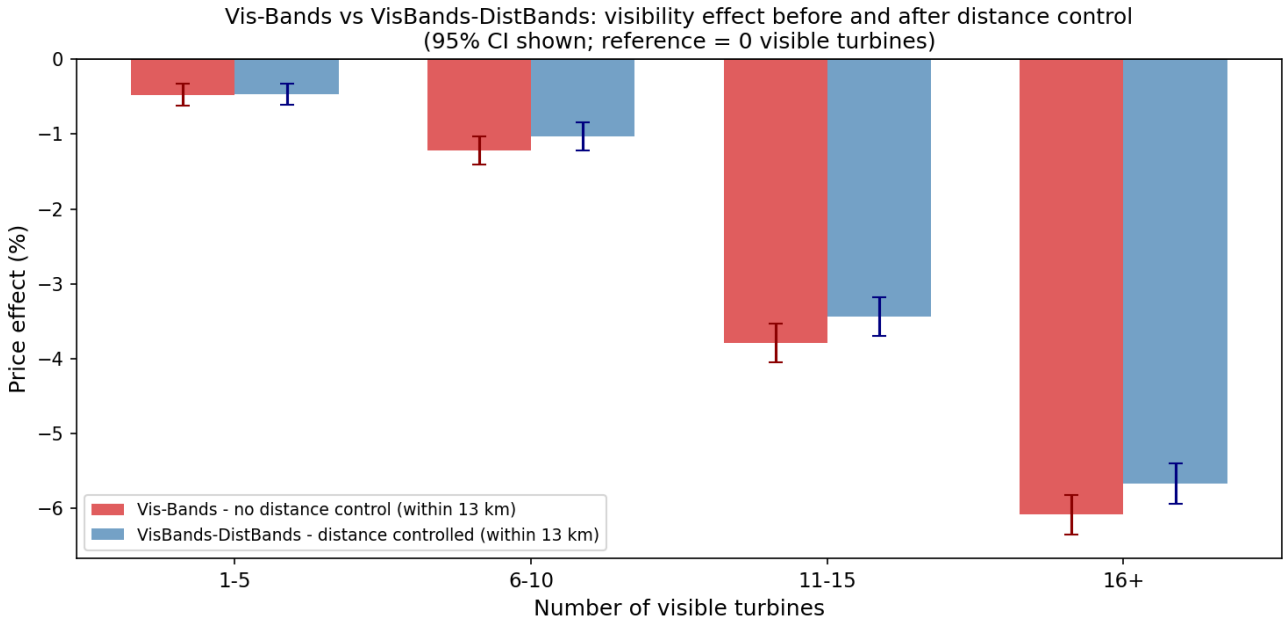

Figure 8.8: Visibility band price effects before (Vis-Bands, red) and after (VisBands-DistBands, blue) controlling for distance, on the same within-13 km sample. Error bars show 95% confidence intervals.

The low-visibility bands are essentially unchanged: the 1–5 turbine estimate almost unchanged, moving from -0.48% to -0.47% once distance is controlled.

The high-visibility bands are slightly attenuated: the 11–15 estimate falls from -3.79% to -3.44%, and the 16+ estimate changes from -6.08% to -5.67%.

This suggests that part of the high-visibility discount overlaps with proximity, since properties with many visible turbines are often also closer to turbines. However, most of the visibility gradient remains after distance is controlled. **The visual channel therefore remains separate from the proximity channel.**

Distance channel:
the distance gradient from the same VisBands-DistBands model after controlling for visibility. The distance coefficients are measured relative to the 12–13 km control ring. The gradient remains negative in the closest distance bands, showing that **proximity still matters even after turbine visibility is controlled**.

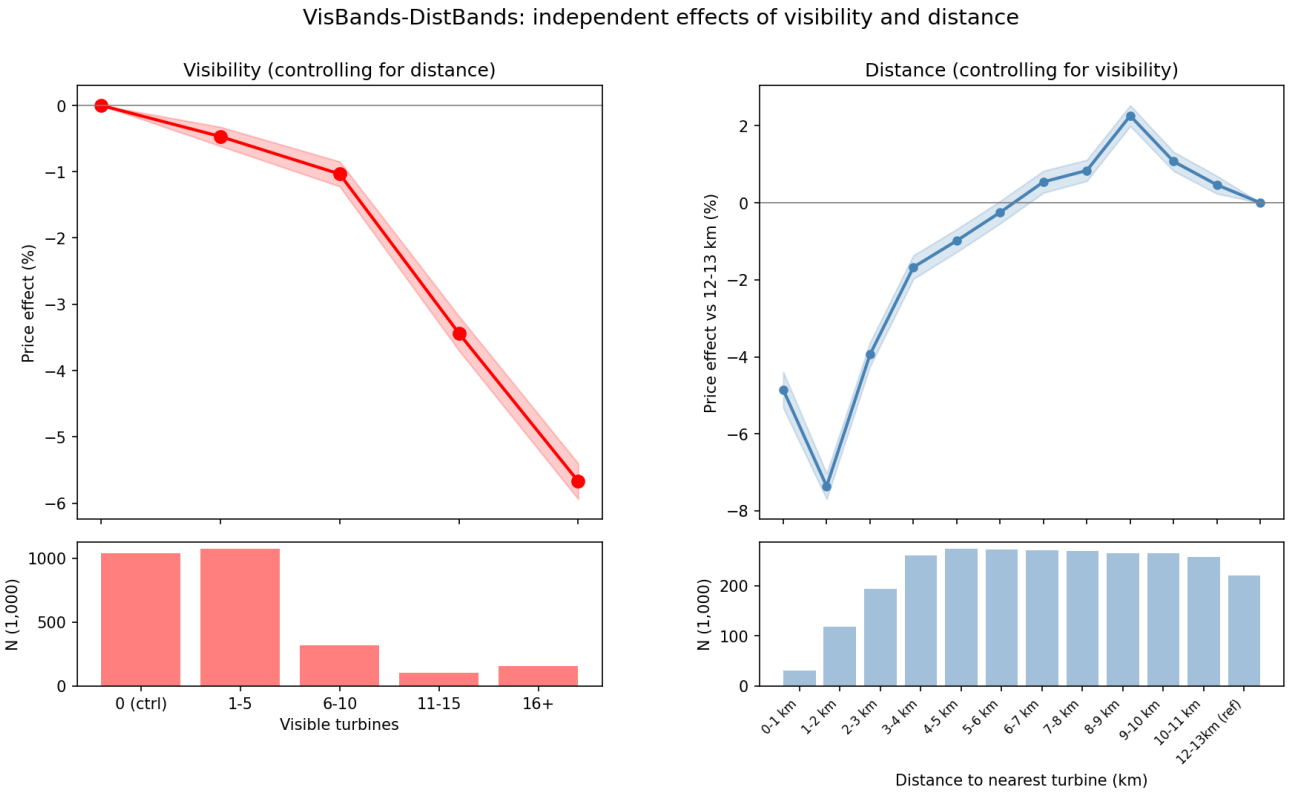

Figure 8.9: VisBands-DistBands results. *Left:* visibility band effects controlling for distance. *Right:* distance band effects controlling for visibility. Both panels come from the same regression.

One small irregularity is that the 0–1 km distance coefficient is *less* negative than the 1–2 km coefficient (-4.86% vs. -7.36%). This does not mean that the closest properties face a smaller total impact. In the combined model, the distance coefficient captures only the remaining distance effect after visibility has been controlled for. Many properties in the closest ring also fall into the highest visibility bands, so part of their overall discount is captured by the visibility coefficient rather than the distance coefficient. The total predicted effect for a close and highly visible property is therefore obtained by combining both channels, as shown in Figure 8.7.

Overall, the combined model shows that distance and visibility are related but separable. Visibility remains important after controlling for distance, and distance remains important after controlling for visibility. Which means that the price effect is not driven by proximity alone, and actual turbine visibility provides an additional and independent source of property price pressure.

8.2.3 Visibility Event Study

Design

This model asks how property prices evolve before and after turbines first become visible from a property location. Unlike the proximity event study, where treatment is based on turbine commissioning near the property, treatment here is based on the first year in which the viewshed rasters show positive turbine visibility at the property's location.

Key design choices:

- **Treatment year:** the first year in which at least one turbine becomes visible from the property location.
- **Event time:** event time is measured relative to the first visible year. The reference period is $k = -1$, meaning the year immediately before turbines first become visible.
- **Event Window:** the model uses $k \in [-4, +5]$. This window keeps enough observations before and after first visibility within the 2011–2019 transaction period.
- **Cohorts:** the analysis includes first-visibility cohorts from 2013 to 2017, since these cohorts have sufficient pre- and post-treatment support.
- **Control group:** the comparison group is made up of properties that never gain turbine visibility but are still located within 13 km of a turbine. This keeps the control group within the broader wind-farm geography.
- **Distance control:** $log(distancetothenearestturbine)$ is included so that the visibility estimates are not simply capturing physical proximity to turbines.
- **Cohort isolation:** treated postcodes are kept only when visible turbine exposure remains stable after the first visible year. This avoids cases where a property first sees one turbine cohort and later gains visibility of additional turbines, which would make event time harder to interpret.

Two specifications are estimated:

VisES-Binary: estimates the average event-time path for properties that gain any turbine visibility.

VisES-Bands: splits treated properties into the same visibility bands used in the static VisBands model (1–5, 6–10, 11–15, 16+ visible turbines). This tests whether the timing and persistence of the visibility effect differ by exposure intensity.

Results

VisES-Binary

Table 8.6 reports the binary visibility event-study estimates, and Figure 8.10 shows the same path visually.

Table 8.6: VisES-Binary results.

Event time k	Effect	95% CI	p-value
-4	**+5.35%**	[+4.88%, +5.81%]	≈ 0
-3	**+2.70%**	[+2.33%, +3.06%]	≈ 0
-2	**+1.20%**	[+0.88%, +1.53%]	< 0.001
-1	0.00% (reference)	—	—
0	**−0.97%**	[-1.26%, -0.67%]	< 0.001
1	**−1.06%**	[-1.36%, -0.76%]	< 0.001
2	**+0.14%**	[-0.17%, +0.45%]	0.381
3	**+1.34%**	[+0.99%, +1.68%]	< 0.001
4	**+1.97%**	[+1.60%, +2.34%]	≈ 0
5	**+2.37%**	[+1.95%, +2.80%]	≈ 0

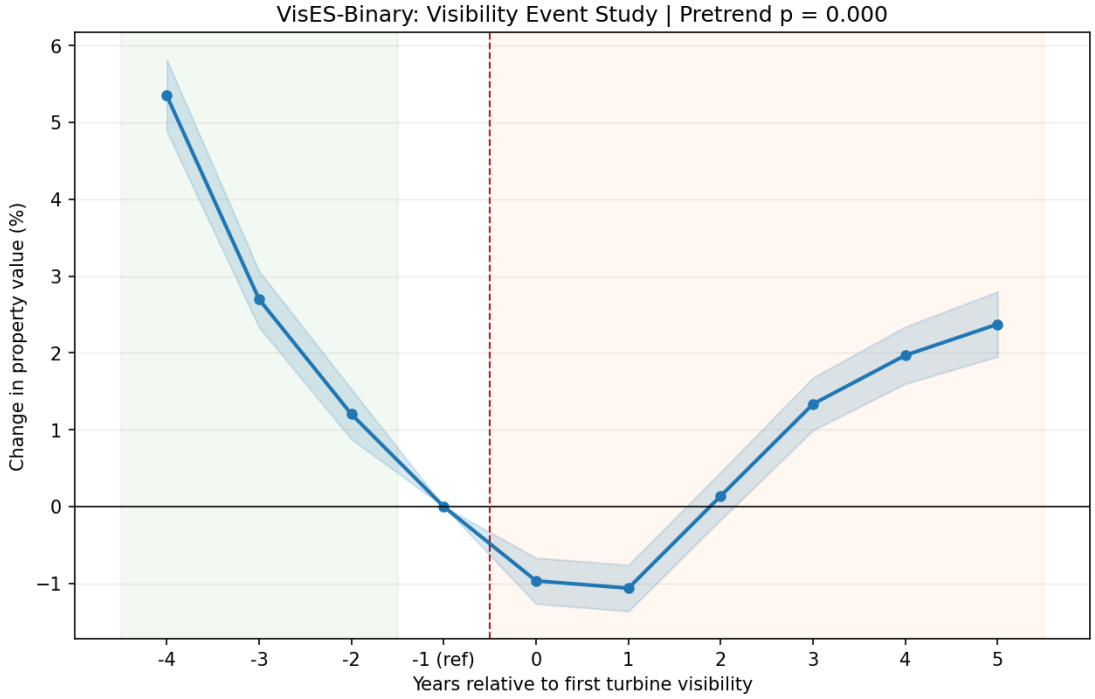

Figure 8.10: VisES-Binary event-study plot. $ATT(k)$ estimates and 95% confidence intervals by event time. Reference: $k = -1$.

The first important result is the pre-treatment pattern. Properties that later gain turbine visibility already have higher prices several years before first visibility. The estimate is starts from +5.35% at $k = -4$, and keeps decreasing until it reaches $k = -1$. This means the treated and control properties were not following parallel price paths before treatment.

This pre-trend is important for interpretation. The visibility event study should therefore not be read as a clean causal recovery path. Instead, it is more useful as descriptive evidence on timing. One *possible interpretation* is that properties that later become visible to turbines are often located in more open or elevated areas, which may have carried a **pre-existing price premium** before turbine visibility appeared. The decline before treatment may also reflect

anticipation, if buyers or sellers expected nearby turbine development before it became visible in the raster data.

After first visibility, the estimates turn negative in the short run, before becoming positive at later event times. This pattern suggests a short-run discount around the moment turbines first become visible, followed by a return toward the earlier premium. However, because the pre-treatment path is already strongly declining, this later increase **should not be interpreted as definitive evidence of market adaptation**. It is safer to describe it as a descriptive recovery pattern rather than a causal recovery effect.

Overall, the binary visibility event study supports the static visibility results in one limited sense: the estimates become negative immediately after first visibility. But the strong pre-trend means that the dynamic path mainly highlights the difficulty of identifying a clean timing effect for visibility.

VisES-Bands

The banded event study in Figure 8.11 adds an important layer to the binary result. The average path in VisES-Binary combines very different types of visibility exposure. Splitting the treated group by the number of visible turbines shows that low and high visibility follow different dynamic patterns, especially that Vis-Bands results showed that these groups behave differently.

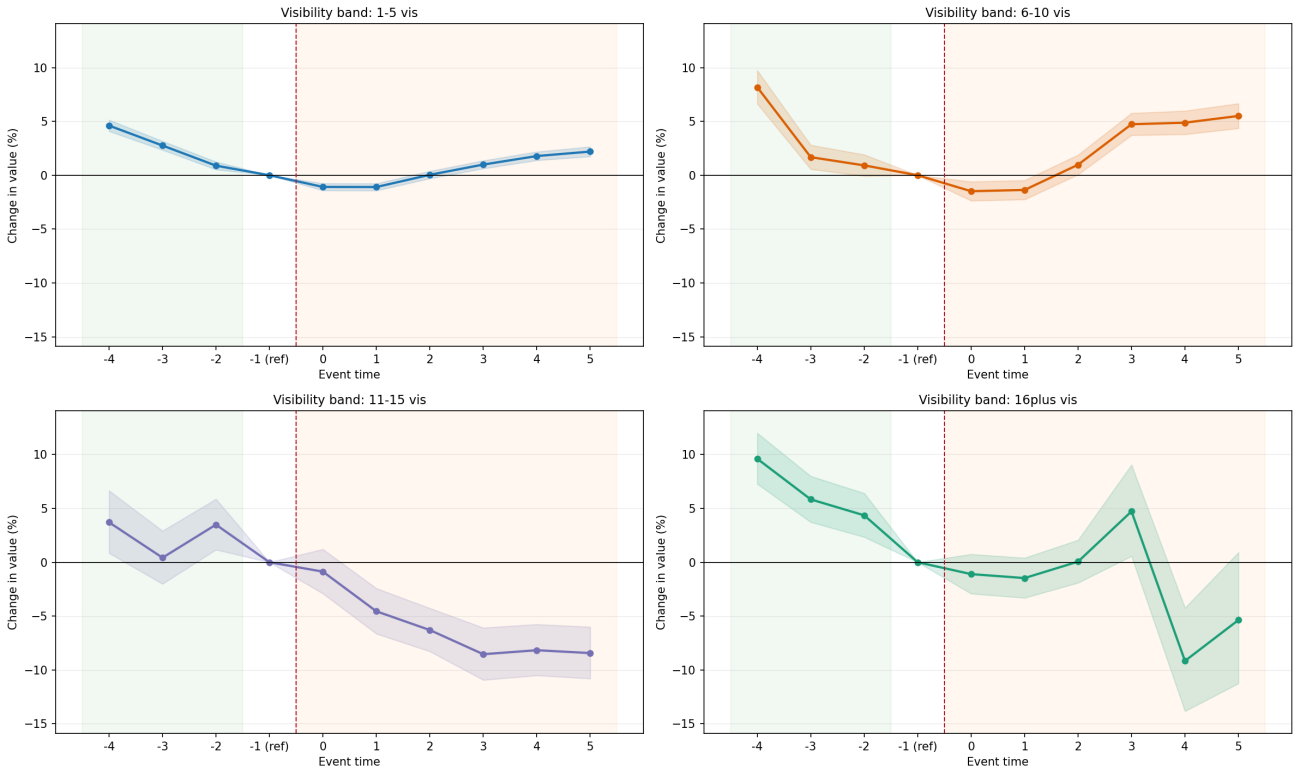

Figure 8.11: VisES-Bands event-study panel. Each sub-panel shows $ATT(k)$ for one visibility band (1–5, 6–10, 11–15, 16+). Reference: $k = -1$.

The low exposure visibility bands show relatively mild and short-lived negative estimates around first visibility. For properties seeing 1–5 or 6–10 turbines, the estimated discount appears mainly around the treatment year and the first year after treatment, before moving back toward

zero or positive values. This suggests that low and moderate visibility exposure is not associated with a persistent negative pattern in the event-study results.

While the high exposure visibility bands show a different pattern. We can clearly see that they have wider confidence intervals, reflecting the smaller number of observations shown in Figure 8.12. For the 11–15 group, the estimates become increasingly negative after first visibility, which is consistent with the static result that the visibility penalty becomes much larger once turbines form a more substantial part of the view. The 16+ group is harder to interpret. Despite it following the same general pattern as the 11–15 group, it jumps to a positive estimate around $k = 3$, followed by a sharp negative estimate at $k = 4$, so it should be treated as instability in the dynamic path rather than as a meaningful one-year recovery.

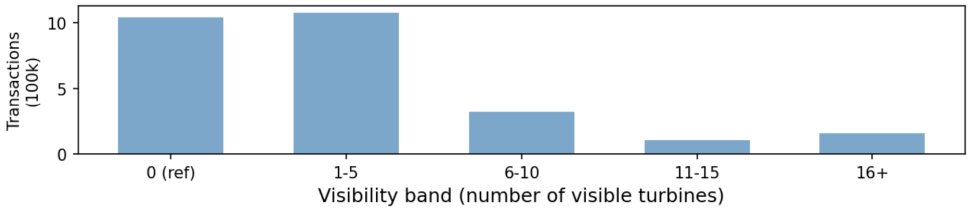

Figure 8.12: Transactions cound histogram based on the visibility bands.

Overall, the banded event study supports the same broad conclusion as the static visibility models: **low visibility exposure is associated with limited temporary price effects**, while **high visibility exposure is associated with larger more persistent price effects.** However, because the pre-trends are strong and the high-exposure bands have less data and wider confidence intervals, **these results should be used as descriptive evidence on timing and heterogeneity rather than as precise causal estimates for each event year.**

8.2.4 Results Summary — Visibility Models

The visibility models show that turbine visibility is associated with lower property prices, and that the effect becomes much stronger as the number of visible turbines increases. The banded model gives the clearest result: low visibility exposure has a small effect, while properties with 11 or more visible turbines face much larger discounts.

The distance-control models show that this is not only a proximity effect. Visibility remains important after controlling for distance, while distance also remains important after controlling for visibility. This suggests that proximity and visibility are related but separate channels.

The event-study results support the same general pattern, but they should be interpreted cautiously because of strong pre-trends. They are useful as descriptive evidence on timing and exposure intensity, rather than as clean causal dynamic estimates.

8.3 Shadow-Flicker Models

The shadow-flicker models add another possible way that turbines may affect nearby properties, beyond proximity and visibility. They test whether estimated annual shadow-flicker exposure is related to transaction prices within the near-turbine sample, and whether the result remains stable when the model controls for smaller local areas.

8.3.1 Shadow-Flicker DiD

Design

The model measures annual shadow-flicker exposure at each property and assigns each transaction the corresponding number of exposure hours for that year.

The model only uses properties within 3 km of a dated turbine because that is where the shadow-flicker exposure is most concentrated. Table 8.7 shows that only 2.3% of all merged transaction rows have positive exposure, compared with 19.4% within 3 km.

Table 8.7: Shadow-flicker sample and exposure summary.

Metric	Value
Merged transaction rows	2,478,283
Rows with positive shadow flicker	66,851
Positive shadow-flicker share, full sample	2.3%
Rows within 3 km	343,817
Positive shadow-flicker share within 3 km	19.4%
Binary TWFE regression observations	186,438
Treated observations in 3 km model	42,625
Treated post-commissioning observations	28,593

Figure 8.13 shows how shadow-flicker exposure varies across distance bands. Exposure is most concentrated in the 0–1 km band, where higher annual shadow-flicker buckets are most common. In the 1–2 km band, exposure is mostly low to moderate, while the 2–3 km band is dominated by very low exposure.

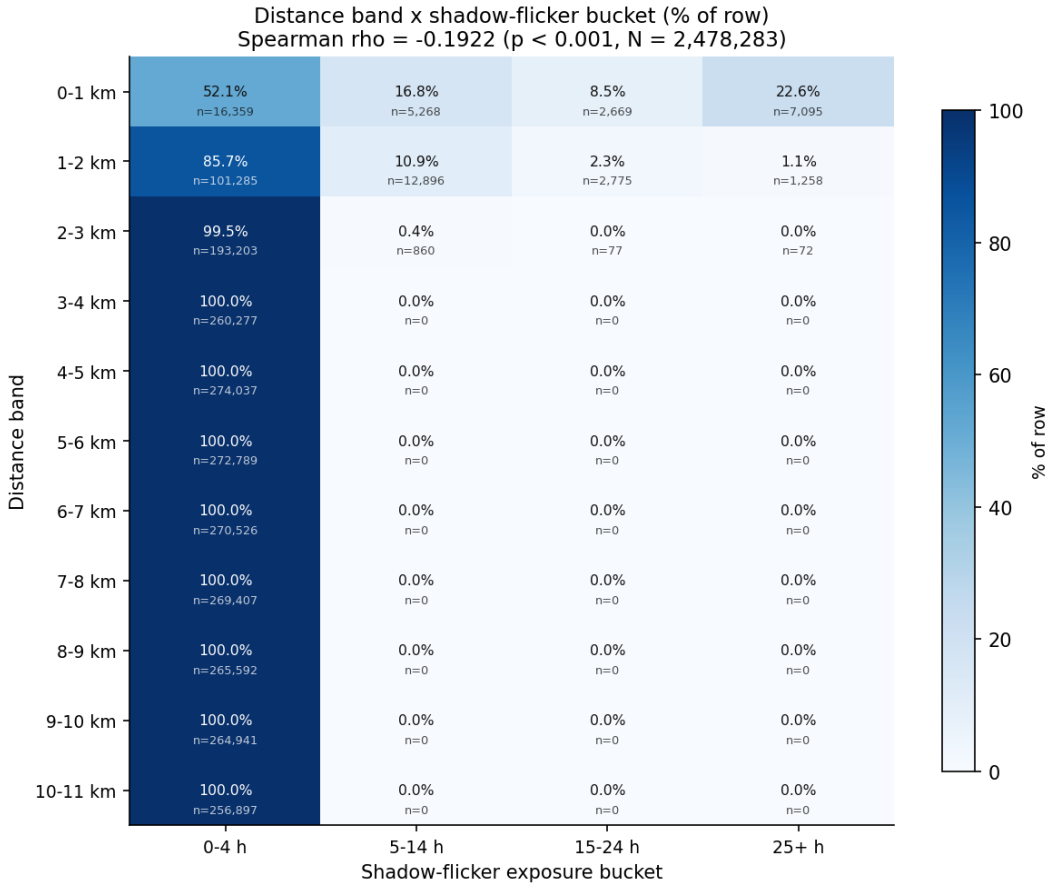

Figure 8.13: Distance band by shadow-flicker exposure bucket. Cells report the within-distance-band share of transactions in each annual shadow-flicker bucket, with transaction counts shown underneath.

Two static TWFE specifications are estimated, using postcode-area fixed effects as the headline and postcode-district fixed effects as a sensitivity check. Hedonic controls, year FE, and the fixed-effect convention follow Chapter 7.

SF-Binary — *Does any post-commissioning shadow flicker reduce prices?*
Compares properties with any post-commissioning shadow-flicker exposure to properties with zero shadow-flicker hours.

Treatment variable: `sf_post` = `(sf_hours > 0)` × `post`.

SF-Bands — *Does more exposure produce a larger penalty?*
Annual shadow-flicker hours grouped into bands of 0–1, 1–5, 5–10, 10–25, and 25+ hours per year, each estimated relative to zero shadow-flicker hours.

Results

The headline postcode-area model estimates that properties exposed to any post-commissioning shadow flicker sell for 2.0% less than otherwise comparable zero-flicker properties within the 3 km sample. The estimate is statistically significant, with a 95% confidence interval of [-2.6%, -1.3%] and p-value < 0.001.

Table 8.8: Shadow-flicker binary TWFE sensitivity to geographic fixed effects.

Fixed effects	Effect	95% CI	p-value	Observations	Treated post
Postcode area + year	-1.97%	[-2.64%, -1.29%]	1.76e-08	186,438	28,593
Postcode district + year	-0.23%	[-0.85%, 0.38%]	0.455	186,430	28,593

The main limitation is sensitivity to the geographic fixed-effect level. When postcode-district fixed effects are used instead of postcode-area fixed effects, the binary estimate falls to -0.2% and is not statistically significant. The area-level model is consistent with a shadow-flicker discount, but the district-level model suggests that part of the area-level effect may reflect local composition differences between exposed and unexposed places.

The SF-Bands model gives a similar but non-monotonic pattern. Very low exposure, 0–1 hours per year, produces an estimate of -0.6% that is statistically indistinguishable from zero. The middle bands are larger: -2.2% for 1–5 hours, -2.8% for 5–10 hours, and -2.1% for 10–25 hours. The 25+ hour band is negative at -1.4% but less precise. Figure 8.14 shows that this pattern is broadly consistent across both fixed-effect specifications.

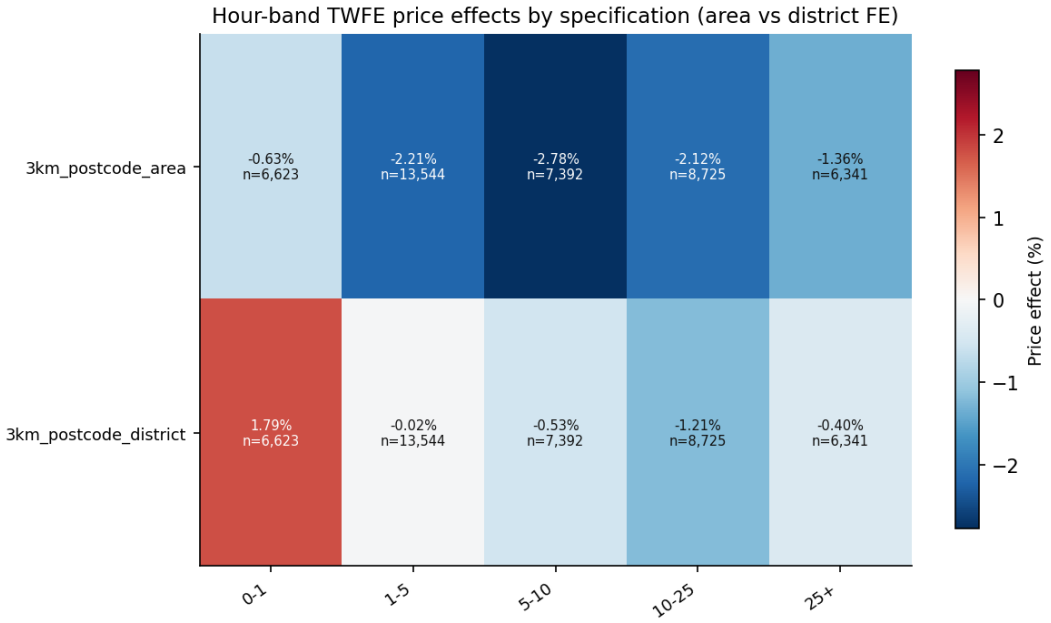

Figure 8.14: Shadow-flicker hour-band estimates by geographic fixed-effect specification. Effects are measured relative to zero shadow-flicker hours within the 3 km sample.

A pooled event-year diagnostic is shown in Figure 8.15. The pre-treatment estimates are already positive at $k = -4$ and $k = -3$, and the path declines toward the reference year before commissioning. This is not the flat pre-treatment pattern that would support a strong causal timing interpretation, so the event-year path is treated as descriptive evidence rather than the main causal estimate.

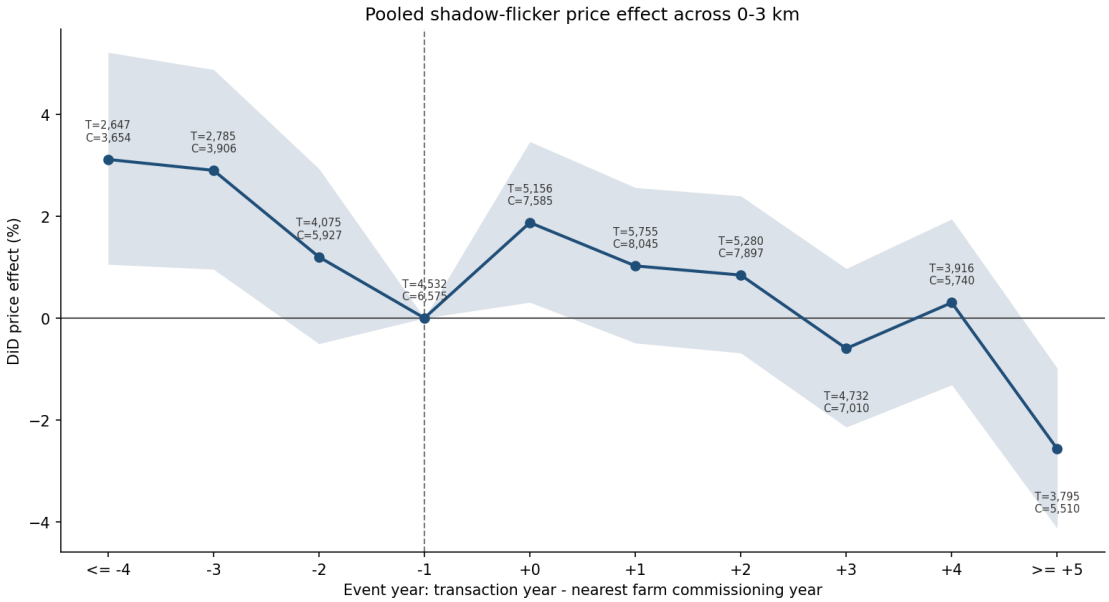

Figure 8.15: Pooled shadow-flicker event-year diagnostic across the 0–3 km sample. Effects are measured relative to $k = -1$, with the shaded band showing the 95% confidence interval.

Results Summary — Shadow Flicker Models

The shadow-flicker analysis provides suggestive evidence of a physical nuisance channel. The postcode-area estimate is economically meaningful and precisely estimated, but it is not robust to postcode-district fixed effects, placing shadow flicker below visibility in evidential strength.

CHAPTER 9

Conclusion

This thesis investigated whether onshore wind turbines are associated with changes in residential property prices in England and Wales, and whether these effects are better explained by simple proximity or by direct exposure channels such as visibility and shadow flicker. Using property transaction data from 2011 to 2019, wind turbine commissioning data, terrain-aware viewshed rasters, and simulated shadow-flicker exposure, the project applied Difference-in-Differences models to study how property prices vary before and after turbine exposure.

Overall, the results suggest that wind turbines are associated with lower nearby property prices, especially for properties closest to turbines and for properties with high visual exposure. The clearest result is that proximity and visibility are related but separate channels. Distance matters, but actual turbine visibility also remains important after controlling for distance. This means that the price effect is not driven by proximity alone. The results also provide suggestive evidence that shadow flicker may matter, although the shadow-flicker findings are less robust than the visibility findings.

9.1 Summary of Main Findings

The proximity models show a clear distance gradient. Properties within 1 km of a turbine experience the largest estimated discount, around -4.29% relative to the 10–11 km control ring. The estimates remain negative and statistically significant through the 5–6 km band, before becoming statistically indistinguishable from zero at 6–7 km. This suggests that living close to a turbine is associated with lower property prices, even before considering whether the turbine is actually visible.

The visibility models provide the strongest contribution of the thesis. Properties that gain any turbine visibility are associated with an average price discount of around -0.83%. More importantly, the effect becomes much stronger as the number of visible turbines increases. Properties with 1–5 visible turbines experience only a small estimated discount, while properties with 16 or more visible turbines are associated with a much larger discount of around -6.08%. This shows that visibility is not simply a binary issue. The intensity of visual exposure matters.

The combined visibility-distance models show that visibility remains important even after distance is controlled. The binary visibility estimate changes only slightly after adding distance controls, from -0.83%to about -0.78%. The visibility-band estimates also remain large for high-exposure properties after controlling for distance. At the same time, the distance gradient

remains negative in the closest bands after visibility is controlled. This supports the conclusion that proximity and visibility are separate mechanisms through which wind turbines may affect property prices.

The event-study models provide useful descriptive evidence about timing, but they must be interpreted cautiously. Both the proximity and visibility event studies show clear pre-treatment movement, meaning that the parallel trends assumption is not supported. For this reason, the dynamic estimates should not be read as clean causal evidence of when the price effect begins or whether prices recover over time. Instead, they show that treated and control properties were already following different paths before treatment, which limits the strength of any causal timing interpretation.

The shadow-flicker analysis provides weaker but still relevant evidence. In the postcode-area fixed-effect specification, properties exposed to post-commissioning shadow flicker are associated with a discount of around 2.0%. However, this estimate becomes small and statistically insignificant when postcode-district fixed effects are used. This means that shadow flicker may represent an additional nuisance channel, but the evidence is less robust than for visibility.

9.2 Answering the Research Questions

The main research question in Section 2.3 asked:

> *How do onshore wind farms affect residential property prices in England and Wales, and how much of that effect is attributable to visibility rather than proximity alone?*

together with the first two supporting questions, is answered most directly by the proximity, visibility, and distance-control models. The results suggest that onshore wind turbines are associated with lower nearby residential property prices, and that this effect is not explained by proximity alone. Properties closest to turbines show the largest proximity discounts, while properties with turbine visibility also experience lower prices. When distance is added to the visibility models, the visibility estimates shrink only slightly and remain statistically significant. This means that actual visibility is an independent channel, not just a proxy for living close to turbines.

The visibility results also show that the effect is not uniform. Gaining any turbine visibility is associated with a modest price discount, but the effect becomes much larger when many turbines are visible. This is especially clear in the visibility-band models, where properties with only a few visible turbines experience relatively small discounts, while properties with high visual exposure face much larger estimated price reductions. Therefore, the answer is not simply that visibility matters, but that the intensity of visibility matters.

The shadow-flicker question is answered more cautiously. The postcode-area model suggests that shadow flicker may be associated with an additional price discount, but the result is not robust when postcode-district fixed effects are used. For this reason, the thesis provides suggestive evidence that shadow flicker may represent an additional physical nuisance channel, but the evidence is weaker than for visibility and should not be treated as a central finding.

The final question asks whether property prices recover over time after turbine installation. The event-study models do not provide a strong causal answer to this question because the pre-treatment trends are not flat. Some estimates show short-run discounts followed by recovery-like patterns, especially in the visibility event study, but these patterns should be interpreted as descriptive rather than causal. The safest conclusion is that there is some evidence of recovery-like movements after exposure, but it cannot confidently establish that property prices recover over time.

9.3 Limitations

The main limitation of the study is the failure of the parallel trends assumption in the event-study models. Treated and control properties often show different price paths before turbine commissioning or before first visibility. This weakens the causal interpretation of the dynamic results and means that the event studies should mainly be used as diagnostics and descriptive evidence.

A second limitation is that wind turbines are not randomly placed. Turbines are often built in specific types of areas, such as rural, open, or elevated locations. These areas may already differ from other housing markets in ways that are difficult to fully capture, even with geographic fixed effects and property controls.

A third limitation concerns exposure measurement. Viewshed rasters provide a useful terrain-aware measure of visibility, but they do not capture every detail of how residents experience turbines, such as exact viewing direction, window placement, landscape quality, or whether a turbine is visually dominant from the property. Similarly, the shadow-flicker model estimates physical exposure, but the actual nuisance experienced by residents may depend on factors such as building orientation, whether people are usually at home during affected hours, and site-specific conditions such as surrounding trees or nearby buildings.

Finally, the study period is limited to 2011–2019. This provides a large number of transactions, but it also restricts how much pre-treatment anticipation and post-treatment long-run recovery or adaptation can be observed.

9.4 Final Remarks

This thesis shows that the local property-market effects of wind turbines cannot be understood through distance alone. While nearby properties show clear price discounts, the analysis also finds that actual turbine visibility remains important even after proximity is controlled. The strongest evidence comes from high-visibility cases, where turbines form a more substantial part of the surrounding landscape.

The findings should still be interpreted cautiously, especially because the event-study models show pre-treatment differences between treated and control properties. For that reason, the thesis does not claim to establish a precise causal recovery path over time. Its main contribution is instead to show that proximity and visibility are related but distinct exposure channels, and

that terrain-aware visibility measures can add important information to the study of wind farm impacts on residential property prices.

9.5 Future Work

Several directions could extend this work. First, the shadow-flicker analysis should be investigated further. The current results suggest that shadow flicker may be associated with lower property prices, but the estimates are sensitive to the geographic fixed-effect specification. Future work could test alternative exposure definitions, improve the shadow-flicker simulation, and examine whether the results become more stable with different samples or more detailed local controls.

Second, future research should examine the source of the parallel trends violations found in the event-study models. The treated properties were already following different price paths before treatment, but the underlying reason is not fully clear. This could reflect anticipation effects, planning-stage information, local housing-market differences, scenic or elevated locations, or other unobserved factors. Understanding this issue would strengthen the causal interpretation of the results.

Third, the analysis could be extended by adding wind turbine noise models. Noise was not included in the maintained empirical results, but it remains one of the main concerns raised in relation to nearby wind farms. Modelling noise exposure would make it possible to compare visual, proximity, shadow-flicker, and noise channels more directly.

A further step would be to estimate the combined effect of all exposure channels in one consistent framework. This would require comparing the proximity, visibility, shadow-flicker, and noise models carefully to avoid double counting, since the channels are related. A combined model could help separate the individual contribution of each channel while also estimating the overall effect of turbine exposure on property prices.

Finally, future work could turn the analysis into a practical visualization tool for planning and policy use. Such a tool could allow planners, developers, or local authorities to estimate the expected total effect of a proposed turbine or wind farm, as well as the contribution of each channel separately. This could support more transparent planning decisions and help estimate fair compensation where local property-market impacts are expected.

CHAPTER 10

Bibliography

[1] WIMBY Consortium, *Confidential wind turbine dataset*, Unpublished, Wind In My Backyard (WIMBY), Horizon Europe Project, 2026. [Online]. Available: https://wimby.eu.

[2] Department for Energy Security and Net Zero. "DESNZ public attitudes tracker: Headline findings, Spring 2025, UK," GOV.UK, Accessed: May 14, 2026. [Online]. Available: https://www.gov.uk/government/statistics/desnz-public-attitudes-tracker-spring-2025/desnz-public-attitudes-tracker-headline-findings-spring-2025-uk.

[3] L. Susskind, J. Chun, A. Gant, C. Hodgkins, J. Cohen, and S. Lohmar, "Sources of opposition to renewable energy projects in the United States," *Energy Policy*, vol. 165, Jun. 2022, Art. no. 112922. DOI: 10.1016/j.enpol.2022.112922.

[4] F. Rankl. "Planning for onshore wind," House of Commons Library, Accessed: May 14, 2026. [Online]. Available: https://commonslibrary.parliament.uk/research-briefings/sn04370/.

[5] S. Gibbons, "Gone with the wind: Valuing the visual impacts of wind turbines through house prices," *Journal of Environmental Economics and Management*, vol. 72, pp. 177–196, Jul. 2015. DOI: 10.1016/j.jeem.2015.04.006.

[6] W. Guo, L. Wenz, and M. Auffhammer, "The visual effect of wind turbines on property values is small and diminishing in space and time," *Proceedings of the National Academy of Sciences*, vol. 121, no. 13, Mar. 2024, Art. no. e2309372121. DOI: 10.1073/pnas.2309372121.

[7] H.-Y. Chen, "Wind turbine noise, shadow flicker, and property prices: An exposure-based study in the Netherlands," MSc thesis, Utrecht University, Utrecht, The Netherlands, 2025.

[8] C. Andersen and T. Hener, "Wind turbines, shadow flicker, and real estate values," *Environmental and Resource Economics*, vol. 88, no. 3, pp. 731–759, Mar. 2025. DOI: 10.1007/s10640-024-00947-x.

[9] S. Jarvis, "The economic costs of NIMBYism: Evidence from renewable energy projects," *Journal of the Association of Environmental and Resource Economists*, vol. 12, no. 4, pp. 983–1022, Jul. 2025. DOI: 10.1086/732801.

[10] M. I. Dröes and H. R. A. Koster, "Renewable energy and negative externalities: The effect of wind turbines on house prices," *Journal of Urban Economics*, vol. 96, pp. 121–141, Nov. 2016. DOI: 10.1016/j.jue.2016.09.001.

[11] M. Schütt, "Wind turbines and property values: A meta-regression analysis," *Environmental and Resource Economics*, vol. 87, no. 1, pp. 1–43, Jan. 2024. DOI: 10.1007/s10640-023-00809-y.

[12] E. J. Brunner, B. Hoen, J. Rand, and D. Schwegman, "Commercial wind turbines and residential home values: New evidence from the universe of land-based wind projects in the United States," *Energy Policy*, vol. 185, Feb. 2024, Art. no. 113837. DOI: 10.1016/j.enpol.2023.113837.

[13] S. Sims, P. Dent, and G. R. Oskrochi, "Modelling the impact of wind farms on house prices in the UK," *International Journal of Strategic Property Management*, vol. 12, no. 4, pp. 251–269, Oct. 2010. DOI: 10.3846/1648-715X.2008.12.251-269.

[14] C. U. Jensen, T. E. Panduro, T. H. Lundhede, A. S. E. Nielsen, M. Dalsgaard, and B. J. Thorsen, "The impact of on-shore and off-shore wind turbine farms on property prices," *Energy Policy*, vol. 116, pp. 50–59, May 2018. DOI: 10.1016/j.enpol.2018.01.046.

[15] University College London, *UCL property transactions dataset, 2011–2019*, Geospatial shapefiles, unpublished, 2019.

[16] DTU Energy. "Wind energy research," Technical University of Denmark, Accessed: Apr. 26, 2026. [Online]. Available: https://www.energy.dtu.dk/.

[17] H.-H. Chen, *wimby-sf: A package for the WIMBY shadow flicker functions*, Python package, 2025.

[18] L. Sun and S. Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, vol. 225, no. 2, pp. 175–199, 2021. DOI: 10.1016/j.jeconom.2020.09.006.