

The Effects of Wind Turbines on Housing Prices in the Netherlands

Assessing Onshore, Offshore, and Heterogeneous Effects

July 8, 2026

Tan Dieltjes



The Effects of Wind Turbines on Housing Prices in the Netherlands

Assessing Onshore, Offshore, and Heterogeneous Effects

Master thesis submitted to Delft University of Technology in partial fulfilment of the requirements for the degree of

MASTER OF SCIENCE

Engineering and Policy Analysis
Faculty of Technology, Policy and Management
Delft University of Technology

Student name: Tan Dieltjes
Student number: 6107028
Thesis defence date: 16 July 2026

Graduation committee

Dr. S.J. Pfenninger-Lee, TBM-ESS-EI
Dr. R. Stöllinger, TBM-VTI-ETI
Dr. T. Chatzivasileiadis, TBM-MAS-BA

Chair
First supervisor
Second supervisor

Executive Summary

The Netherlands is rapidly expanding its wind energy capacity both onshore and offshore as part of its broader energy transition. Although wind capacity has more than doubled from 2019 to 2024, residents experience localised externalities from wind turbines, including visual impact, noise, and shadow flicker. These costs are often capitalised into housing prices. Therefore, the placement of turbines involves balancing the progress of the energy transition with the burdens imposed on local communities, emphasising the importance of reliable evidence on the extent and distribution of these effects to inform spatial planning strategies.

Among the many externalities of wind turbines, ranging from the effects on landscape to impacts on nature, housing prices are particularly relevant. They can be directly observed in housing transaction records and combine various perceived local disamenities into a single, quantifiable measure. Furthermore, house prices are linked to Dutch compensation policy: the Omgevingswet, or the Dutch Environment and Planning Act, sets a societal risk threshold of four per cent of the property's value. If the depreciation falls below this threshold, it is not eligible for any compensation. Therefore, analysing house prices provides a tangible, policy-relevant insight into who bears the local costs of the energy transition. This aligns with the Engineering and Policy Analysis (EPA) programme's goal of employing quantitative methods to inform decision-making regarding grand societal challenges.

While prior studies have examined how turbines impact housing prices in the Netherlands, two important gaps still persist. First, the impact of offshore turbines on housing prices has not been examined in the Dutch context. Second, existing studies capture spatial variation through geographic categories but leave the role of socioeconomic factors almost entirely unexplored. Additionally, the Dutch estimates are based on methods that could introduce bias as turbine adoption is staggered. This raises questions about the current magnitude of these effects when staggered adoption is taken into account, and how these effects are distributed among communities exposed to turbines.

This thesis estimates the effects of onshore wind turbines using a Difference-in-Differences (DiD) design, based on a sample of 1.83 million housing transactions from 2010 to 2023 from the Nederlandse Vereniging van Makelaars (NVM), merged with WindStats turbine data and neighbourhood-level Socioeconomic Status (SES) data. To extend current literature and to account for the staggered nature of Dutch wind turbine construction, the analysis uses the [Sun & Abraham \(2021\)](#) estimator and systematically controls for socioeconomic composition, which prior Dutch studies did not. A new addition to current literature is the offshore component, which is deliberately descriptive: It documents coastal price trends associated with offshore wind farm development and clarifies why the Dutch data do not support credible causal identification.

The results show that the effects of onshore wind turbines are small, not robust across different control group definitions, highly localised, and statistically indistinguishable from zero under most alternatives. In the preferred specification, they are found to slightly lower nearby property prices by about 0.8% within two kilometres. This is roughly a fifth of the four per cent compensation threshold of the Omgevingswet, so most affected homeowners do not qualify for any compensation. The estimated price effect is smaller than previous Dutch estimates, primarily because the sample from 2010 to 2023 limits the identifying variation needed to estimate causal effects. The causal interpretation is further limited by imperfect pre-trends, though robustness checks suggest these do

not drive the estimated effect. These effects are unlikely to be lower due to changes in perceptions of wind energy, as [Dröes & Koster \(2019\)](#) have shown that willingness to pay to live next to turbines is generally stable over time.

The most robustly tested finding in the heterogeneity analysis is that additional turbines do not produce larger effects, meaning they do not lead to higher price effects, while also reducing the number of properties exposed to turbines. Subgroup analyses indicate effects are strongest near turbines and in densely populated regions, although these differences between subgroups are not statistically significant. Additionally, while turbines are disproportionately sited in neighbourhoods with lower transaction prices, associated with lower SES, whether the price effect itself differs by SES could not be established. In all these analyses, none of the estimated effects approaches the four per cent compensation threshold set by the *Omgevingswet*. The offshore wind analysis remains descriptive rather than causal due to the small number of offshore parks, their recent clustered construction, a limited coastal sample, and the lack of visibility data, all of which prevent reliable causal identification.

These findings inform a siting strategy focused on reducing housing-market costs of turbines. The only well-supported principle is that clustering turbines concentrates price effects on fewer properties, without increasing their magnitude relative to single-turbine installations. The weaker, non-significant patterns suggest, cautiously, that turbines should be placed away from urbanised areas. Since onshore turbines are already mostly located in communities with lower prices, associated with lower SES, an efficiency-based siting strategy that ignores distributional risks will exacerbate existing inequalities in turbine placement. Where turbines are sited is therefore ultimately a political choice that should weigh minimising housing market impacts against the fair distribution of turbine placements; this thesis explicitly highlights this trade-off by analysing both the magnitude of effects and the populations most affected. More broadly, the small estimated effects imply that turbines have a limited impact on the Dutch housing market, inviting a shift in the debate toward more difficult-to-measure externalities, such as landscape and health concerns.

Contents

Executive Summary	i
List of Figures	v
List of Tables	vii
Glossary	ix
1 Introduction	1
1.1 Background and Context	1
1.2 Problem Statement and Knowledge Gaps	2
1.3 Research Objective and Questions	2
1.4 Research Approach	3
1.5 Relevance and Link to EPA	3
1.6 Reading Guide	4
2 Literature Review	5
2.1 Hedonic Pricing Evidence on Wind Turbines	5
2.1.1 Mixed International Findings	5
2.1.2 Limits of Hedonic Pricing	6
2.2 Difference-in-Differences: Onshore Evidence	7
2.2.1 Baseline Dutch Effects	7
2.2.2 What Drives the Effect?	8
2.2.3 Temporal Dynamics	9
2.2.4 Modern Estimators	9
2.2.5 Spatial Heterogeneity	10
2.2.6 Socioeconomic Heterogeneity	11
2.3 Difference-in-Differences: Offshore Evidence	11
2.4 Synthesis and Research Gaps	13
2.5 Research Questions and Hypotheses	14
3 Data Sources and Analysis	17
3.1 Data Sources	17
3.1.1 Housing Transactions	17
3.1.2 Wind Turbines	19
3.1.3 Socioeconomic Elements	20
3.1.4 Geometries	22
3.2 Construction of Analysis Dataset	22
3.3 Exploration of the Working Dataset	24
3.3.1 Treatment Exposure	24
3.3.2 Identifying Variation	26
3.3.3 Socioeconomic Composition of Dataset	29

4	Research Design and Methodology	37
4.1	Identification Strategy	37
4.2	Baseline Specification	38
4.3	Robustness Analyses	41
4.4	Heterogeneity Analyses	41
4.5	Socioeconomic Score Extensions	42
4.6	Descriptive Offshore Analysis	42
5	Results	45
5.1	Onshore Effects	45
5.1.1	Treatment Effects Under the Baseline Specification	45
5.1.2	Robustness Analyses	47
5.1.3	Heterogeneity Analyses	52
5.1.4	Socioeconomic Score Extensions	57
5.2	Descriptive Offshore Analysis	60
6	Discussion	67
6.1	Main Findings in Context	67
6.2	Differences From Prior Dutch Studies	68
6.3	Heterogeneity and Policy Implications	70
6.4	Distributive Implications	71
6.5	The Effects in Perspective	72
6.6	General Limitations	74
7	Conclusion	77
7.1	Answers to Research Questions	77
7.2	Reflection	78
7.3	Future Research	79
	References	81
A	AI Reflection	85
B	Additional Descriptive and Socioeconomic Data Analysis	87
B.1	Socioeconomic Data	87
C	Additional Specifications	95
C.1	Two-Way Fixed Effects Replication	95
C.2	Heterogeneity Analyses	95
C.3	Socioeconomic Score Additional Specifications	100
D	Additional Offshore Descriptive Elements	103
D.1	Offshore Descriptives	103
E	Housing Attributes	107
E.1	Housing Attributes of NVM Data	107
E.2	Baseline Specification	108

List of Figures

1.1	Installed onshore and offshore wind power capacity in the Netherlands.	1
3.1	The national spatial coverage of all cleaned NVM transactions between 2010 and 2023.	18
3.2	The annual transactions recorded by the NVM between 2010 and 2023.	19
3.3	The number of turbines constructed each year by turbine type during the sample period of the housing transactions data.	20
3.4	The locations of all turbines ever constructed in the Netherlands categorised by turbine type.	21
3.5	The tip height of newly constructed wind turbines over the years, organised per type.	21
3.6	The temporal variation in the transaction-weighted SES scores from 2014 onwards.	22
3.7	The coastline and water layers used to compute coastal distance and classify water-related geography.	23
3.8	A heatmap showcasing the distribution of the share in properties exposed to onshore turbines.	25
3.9	The percentage of housing transactions that have 1 or more turbines within 2 kilometres over the years.	26
3.10	The number of onshore turbine constructions compared to the number of PC4 areas, newly exposed to onshore turbines.	27
3.11	Onshore turbines per construction period scattered over the Netherlands.	28
3.12	Distributions of the number of pre- and post-turbine construction observations. . .	29
3.13	Decompositions of the merged sample regarding the effective identifying sample for causal inference.	29
3.14	A comparison between the number of observations before, after, and on the minority side of the turbine exposure for identifying PC4 regions.	30
3.15	The distribution of socioeconomic status across the Netherlands.	31
3.16	The average socioeconomic status in the Netherlands demonstrates an s-curve over time.	32
3.17	The socioeconomic score quintiles highlight that a higher score is linked to a higher transaction price, conditioned on living area.	32
3.18	The median price index of socioeconomic score quintiles mostly highlights similar trends for ever- and never-treated properties, but price growth is more pronounced in neighbourhoods with lower socioeconomic scores.	33
3.19	The standardised differences between treated and non-treated properties exhibit that differences in socioeconomic status exist.	34
3.20	The socioeconomic score, adjusted for national trends, shows diverging socioeconomic trends over time.	34
3.21	The socioeconomic score, adjusted for urbanisation trends, shows diverging socioeconomic trends over time.	35
4.1	An illustration of the treatment effect of the Difference-in-Differences identification strategy.	38

5.1	Event study of the baseline specifications with the Sun & Abraham (2021) estimator with $t = -1$ as the reference year.	47
5.2	The turbine placement and construction numbers over the years in the Netherlands.	62
5.3	The price differences between ever-treated and control units and their robustness.	62
5.4	The aggregate socioeconomic score over the years for Dutch coastal communities.	63
5.5	The descriptive differences in prices for coastal communities between treated and control properties relative to their treatment year, where $t = -1$ is the reference year. It also highlights the total observations per event time.	64
5.6	The bars report the share of 6-digit postal code (PC6) areas with treated observations with no pre-treatment observations, no post-treatment observations, and often no more than 2 post-treatment years.	65
5.7	The differences between treated and control units after accounting for different time effects, comparing global versus municipal time effects.	65
6.1	Time losses caused by onshore turbines, with price effects of -0.8% , graphed over time.	68
6.2	The year-on-year growth of housing prices across the Netherlands over time.	73
B.1	The financial SES scores distributed across the Netherlands.	88
B.2	The educational SES scores distributed across the Netherlands.	89
B.3	The change in the partial and the aggregate SES scores between 2014 and 2023 shown across the Netherlands.	90
B.4	The change in additional SES scores, related to finance, education, and the aggregate score, between 2014 and 2023 across the Netherlands.	91
B.5	The spread of SES variables highlighting heterogeneous neighbourhoods and the 4-digit postal code (PC4) areas that do not contain SES data.	92
B.6	The correlation matrix of time-varying SES in the Netherlands between 2014 and 2023.	93
B.7	The median transaction price growth over the years in the Netherlands with base year 2010.	94
D.1	The number of properties ever exposed to offshore wind turbines in the Netherlands by province.	103
D.2	The price differences between ever-treated and control units and their robustness, now adjusted at the PC4 level.	104
D.3	The descriptive differences in prices between treated and control properties at the PC4 level, relative to their treatment year, where $t = -1$ is the reference year. It also highlights the total observations per event time.	104
D.4	The bars report the share of PC4 areas with treated observations with no pre-treatment observations, no post-treatment observations, and at most two post-treatment years.	105
D.5	The differences between treated and control units after accounting for different time effects, comparing global versus municipal time effects at the PC4 level.	105

List of Tables

3.1	The number of observations of the analysis sample after each step of data merging and cleaning.	24
3.2	Distribution of pre-treatment, post-treatment, and minority side observations within identifying PC4 areas.	30
4.1	Descriptive statistics: treated and control transactions.	40
5.1	The baseline onshore effects of wind turbines on housing prices.	46
5.2	The onshore effects of wind turbines on housing prices, excluding potential anticipation effects up to three years prior to treatment.	48
5.3	The onshore effects of wind turbines on housing prices, excluding event times with pre-trend deviations.	49
5.4	The onshore effects of wind turbines on housing prices with alternative control group definitions.	50
5.5	The onshore effects of wind turbines on housing prices if the treatment range is increased to 2.5 kilometres.	51
5.6	The onshore effects of wind turbines on housing prices depending on the number of turbines during the treatment period.	52
5.7	The interaction effects of all the presented subgroups in heterogeneity analyses using a Two-Way Fixed Effects (TWFE) specification.	54
5.8	The onshore effects of wind turbines on housing prices grouped by turbine height groups.	55
5.9	The onshore effects of wind turbines on housing prices grouped by urbanisation category.	56
5.10	The onshore effects of wind turbines on housing prices depending on the distance to the closest turbine during the treatment period.	57
5.11	The baseline onshore effects of wind turbines on housing prices when controlling for SES variables.	58
5.12	The onshore effects of wind turbines on housing prices grouped by urbanisation category, with and without the standardised aggregated SES as a control variable.	59
5.13	Descriptive statistics regarding coastal transactions with turbine exposure.	60
C.1	The TWFE replication of the baseline specification, highlighting the effects of wind turbines on housing prices with a similar method to prior literature.	95
C.2	The onshore effects of wind turbines on housing prices depending on the number of turbines during the treatment period with PC4 fixed effects.	96
C.3	The baseline estimate of onshore turbines divided by sale year.	97
C.4	The onshore effects of wind turbines on housing prices grouped by urbanisation category.	98
C.5	The onshore effects of wind turbines on housing prices grouped by turbine height groups, with PC4 fixed effects.	99
C.6	The onshore effects of wind turbines on housing prices grouped by urbanisation category with SES controls.	100

C.7	The effects of wind turbines on housing prices with SES as control variables, with year-municipality and PC4 fixed effects.	101
C.8	The onshore effects of wind turbines split per aggregate SES quartile.	102
E.1	The types of housing characteristics used as controls in DiD analyses.	107
E.2	The full regression table of the baseline results that include housing characteristics.	109

Glossary

Average Treatment Effect of the Treated (ATT) The average causal effect of treatment among units that were treated. 20, 38, 39, 46–60, 69, 74, 95–102, 111

Central Bureau of Statistics (CBS) The independent national statistical office of the Netherlands, responsible for collecting, processing, and publishing official statistical information about the country. 3, 20, 42, 57, 59

Difference-in-Differences (DiD) A quasi-experimental method that estimates the causal effect of a treatment by comparing changes in outcomes over time between treated and untreated units, assuming that their untreated trends would have evolved in parallel. i, viii, 2–10, 12, 13, 26, 37–40, 43, 45, 48, 53, 67, 69, 79, 80, 107

Engineering and Policy Analysis (EPA) An interdisciplinary MSc programme taught at Delft University of Technology focused on grand societal challenges using both qualitative and quantitative methods to inform policy. i, 4

Independent and Identically Distributed (IID) An assumption in statistics that observations are drawn independently from the same distribution. 29, 40, 74

Minimum Detectable Effect (MDE) The smallest true effect size that a statistical test can detect with a given power level, significance level, and the estimated standard errors. 45

Nederlandse Vereniging van Makelaars (NVM) The Dutch association of real estate agents, which is the source of the housing transaction data with an approximately 70% market share. i, v, 3, 8, 17–19, 22, 29, 51, 69, 74, 78, 85, 98, 107

4-digit Postal Code (PC4) Postal code regions with 4 digits, such as 1234. These are less granular than postal code regions with more digits, such as PC5 or PC6. vi, vii, 7, 24, 26–30, 32, 33, 39, 40, 42, 45, 46, 48, 49, 53, 54, 58, 60, 61, 67, 70, 75, 77, 87, 92, 95, 99, 100, 103–105, 108

6-digit Postal Code (PC6) Postal code regions with 6 digits, such as 1234AB. These are the most fine-grained postal codes for data analysis. vi, 7, 17, 39–41, 46–51, 53, 54, 61, 62, 64, 65, 67, 95, 100, 108

Socioeconomic Status (SES) A wide variety of socioeconomic variables measuring the socioeconomic status of a neighbourhood, including variables on income, education, and labour market history. i, ii, v–viii, 3, 11, 14, 15, 20, 22, 24, 29, 30, 32, 33, 42, 57–61, 63, 67, 71, 72, 75, 77, 78, 80, 85, 87–93, 100–102

Two-Way Fixed Effects (TWFE) a Difference-in-Differences specification including both unit and time fixed effects, the conventional estimator used in prior wind-housing literature. vii, 2, 9, 10, 13, 37, 38, 42, 46, 52, 54, 69, 79, 80, 95, 100

Introduction

1.1. Background and Context

The European Union's transition to a low-carbon energy system has accelerated the deployment of renewable energy infrastructure across Europe, with wind power playing a central role in decarbonisation initiatives and offshore planning goals (European Commission, 2025). In the Netherlands, ambitious climate goals, concerns about energy security, and efforts to reduce fossil-fuel dependence are driving a substantial expansion of onshore and offshore wind capacity (Rijksoverheid, n.d.-c). As part of the broader grand challenges of climate change mitigation and sustainable energy transitions, recent years have seen substantial growth in installed wind turbine capacity, as shown in Figure 1.1. Between 2019 and 2024, the installed wind capacity increased at an average annual rate of 21.2%, more than doubling the installed capacity (Centraal Bureau voor de Statistiek, 2025a). The total share of electricity production by wind energy was 26.9% in 2024, of which 40.5% was already provided by offshore turbines (Centraal Bureau voor de Statistiek, 2025a; Centraal Bureau voor de Statistiek, 2025b). This is accompanied by the development of large offshore zones such as the Hollandse Kust (West) and IJmuiden Ver areas (Rijksoverheid, n.d.-b), and by a recent trend toward increasingly large turbines. As a result, wind energy infrastructure increasingly interacts with densely populated neighbourhoods, coastal landscapes, and land-use pressures that define the Dutch spatial planning system. With several large offshore projects underway and more planned, understanding potential coastal impacts has become increasingly urgent.




Figure 1.1: Installed onshore and offshore wind power capacity in the Netherlands.

Note: Data from Centraal Bureau voor de Statistiek, 2025a.

1.2. Problem Statement and Knowledge Gaps

The rapid expansion of wind energy also raises policy challenges. While wind turbines contribute significantly to national climate objectives, their siting often involves complex trade-offs between environmental benefits and perceived local externalities. Residents may express concerns about landscape changes, noise, shadow flicker, and impacts on property values (Mattmann et al., 2016; Onakpoya et al., 2015). Qualitative work on wind farms highlights concerns regarding noise, visual impacts, and scenic vistas, often discussed in terms of various stigmas around nuisance and area image (Torzewski, 2016). Similar dynamics emerge in coastal contexts, where debates about offshore wind frequently centre on impacts on tourism, fisheries, and the visual perception of seascapes (Russell et al., 2021). This thesis focuses on housing prices as the primary outcome. House prices, directly observable through market transactions, aggregate various local disamenities into a single, measurable indicator that can be compared across the Netherlands. Additionally, they are closely linked to Dutch compensation policy, as depreciation influences residents' eligibility for any compensation.¹ Housing prices provide a concrete and policy-relevant view of local wind energy costs. This makes it possible to examine not only whether turbines affect housing prices, but also whether these costs are equitably distributed across the Netherlands. Against this background, the core problem this thesis addresses is twofold: existing Dutch studies have documented adverse effects of onshore turbines on nearby housing markets, but they rely on conventional Two-Way Fixed Effects (TWFE) Difference-in-Differences (DiD) designs that are ill-suited to staggered adoption; and it remains unclear how housing-market impacts from today's deployment of onshore and offshore wind are distributed across regions and socioeconomic groups.

Despite the policy relevance of these questions, the scientific evidence base remains unbalanced. Research on the housing-market effects of onshore wind turbines is extensive, and recent DiD studies in the Netherlands provide credible estimates of localised price impacts (Dröes & Koster, 2016; Dröes & Koster, 2021; Eichholtz et al., 2023). However, two essential gaps persist. First, evidence on the effects of offshore wind turbines on coastal housing markets is limited, and none exists for the Netherlands, despite the rapid expansion of offshore capacity. Second, the literature provides only a limited understanding of what drives heterogeneity in treatment effects, particularly along socioeconomic dimensions; it remains largely unknown whether and how socioeconomic conditions affect the impact of wind turbines on housing prices, and for whom the costs of turbine expansion are largest. Together, these gaps limit our understanding of how onshore and offshore wind turbines are capitalised into housing markets, and how spatial context and their distribution across different communities shape these effects (Jenkins et al., 2016). From an energy-justice and distributive-fairness perspective, unequal distributions of benefits and burdens are only acceptable if local communities are meaningfully involved in the decision-making process and compensated for the incurred losses (Jenkins et al., 2016; Vallor et al., 2022).

1.3. Research Objective and Questions

This study addresses these gaps by examining how onshore and offshore wind turbines affect housing prices in the Netherlands and whether these effects vary across regions and socioeconomic contexts. By combining nationwide housing transaction data with detailed turbine characteristics and econometric methods, the analysis contributes to ongoing scientific debates in environmental and urban economics on the valuation of energy infrastructure. Specifically, the study expands the descriptive evidence on offshore wind associations in a densely populated European context and determines heterogeneity across Dutch contexts. In doing so, it provides an empirical basis for more equitable spatial planning and highlights whether local homeowners are likely to be compensated to offset turbine-related costs, while the wider population benefits from the Dutch energy transition.

¹The Omgevingswet, or the Dutch Environment and Planning Act, sets a societal risk threshold of four per cent of the property's value. If the depreciation falls below this threshold, it is not eligible for any compensation.

Instead of assessing fairness through normative standards, the analyses will quantify where and for whom losses appear in transaction prices. This will help policymakers identify potential localised burdens on nearby residents and support the design of siting policies that reduce turbine-related externalities. These policies should not only be efficient but also perceived as equitable. Given these objectives, the main research question of this thesis is: ***How do onshore and offshore wind turbines affect housing prices in the Netherlands, and how do these effects vary across space and socioeconomic contexts?***

To address this main question, the analysis is structured around three sub-questions:

- RQ1:** To what extent do onshore wind turbines affect nearby housing prices in the Netherlands?
- RQ2:** What descriptive patterns can be observed in coastal housing prices around offshore wind farm development in the Netherlands?
- RQ3:** How do the effects of onshore wind turbines vary across regions and socioeconomic groups in the Netherlands?

The first question establishes a causal benchmark for onshore effects using modern data and an estimator suited to staggered adoption, building on existing literature. The second research question addresses the complete lack of Dutch offshore evidence in a descriptive rather than causal manner. The third examines the heterogeneity of onshore effects across spatial and socioeconomic contexts, showing under what conditions wind turbines cause the largest effects.

1.4. Research Approach

To answer research questions one and three, a Difference-in-Differences (DiD) specification will be adopted. DiD is an econometric method that estimates the causal effects of an event, for instance a turbine becoming operational, by comparing how the prices of properties with turbine exposure evolve relative to those not exposed. The key assumption is that prices would have *parallel trends* if no turbine was constructed, meaning that any price effect is attributable to wind turbines.

The analysis combines Dutch housing transaction data from the Nederlandse Vereniging van Makelaars (NVM) with turbine characteristics from WindStats and neighbourhood-level Socioeconomic Status (SES) data from the Central Bureau of Statistics (CBS). In the analyses, properties will be considered treated based on their distance to the nearest turbine. As turbine construction is staggered, the analyses employ a modern estimator suited for the staggered nature of turbine adoption. This ensures that the results are free from the potential bias that conventional specifications used in earlier studies may incur when properties are exposed to turbines at different times.

In contrast, the offshore analysis is primarily descriptive because the small number and recent development of offshore parks prevent reliable causal conclusions. The second research question is addressed descriptively rather than via a causal DiD estimate. The offshore analysis examines coastal price trends related to offshore wind farm development and notes data limitations that prevent reliable causal conclusions.

1.5. Relevance and Link to EPA

The findings of this study are relevant to municipal and provincial spatial planners and to the ministries responsible for climate and spatial policy. The estimated price effects for onshore turbines can inform siting decisions during the energy transition. For instance, if two projects are proposed simultaneously that produce similar levels of electricity, this research can help spatial planners choose the option that minimises the estimated depreciation of property. The results also bear on compensation policy. The Omgevingswet (the Dutch Environment and Planning Act) applies a societal risk threshold of four per cent of property value, below which depreciation does not qualify

for compensation ([Rijksoverheid, n.d.-a](#)).² By estimating where effects occur under various conditions, the analysis can identify who bears the costs of turbine externalities and help policymakers weigh whether to mitigate such impacts through siting strategies. This will thus make trade-offs more transparent, supporting more equitable public debates regarding new wind projects and their potentially localised adverse effects.

This focus aligns the thesis with the Engineering and Policy Analysis (EPA) programme at Delft University of Technology, an interdisciplinary programme that addresses grand societal challenges using quantitative and analytical methods to inform policy. This research project embodies this in three ways. Analytically, it employs recent econometric methods suited to staggered adoption, drawing on a large housing transaction dataset and combining programming and modelling to infer causality between turbine placements and property devaluations. In terms of policy, it highlights the estimated effects of alternative turbine siting strategies on property prices, which planners and ministries can use directly. Finally, it can inform public debate on turbine siting by showing how turbine-related losses are distributed across different locations and socioeconomic groups. This could shift the discussion from a qualitative debate about who bears the costs of the energy transition through wind expansion to quantitative evidence to support more transparent and equitable policymaking.

1.6. Reading Guide

The remainder of this thesis is structured as follows. Chapter 2 presents the empirical literature on wind turbine effects on housing prices and identifies the gaps which this thesis aims to address. Chapter 3 introduces the available data sources and illustrates how the datasets are merged for further analysis. Chapter 4 introduces the Difference-in-Differences (DiD) specifications, followed by the robustness and heterogeneity analyses. Chapter 5 reports the results for onshore and offshore turbines following the methodology. Chapter 6 synthesises the findings and discusses the policy implications that follow. Finally, chapter 7 presents the conclusion, which addresses the research questions, reflects on the research process, and outlines directions for future work.

²It should be noted that the threshold used to be two per cent but was increased to four per cent since January 1st, 2024, making compensation due to turbine exposure less likely.

Literature Review

This chapter reviews the empirical literature on how wind turbines affect housing prices. It begins by summarising hedonic pricing models and DiD studies on onshore turbines, then reviews the limited offshore evidence, and concludes by identifying knowledge gaps that will lead to the research questions. Having a clear grasp of this literature's strengths and limitations is crucial for generating evidence that is both methodologically credible but also informative for policy decisions on wind energy deployment within the Dutch energy transition.

2.1. Hedonic Pricing Evidence on Wind Turbines

To determine the effects of turbines on property prices, some studies used hedonic pricing models to estimate these effects. In essence, hedonic pricing models treat properties as bundles of attributes, with transaction prices reflecting structural characteristics, neighbourhood amenities, and environmental disamenities, including proximity to or exposure to wind turbines. In this framework, some form of turbine exposure is included in a price equation as one of many characteristics, and the estimated coefficient is interpreted as the marginal willingness to pay to avoid that exposure ([Skenteris et al., 2019](#)). However, because traditional hedonic models primarily use cross-sectional variation in prices, they often struggle to account for unobserved neighbourhood quality or endogenous siting, which is the tendency for wind turbines to be built in areas where housing prices are already lower.

Empirically, the international hedonic literature yields mixed findings. While some studies report significant price depreciation for properties near turbines, an equal number of large-sample studies find no systematic or statistically significant impacts after applying standard housing and spatial controls. Recent meta-regression evidence by [Schütt \(2024\)](#) demonstrates that this apparent inconsistency is largely driven by study design: after correcting for publication bias and methodological misspecifications, the true average price effect is modest, highly localised, and highly sensitive to whether studies control for pre-existing price differentials. The remainder of this section first synthesises the mixed international findings from primary studies. Subsequently, the meta-analysis by [Schütt \(2024\)](#) is used to explain this variance, ultimately highlighting the inherent limitations of standard hedonic pricing models and providing the methodological justification for a Difference-in-Differences (DiD) design.

2.1.1. Mixed International Findings

Empirically, the international hedonic literature on the effects of wind turbines on housing prices is highly fragmented, dividing largely into two camps. On one side, several studies document statistically significant, localised price reductions. For example, [Westlund & Wilhelmsson \(2021\)](#) find a significant negative effect for properties located up to 8 kilometres from wind farms in Sweden. Using nearly 100,000 transactions and over 4,000 turbines, they find that larger turbines or large clusters of turbines lead to particularly high property devaluation. Similar adverse effects are reported across diverse geographies, such as in Germany, where [Sunak & Madlener \(2016\)](#) apply a

spatially weighted hedonic model to show that price depressions are highly dependent on direct visual exposure. Heterogeneity is also evident on a regional scale: [Heintzelman & Tuttle \(2012\)](#) observe significant value reductions in two out of three examined counties in New York State. They propose several explanations, including differences in turbine age, siting characteristics, and heterogeneity in local preferences. Additionally, [Skenteris et al. \(2019\)](#) report a 12.4% price discount for dwellings within 2 kilometres of turbines on the Greek island of Evia, but no measurable effect on the neighbouring island of Kefalonia.

In stark contrast, a comparable strand of literature consistently fails to find any widespread or statistically significant price impacts. In the United States, a comprehensive multi-site study by [Hoen et al. \(2009\)](#) analysed nearly 7,500 sales across 24 wind facilities, concluding that neither the distance to the turbines nor their visibility had a measurable effect on property prices. This absence of a statistically significant property devaluation is also echoed internationally. In Canada, [Vyn & Mccullough \(2014\)](#) analysed over 5,414 rural residential and 1,590 farmland sales in Ontario and found no significant impacts from proximity or visibility. Similar statistically insignificant findings have been documented in New Zealand for properties located between 2.5 and 6 kilometres from turbines [McCarthy & Balli \(2014\)](#), and in the United Kingdom, where [Sims et al. \(2008\)](#) could not establish a relation between wind turbines and property prices.

However, it is important to note that some of these statistically insignificant findings may reflect limited statistical power, given the small sample sizes. For instance, the findings of [McCarthy & Balli \(2014\)](#) and [Sims et al. \(2008\)](#) are based on just 945 and 201 housing transactions, respectively. Such small, geographically confined samples often lack the statistical power to detect modest, highly localised effects. This could mean that, even though no such effects were found, they may exist but are not observable with the data used.

2.1.2. Limits of Hedonic Pricing

Rather than reflecting fundamental differences in how housing markets respond to wind turbines, recent research shows that the perceived inconsistency in the global literature is mainly due to variations in study design and methodological limitations. These studies differ considerably in their definition of turbine exposure, ranging from distance-based measures to various viewshed analyses. Additionally, standard cross-sectional hedonic pricing models often struggle to distinguish turbine effects from pre-existing neighbourhood characteristics and unobserved spatial dynamics. These methodological choices, along with the limited statistical power of some studies focused on specific regions, likely explain the mixed results seen in the international literature.

A synthesis of the fragmented literature can be derived from a comprehensive meta-regression analysis of 720 estimates from 25 hedonic pricing studies, conducted by [Schütt \(2024\)](#). He demonstrates that the variance in reported effect sizes can be attributed to publication selection bias and methodological misspecifications. Specifically, studies that fail to control for unobserved spatial effects and pre-existing price differentials systematically overstate the negative impacts of wind turbines on property prices. After correcting for publication bias and misspecifications, the meta-analysis highlights that the true average effect is modest: property values decline by 0.68% at a distance of 1.89 miles, whereas this effect turns to zero beyond 2.8 miles (or 4.5 kilometres).

To overcome the limitations of regular hedonic pricing models, specifically omitted variable bias and endogenous siting (the tendency for turbines to be built on already cheaper land), some researchers note that DiD specifications could help reduce these issues. For example, [Sunak & Madlener \(2017\)](#) note that utilising DiD designs with repeat sales could further reduce concerns about omitted variable bias and thus increase confidence in the results found. Other hedonic pricing model studies attempt to address these issues within the hedonic framework itself: [Heintzel-](#)

man & Tuttle (2012) apply both location and time fixed effects to control for unobserved spatial heterogeneity and temporal price trends. Meanwhile, studies such as Hoen et al. (2009) incorporate specifications with repeat-sales models to ensure that all time-invariant structural and location characteristics of a property are eliminated.

As Schütt (2024) highlights, model specifications that omit these types of controls systematically overstate the price depreciation of properties exposed to turbines. Ultimately, this means that to obtain credible, causal estimates free from endogeneity and omitted-variable bias, identification strategies must move beyond standard hedonic pricing models and adopt designs that explicitly account for spatial unobservables.

2.2. Difference-in-Differences: Onshore Evidence

To address endogeneity and omitted-variable bias inherent in cross-sectional hedonic models, recent literature has increasingly turned to Difference-in-Differences (DiD) designs. By comparing the price trends of properties exposed to wind turbines with those of unexposed properties before and after construction, these approaches yield more credible causal estimates of the effects of turbine exposure. The causal effect, in addition to fixed effects, relies on the *parallel trends assumption*, under which, in the absence of turbine exposure, treated and control properties would have followed identical price trends. Under this assumption, any post-treatment differences between the two groups are attributable to the turbine exposure, yielding a causal interpretation of the treatment effect.

This section synthesises the empirical evidence with a particular focus on the Dutch context. The remainder of this section is structured as follows. First, it outlines the baseline causal effects of onshore wind turbines on local property values in the Netherlands. Second, it explores the underlying drivers of these effects, such as proximity and visibility, and then examines temporal dynamics, including anticipation and post-construction adaptation. Next, it briefly discusses the methodological implications of staggered adoption and the necessity of modern estimators. Finally, the section delves into spatial and socioeconomic heterogeneity to understand how these price effects are distributed across different communities, leading to a subsequent discussion of the much scarcer evidence on offshore wind.

2.2.1. Baseline Dutch Effects

Dutch DiD evidence on onshore wind turbines is provided by three closely related studies. Dröes & Koster (2016) provided the first causal estimate in 2016, estimating that properties within two kilometres of an onshore turbine decline in price by approximately 1.4%. Their follow-up study, Dröes & Koster (2021), extends the sample period and introduces turbine-height categories, revealing that the effects are substantially larger for taller turbines. Properties near turbines exceeding 150 metres in tip height decline by up to 5.4%, those near medium-height turbines (50-150 metres) by approximately 3.0%, and those near smaller turbines by up to 2% within one kilometre. Eichholtz et al. (2023) find similar adverse effects, reporting effects between 1.8 and 2.1%, which increase to 2.7% when the sample is restricted to repeat sales of the same property.

These estimates are considered credible for several reasons. All three studies exploit fine-grained location fixed effects at the 4-digit postal code (PC4) and 6-digit postal code (PC6) levels, controlling for time-invariant neighbourhood characteristics at a very local scale. The sample sizes are very large, ranging from approximately 1.2 to 3 million housing transactions, reducing the risk of bias from outliers in the causal estimates. Dröes & Koster (2021) strengthen causal claims by using future-treated areas, locations where turbines have been planned but not yet constructed, as the control group. This reduces concerns that treated and control areas differ systematically, such that only the trends of comparable groups are compared. Additional evidence for causality comes

from [Eichholtz et al. \(2023\)](#), who show that when turbines are decommissioned and thus no longer treated, nearby property prices recover by approximately 6.7%. The reverse effect, whereby property prices increase after turbine decommissioning, highlights the adverse impact of turbines on property prices in the Netherlands.

Despite their methodological strengths, it should be noted that all three papers rely on the same datasets. Transaction data is provided by the NVM Brainbay, a dataset of the Nederlandse Vereniging van Makelaars (NVM). These housing transactions were then linked to turbines from the Wind-Stats database. These shared datasets imply that the estimates are directly comparable across studies, but they also imply that any limitations of the NVM data, most notably its approximately 70% coverage of Dutch transactions and slight overrepresentation of the Randstad, carry through to all findings of turbine effects. Moreover, all three studies focus exclusively on onshore wind turbines, fully omitting the effects of offshore turbines on coastal housing markets.

2.2.2. What Drives the Effect?

While it is clear that onshore turbines depreciate property values in the Netherlands, these studies did not address the drivers of these effects. Housing prices capitalise a range of perceived externalities into a single transaction value, and the literature identifies several channels through which wind turbines may affect local housing markets. These include visual obstruction, noise, shadow flicker, and broader concerns about landscape industrialisation ([Mattmann et al., 2016](#); [Onakpoya et al., 2015](#)). Qualitative work on community energy projects shows that renewable energy developments can both mobilise and disrupt place attachment, with opposition to wind turbines often stemming from emotional ties to unspoiled landscapes ([van Veelen & Haggett, 2017](#)). Because DiD models estimate reduced-form price effects, they cannot determine which of these elements generate the observed property depreciation. Nevertheless, several studies have attempted to disentangle proximity effects from visibility effects, leading to mixed evidence.

[Dröes & Koster \(2016\)](#) test whether turbine visibility and shadow casting contribute to price effects. The interaction between the treatment variable and turbine view is positive but not statistically significant, and shadow casting is also insignificant. While this analysis is noisy, it suggests that, conditional on proximity, having a direct view of a turbine does not generate additional price depreciation. In contrast, proximity alone consistently produces significant negative estimates. Another possible explanation is that the binary view measure used by [Dröes & Koster \(2016\)](#) is too coarse to capture meaningful variation in visual exposure. However, evidence from the United States complicates this hypothesis. [Lang et al. \(2014\)](#) conducted 1354 site visits to physically verify whether turbines were visible from properties and categorised the degree of dominance the turbine has on the landscape in 5 categories. Despite this effort, they find no statistically significant relationship between turbine visibility and property prices. This suggests that visibility may not fully capture the causal effects of turbines, given that these site visits are considered the gold standard for assessing actual visibility.

More recently, [Dong et al. \(2024\)](#) revisited visibility exposure using digital surface model (DSM) data derived from LiDAR, accounting for buildings, trees, and other elements that obstruct direct line of sight to turbines. They find that conventional digital elevation models, which ignore these obstructions, misclassify the visibility status for about 75% of properties compared to DSM. When visibility is measured using the more precise surface model, statistically significant adverse price effects of approximately 2.2% to 2.5% were found. Although this suggests that measurement quality is highly important, the contrast with [Dong et al. \(2024\)](#) is noteworthy. Physical site visits, which should be highly accurate, such as the digital surface model, did not detect effects that differed statistically from zero. This discrepancy suggests that whether visibility drives price effects may depend on other factors, like the geographic or housing-market context, rather than on measurement

precision alone. Additionally, other turbine characteristics may also shape the severity of exposure. (Dröes & Koster, 2016) find that turbines with large blade diameters produce larger price effects than those with tall masts alone, but both effects are statistically significant. This could imply that the tip height will best capture exposure, combining both blade diameter and mast height.

Taken together, the literature suggests that proximity to turbines is the most consistent predictor of price effects, whereas visibility effects require high-quality measurement and may differ by context. This has implications for how to best capture turbine exposure in empirical studies. Since proximity likely captures a set of externalities, including but not limited to visual obstruction, and since turbine height amplifies both visual obstruction and other proximity-related effects, such as noise, it is important to conduct a wide range of heterogeneity analyses to test the effects of various turbine attributes, instead of relying solely on distance-based treatment definitions.

2.2.3. Temporal Dynamics

The effects of wind turbines on housing prices are not limited to periods after construction. Several studies document anticipation effects, where prices begin to decline before turbines become operational. Dröes & Koster (2016) find that property prices already declined two years in advance at a statistically significant level by 1.7%, which highlights that prices move before turbine construction. Internationally, similar findings have been reported. Brunner et al. (2024) report that price effects in the United States begin after the announcement rather than during construction, with 11% declines. Furthermore, Hoen et al. (2025) synthesises that post-announcement effects in the US are a decline of 5% to 10% in property prices. These papers thus imply that anticipated turbines already affect property prices before turbines begin operating.

There also exists evidence on whether price effects persist, recover, or intensify after construction. Dong et al. (2023) document a U-shaped pattern in New England, which they attribute to residents getting used to turbines, following the argument of Devine-Wright (2005), who found that local attitudes toward wind energy can evolve once people get more familiar with turbines. Brunner et al. (2024) find a comparable rebound after an initial drop in prices, and Eichholtz et al. (2023) estimate that the treatment effect diminishes by approximately 0.4% per year. On Cape Cod, Dong et al. (2023) find that while prices do rebound slightly, they remain at -8% after nine years post-construction. However, not all evidence points towards a rebound post-construction. Dröes & Koster (2021) find that willingness-to-pay to live near turbines has remained constant between 1985 and 2019, suggesting this rebound effect does not exist in Dutch contexts.

2.2.4. Modern Estimators

The preceding sections documented that turbine effects vary in magnitude over time, with anticipation effects emerging before construction and partial recovery occurring in some contexts. Another element that varies over time is the staggered adoption of turbines, resulting in variation in treatment timing. Goodman-Bacon (2021) shows that under staggered adoption, the conventional Two-Way Fixed Effects (TWFE) estimator can produce biased treatment effect estimates because the coefficient of interest is now a weighted average of multiple DiD estimators, potentially leading to biases and less certainty regarding causal claims.

Modern DiD estimators have been developed to address this issue. Callaway & Sant'Anna (2021) propose an estimator that separately identifies cohort-specific treatment effects before aggregating them, and Sun & Abraham (2021) introduce interactions between cohort indicators and relative time periods in a single specification to obtain clean cohort-by-event-time treatment effects. These effects are then aggregated using weights based on cohort shares at each event time to construct an average of the estimand of interest. Despite the availability of these methods, only two of the onshore studies reviewed employ them: both Brunner et al. (2024) and Hoen et al. (2025) use a

stacked design combined with the Sun and Abraham estimator for event studies. All three Dutch studies rely exclusively on conventional TWFE models. While the Dutch findings are supported by extensive robustness checks, the exact estimated effects may be imprecise due to the issues identified by [Goodman-Bacon \(2021\)](#). This means that applying a modern estimator to Dutch contexts could verify whether prior findings of estimated effects still hold.

2.2.5. Spatial Heterogeneity

While it is now clear what the estimated effects are in the Dutch context, a large body of literature also documents how these effects vary across space. Aggregate estimates can mask this variation, and understanding where the effects are concentrated is essential for interpreting the policy relevance of average treatment effects.

Credible findings are presented by [Dröes & Koster \(2016\)](#), who interact their treatment indicator with province dummies and find statistically significant adverse effects in 7 out of 12 provinces, meaning that in many provinces, the adverse effects are higher than the documented average of 1.4%. They further document an urban penalty, in which properties in urban areas experience an additional price decline of approximately 3.2% relative to those in non-urban areas. [Eichholtz et al. \(2023\)](#) find similar effects, highlighting that additional adverse effects are present in urban areas across all types of energy facilities they examine. Taken together, this means that turbine effects vary spatially across the Netherlands and are more negative in urbanised areas than in rural areas.

International evidence further supports this pattern. [Brunner et al. \(2024\)](#) analyse onshore wind projects in the United States and find that the entire national result is driven by urban counties, which exhibit adverse effects of approximately 14.9%, while rural counties show adverse effects of 1.6%, which are not statistically significant. In New England, [Dong et al. \(2023\)](#) find that negative turbine effects are mainly concentrated in Cape Cod and Nantucket, where estimated price declines within a 1 km range from about 7% to 10.8% and remain around 9%, up to nine years post-construction. Meanwhile, estimates for other regions in Massachusetts and Rhode Island are small and mostly statistically insignificant, suggesting substantial spatial heterogeneity in treatment effects.

While geographic groupings and urban-rural divisions suggest that effects vary, understanding these spatial patterns requires examining local physical impacts, such as visibility, noise, and shadow flicker, as well as unobserved neighbourhood dynamics, often referred to as spatial dependence. To address the similarity in unobserved price changes among nearby properties, some researchers employ hedonic pricing models that incorporate DiD estimators and spatial econometric methods. For instance, [Sunak & Madlener \(2016\)](#) find that spatial variation in property devaluation is heavily driven by visual exposure. Applying such specifications in Germany, they find price drops of 9% to 14% for properties whose view was strongly affected, whereas minor views show no such effects.

The local social and political context also appears to shape treatment effects. [Dong et al. \(2023\)](#) argue that the persistent effects on Cape Cod and Nantucket reflect the legacy of the failed Cape Wind project, which could have spilt over into residents' dislike of offshore turbines. If true, that would mean the effects of turbines on property prices do not stem solely from noticeable characteristics but are also influenced by local attitudes toward wind energy. These spatial patterns are well-documented but rely on geographic classifications, such as urban-rural splits, individual provinces, or regional groupings. Even when local attitudes are incorporated, as in [Vyn \(2018\)](#), it remains largely unclear which underlying socio-economic characteristics drive these attitudes toward turbines, leading to the observed variation in treatment effects.

2.2.6. Socioeconomic Heterogeneity

The literature so far has shown that treatment effects vary substantially across space, with urbanisation and local context emerging as strong moderators. This raises the question of which underlying socioeconomic characteristics, if any, lead to the observed variation. Despite the policy relevance of this question, the empirical literature offers very little evidence.

Only a small number of studies have incorporated socioeconomic variables into their analysis of wind turbine effects on housing prices, and they do so in varying ways. While this helps isolate the effect of turbines by accounting for socioeconomic confounders, it does not reveal whether socioeconomic factors themselves moderate the impact of turbines on housing prices. For instance, [Hoen et al. \(2025\)](#) include school-district demographic and socioeconomic elements as part of housing characteristics (e.g., median income, share of college-educated people, age structure, racial composition, rural status, etc.) that are interacted with linear time trends. However, as discussed before, this does not indicate whether these socioeconomic factors moderate the effects of turbines on housing prices, thereby creating heterogeneity in treatment effects.

To date, [Dröes & Koster \(2016\)](#) appears to be the only study to interact some socioeconomic variables with the treatment indicator. They test three interactions individually: (i) an urban dummy, (ii) the share of high-income households, and (iii) a residential area dummy. The income result is positive but modest and only weakly significant at the 10% level, moderately suggesting that high-income areas experience less negative turbine effects. The results for urban and residential heterogeneity were both negative and statistically significant, highlighting that urban and residential areas suffer greater effects. Importantly, these interactions were tested one at a time rather than simultaneously, so it remains unclear what the effects are when moderating for other socioeconomic factors. [Dröes & Koster \(2021\)](#) take a different approach in their appendix, comparing whether treated properties have different socioeconomic backgrounds on average compared to non-treated properties. They find that turbines tend to be sited in areas with a lower population density of 30% and a lower median income by 2%. Additionally, areas with a higher share of foreigners also exhibit larger adverse effects, but these may correlate with median income. These statistics indicate that areas where turbines are placed have different socioeconomic backgrounds, but they do not address whether socioeconomic factors moderate the treatment effect itself. [Eichholtz et al. \(2023\)](#) include only an urban location-control dummy and an urban interaction with the treatment indicator, finding negative price effects in urban areas for all facility types studied, including wind turbines. Finally, after testing their main specification, [Brunner et al. \(2024\)](#) split their sample into urban and rural subgroups, but did not explore treatment-effect heterogeneity with other socio-economic elements.

No study has examined whether Socioeconomic Status (SES) variables such as labour history, financial components, or educational levels jointly moderate the treatment effect on housing prices. The existing findings are three individually tested interactions from a single study. Consequently, it remains largely unclear for whom turbine effects are largest, limiting the ability to assess how the costs of wind turbine expansion are distributed across the Netherlands.

2.3. Difference-in-Differences: Offshore Evidence

While there is a substantial body of literature on the effects of onshore turbines on housing prices, empirical evidence on offshore turbines is extremely limited. Only two studies have estimated the effect of offshore wind turbines on coastal housing prices, and both find no statistically significant effect. However, these studies differ substantially in context, scale, and distance from the shore, which could limit the extent to which these findings could apply to Dutch coastal settings.

[Jensen et al. \(2018\)](#) analyse the Nysted and Rødsand II offshore wind farms on the southern coast

of the Danish island of Lolland. The two farms, completed in 2003 and 2010, respectively, comprise 72 and 90 turbines, with hub heights of approximately 69 and 80 metres, and are located only several kilometres from each other, at distances of 9.5 and 3.5 kilometres from the coast. They employ a Difference-in-Differences (DiD) design that utilises variation in time and treatment, and they compare prices of properties with and without views of the offshore farms, both from the property itself and/or from the nearest beach. Across all model specifications, for both residential and vacation properties, they find no statistically significant effect of wind farm visibility on property prices.

[Dong & Lang \(2022\)](#) study the Block Island Wind Farm in Rhode Island, which is the first offshore wind farm in the United States, completed in August 2016. The farm consists of only five turbines producing 30 MW, located approximately 4.8 kilometres from Block Island, and 27 to 44 kilometres from the Rhode Island mainland. Using a hedonic DiD design with LiDAR-based viewshed exposure, they find no significant effects on mainland coastal property prices across a wide range of specifications. A robustness analysis of properties on Block Island itself, where turbines are considerably more visually prominent due to their closer proximity, also yields insignificant effects, although both turbine view estimates are negative. While the use of high-quality viewshed exposure strengthens the identification strategy, the relatively small scale of the offshore wind farm may limit the generalisability of these statistically insignificant findings to larger, closer-to-shore installations.

Several factors suggest that these international findings may not directly transfer to the Netherlands.

1. First, planned Dutch offshore zones, such as Hollandse Kust, are located at varying distances from the coastline. The distances from relevant offshore wind sites to the coastline are 18.5 and 24 kilometres for Borssele and Hollandse Kust (North and South), respectively ([Rijksoverheid, n.d.-b](#)). This would mean that only the study by [Dong & Lang \(2022\)](#) would be representative with respect to distance. However, it could be hypothesised that insignificant effects at closer ranges, found by [Jensen et al. \(2018\)](#), also highlight that effects may not exist when turbines become less visible.
2. Second, the scale of Dutch offshore wind farm size differs substantially, at least compared to the study regarding Rhode Island (USA) by [Dong & Lang \(2022\)](#), where the effects of only five offshore turbines were studied. The study in Denmark by [Jensen et al. \(2018\)](#) is slightly more representative in terms of park size, with 72 and 90 offshore wind turbines at the Nysted and Rødsand II wind parks, respectively. However, this count is still quite low for the Dutch offshore wind parks, given that Borssele, Hollandse Kust North, and Hollandse Kust Zuid comprise 174, 68, and 138 offshore turbines, respectively.

In addition, the US and Danish turbines are smaller, reducing visibility and thus visual obstruction. [Jensen et al. \(2018\)](#) documents hub heights of 72 and 80 metres for the Nysted and Rødsand II windparks, whereas [Dong & Lang \(2022\)](#) documents a hub height of 100 metres and a blade length of 75 metres. However, this pales in comparison to Dutch offshore turbines near the coast, where tip heights of 191 to 251 metres were documented by WindStats.

3. Third, existing studies on offshore wind primarily examine areas in the United States and Denmark, which may not reflect Dutch coastal housing markets. For instance, [Dong & Lang \(2022\)](#) analyse the Block Island Wind Farm using property transactions in small coastal communities in Washington County, Rhode Island, within three kilometres of the shoreline. These properties are generally quite large, with an average living space of 334 square metres and a lot size of 2,521 square metres, which is not a good proxy for Dutch coastal

housing markets. The Danish coastal market consists of smaller houses more comparable to Dutch coastal houses, but they are still quite large, averaging 140 square metres (Jensen et al., 2018). However, they only analyse single-family residential and vacation homes, which, in the Netherlands, would omit terraced housing. For comparison, a Dutch recreational property averaged 70 square metres of living space (NVM, 2024).

4. Fourth, the preceding research is limited to viewshed exposure measures, omitting the possibility that offshore turbines could impact coastal communities, even without direct view of them from the properties or the nearest beach. This aligns with findings by Russell et al. (2021) and Vyn (2018), who stated that local attitudes also drive treatment effects. It could thus be hypothesised that offshore turbines affect the entire coastal community near the turbines, as attitudes are mostly community-wide rather than limited to properties with a direct view, which may explain the reported statistically insignificant findings, as both properties with and without a direct view are affected simultaneously. Focusing solely on viewshed exposure precludes testing whether this hypothesis holds; although statistically insignificant effects were reported, they may not reflect actual effects on housing markets.
5. Lastly, Russell et al. (2021) show that offshore wind development affects coastal communities through multiple socio-cultural dimensions, including concerns about tourism, visual amenity, place attachment, and local identity. They note that evidence on these aspects in the US context still largely comes from case studies, perception and stated-preference research, and has not yet been systematically linked to housing market impacts.

If the initial price effects of onshore turbines partly reflect an uncertainty premium, then offshore turbines, which are distant, undergo long and well-known planning processes, and involve minimal direct noise or shadow casting, as mentioned in subsection 2.2.2, may never generate sufficient uncertainty to depress coastal property prices. After all, next to the place attachment elements, only visual obstruction remains, which is quite predictable. This interpretation is consistent with the statistically insignificant findings in the existing literature but remains untested in the Dutch context.

To date, no study has estimated the effect of offshore wind turbines on housing prices in the Netherlands, neither causal nor descriptive, despite the rapid expansion of offshore capacity documented in the introduction. This gap motivates the descriptive offshore analysis conducted in this thesis.

2.4. Synthesis and Research Gaps

The reviewed literature highlights several key findings. Dutch Difference-in-Differences (DiD) studies consistently estimate significant adverse effects of onshore wind turbines on nearby property prices, ranging from approximately 1.4% to 2.1% for baseline estimates, increasing to 5.4% for taller turbines, and a penalty of approximately 3.2 percentage points in urban areas. These effects are mostly driven by proximity rather than by any single mechanism, emerge before construction through anticipation, and may partially recover over time in some contexts. Furthermore, treatment effects vary across space, with urbanisation and local attitudes moderating them, yet the socioeconomic characteristics driving the observed variation remain largely unexplored. Offshore evidence is limited to only two international studies, both of which find statistically insignificant effects in contexts that differ from the Netherlands in scale, turbine size, and housing-market structure. Methodologically, all Dutch studies utilise conventional Two-Way Fixed Effects (TWFE) designs, which may produce imprecise estimates due to the staggered adoption of Dutch turbine construction. Despite these findings, two gaps limit the current understanding of the effects of turbines on housing markets.

1. No study has examined the relationship between offshore wind development and coastal housing prices in the Netherlands, whether causal or descriptive, even though offshore capacity has expanded rapidly in recent years through large-scale wind parks such as Borssele, Hollandse Kust, and IJmuiden Ver. The two existing international studies find no statistically significant effects but generally consider only smaller farm scales, shorter turbines, and housing markets that differ from Dutch coastal settings. As documented in section 2.3, Dutch offshore farms have more turbines and greater tip heights than those studied in Denmark and the United States. So, while prior research resulted in non-significant effects, these studies are not directly comparable to Dutch housing markets. This means it remains unknown whether offshore wind development is associated with observable coastal housing price patterns in the Dutch context.
2. The literature only provides a limited understanding of what drives heterogeneity in treatment effects. Spatial variation is well-documented through urban-rural splits, provincial interactions, and regional groupings, but these are descriptive geographic categories rather than explanatory variables. The socioeconomic dimension is almost entirely unexplored. Only one study has interacted socioeconomic variables with the treatment indicator, testing three elements individually rather than simultaneously. No study has examined whether Socioeconomic Status (SES) variables and other factors moderate treatment effects when tested simultaneously. This limits the ability to understand why, and for whom, the costs associated with turbine adoption are highest. This question is directly relevant to the design of equitable siting policies and to mapping whether the Omgevingswet threshold of 4% is exceeded, thereby allowing homeowners to be financially reimbursed for the depreciation of their properties.

2.5. Research Questions and Hypotheses

The previous synthesis highlighted two gaps in the current research: the lack of evidence on the impact of offshore turbines on Dutch housing markets, and the limited understanding of how socioeconomic factors drive heterogeneity in treatment effects. Additionally, there is a need to establish a modern causal benchmark for onshore effects. These gaps lead to three research questions, as outlined in section 1.3, that aim to address the thesis's main question. This section details how each question arises from existing literature and presents the related hypotheses.

The first question concerns the onshore baseline. Although prior Dutch research has estimated onshore effects, this analysis revisits them using more recent transaction and turbine data, and will include a modern staggered-adoption estimator, establishing a causal benchmark that also serves as a reference for interpreting the offshore and heterogeneity analyses. Based on the consistent findings of Dröes & Koster (2016), Dröes & Koster (2021), and Eichholtz et al. (2023), it is hypothesised that onshore turbines generate statistically significant adverse effects on nearby property prices, with prior estimates ranging from -1.4% to -5.4% depending on turbine height and distance. It is further hypothesised that adverse turbine effects are greater as turbines have become larger in recent years (see Figure 3.5), meaning negative impacts such as shadow casting and visual obstruction are more prominent, potentially increasing the range over which adverse effects are observed.

The second question addresses the first gap, the complete lack of Dutch offshore evidence, and aims to provide a first descriptive analysis of how coastal price patterns offshore wind farm development. Since studies from Denmark and the US often find offshore effects to be minimal or statistically insignificant (Dong & Lang, 2022; Jensen et al., 2018), any impact in the Dutch context is probably small. However, the Dutch scenario differs with larger farms, taller turbines, and a unique coastal housing market.

The third question addresses the second gap by exploring the socioeconomic factors behind heterogeneity in the Netherlands. Building on the urban-rural divisions identified by [Dröes & Koster \(2016\)](#) and [Eichholtz et al. \(2023\)](#), the analysis extends to examine whether socioeconomic factors moderate the treatment effect. This helps clarify what causes the observed heterogeneity and whether the urban penalty is actually related to socioeconomic composition. Based on earlier Dutch research, negative impacts are expected to be greater in urban areas. Because Socioeconomic Status (SES) may influence political power and resistance to siting decisions, it is also hypothesised that negative price effects are stronger in regions with lower SES in finance, labour history, and education. If true, some of the urban penalty might actually stem from socioeconomic differences rather than urbanisation itself.

Data Sources and Analysis

This chapter presents the data used in the analysis. It describes the four data sources: housing transactions, the turbine data, socioeconomic data, and geospatial layers, and explains how they are merged. Next, it examines the resulting dataset, analysing the identifying variation required for causal inference and the socioeconomic composition of turbine-exposed areas. This is important because it will determine whether the data can detect causal effects. Furthermore, the identified socioeconomic patterns will justify further analyses of heterogeneity.

3.1. Data Sources

3.1.1. Housing Transactions

The primary data source for this analysis, shown in chapter 5, is the housing transaction dataset from Brainbay, by the Nederlandse Vereniging van Makelaars (NVM), the Dutch Association of Realtors. The dataset covers property transactions in the Netherlands from 2010 to 2023 and includes transaction prices, transaction dates, detailed property characteristics, and precise location characteristics, such as 6-digit postal code (PC6) data and Rijksdriehoek coordinates. This level of detail will be highly important as the transaction-level data with exact coordinates permit calculation of the distance from each property to nearby operational turbines. Furthermore, the detailed locational data also allows for location fixed effects, which will be introduced in chapter 4. While broader sources like CBS/Kadaster offer broader coverage, they lack variables such as coordinates or housing characteristics (e.g., PC6, living space), making them less suitable for designs aiming to infer causality.

The NVM data cover approximately 70% of all housing transactions in the Netherlands. Contact with a Brainbay representative indicated that this coverage is somewhat higher in the Randstad, the north, and the centre of the country, while it is somewhat lower in the provinces of Zeeland and Limburg, with no apparent bias by property type or transaction price. The dataset generally reflects the national housing market, although coverage in Zeeland is limited, which is relevant for offshore analysis, where the coastal sample size is already small. Figure 3.1 illustrates the geographic distribution of transactions across the Netherlands, confirming that housing transactions are spread across all regions rather than only concentrating in specific areas. It further highlights that some highly urbanised areas have a considerable number of transactions, which is to be expected of cities like Amsterdam or Rotterdam.

Figure 3.2 showcases the number of transactions over the years. It is clear that the annual number of property transactions varies over the sample period, mirroring broader trends in the Dutch housing market. This includes relatively few transactions between 2010 and 2013, after which transaction volumes increase. These fluctuations should be kept in mind when interpreting the results presented in chapter 5, as fewer transactions near the beginning and end of the sample may affect the analyses.

The NVM dataset also includes a rich set of housing characteristics that will serve as control vari-

National coverage of NVM housing transactions




Figure 3.1: The national spatial coverage of all cleaned NVM transactions between 2010 and 2023.




Figure 3.2: The annual transactions recorded by the NVM between 2010 and 2023.

ables, including the living area, number of rooms, construction period, property type, maintenance condition, and whether the property has a monument status. These control variables are presented in Table E.1, while Table 4.1 compares these characteristics between treated and non-treated properties within the NVM dataset.

Lastly, it is important to note that the NVM data is a repeated cross-section rather than structured as a panel. This means each row represents a separate housing transaction, with columns containing information on its price, location, sale date, and other characteristics. Houses are thus not followed over time, unlike in panel data, but the same house can occur multiple times in the dataset. This also means that many observations are required for each treated property sale, because most property sales involve non-treated properties.

3.1.2. Wind Turbines

Wind turbine data are obtained from WindStats, which records all onshore and offshore wind turbines in the Netherlands. For each turbine, they report its location in Rijksdriehoek coordinates (among other coordinate systems), its tip height, power capacity, and the dates of commissioning and decommissioning. This supports the construction of distance-based treatment measures and will allow heterogeneity analyses by turbine characteristics. WindStats is thus preferred over alternative turbine datasets for its completeness, providing full national coverage of both onshore and offshore turbines.¹ Other sources are often less comprehensive, with imprecise offshore locations or a lack of national coverage.

The most important turbine characteristic documented by WindStats, next to their location, is the dates on which turbines become operational, as many identification strategies rely on timing to determine whether a property is exposed to turbines. These commissioning dates for all turbines are summarised in Figure 3.3.

Figure 3.4 shows the location of all turbines ever constructed in the Netherlands by type, categorised by turbine type into onshore, offshore at sea, and offshore inland. The map illustrates that

¹Contact with a representative of WindStats in December 2025 has confirmed that, to their knowledge, 100% of all Dutch turbines are included in the turbine dataset.




Figure 3.3: The number of turbines constructed each year by turbine type during the sample period of the housing transactions data.

wind parks, both onshore and offshore, tend to cluster rather than be evenly distributed across the Netherlands. This needs to be taken into account when designing specifications and interpreting the results.

Figure 3.5 shows the development in tip heights of newly constructed turbines. Across all turbine types, it reveals a clear upward trend in the sample period, with more recently constructed turbines being substantially taller on average. This trend is relevant for two reasons. First, it motivates analysing heterogeneity by turbine height, as [Dröes & Koster \(2021\)](#) already found that larger turbines cause greater adverse effects, implying that the magnitude of these effects is expected to increase. Secondly, it could indicate that the Average Treatment Effect of the Treated (ATT) may shift per cohort (the year of initial treatment), since later-treated cohorts are systematically exposed to taller turbines.

3.1.3. Socioeconomic Elements

Socioeconomic data are obtained from [Centraal Bureau voor de Statistiek \(2024\)](#), which publishes Socioeconomic Status (SES) scores at the neighbourhood level. The composite score summarises the overall socioeconomic position of a neighbourhood and is constructed from three elements: financial wealth, educational levels, and labour-market history.² Each socioeconomic dimension is also provided as an individual sub-score, enabling the analysis to identify which factors influence the relationship between turbines and housing prices. Note that the average SES scores are defined to be zero in the Netherlands, so a higher score could indicate more financial welfare, a higher educational level, or more stable employment.

The SES data are available from 2014 onward and are recorded at the neighbourhood level. Two features of the data are relevant for the analysis. First, the Central Bureau of Statistics (CBS) only records SES values for neighbourhoods with a sufficient number of households, often excluding neighbourhoods with fewer than 100 households, so very small neighbourhoods are not covered. This means the SES analyses are restricted both temporally, the 2014-2023 period, and spatially, to

²The labour-history component of the SES score summarises the recent employment history of households within a neighbourhood over the past years ([Centraal Bureau voor de Statistiek, n.d.](#)). It categorises households on a scale from consistently employed, through various intermediate stages of partial employment based on duration, to continuously unemployed or retired. A higher neighbourhood score indicates more sustained participation in the labour market.




Figure 3.4: The locations of all turbines ever constructed in the Netherlands categorised by turbine type.




Figure 3.5: The tip height of newly constructed wind turbines over the years, organised per type.

areas with available SES data. Second, SES scores are not constant over time. Figure 3.6 shows SES scores over time, highlighting that socioeconomic conditions evolve rather than being a time-invariant neighbourhood feature.




Figure 3.6: The temporal variation in the transaction-weighted SES scores from 2014 onwards.

A more detailed analysis of the dataset’s socioeconomic composition, including spatial variations of SES, its connection to transaction prices, and differences between turbine-exposed and unexposed regions, is provided in subsection 3.3.3.

3.1.4. Geometries

A spatial map is provided by publicly available Dutch geospatial layers. More importantly, two sources are used to define water-related elements. First, sea areas and coastline geometries are taken from the [Kadaster \(2025\)](#), and additional inland surface waters are provided by the [Informatiehuis Water \(2025\)](#). Using these geometries serves multiple purposes during the analyses. First, they allow determining whether a turbine is onshore, offshore at sea, or offshore inland (located in water bodies such as lakes). Secondly, they allow determining whether properties belong to coastal communities based on their distance to the coastline, which is used in the offshore analysis. Both of these are also helpful when plotting the turbine placements and properties exposed to turbines.

Figure 3.7 shows the coastline and water layers used to compute coastal distance and to classify whether turbines are placed onshore or offshore. More details about how the data merging process went will be given in section 3.2.

3.2. Construction of Analysis Dataset

The final analysis dataset is constructed in 4 steps by (i) linking cleaned Nederlandse Vereniging van Makelaars (NVM) transaction records to turbine data, (ii) defining coastal and geographic variables, (iii) adding socioeconomic variables to each housing transaction, and (iv) removing observations that do not align with the methodology presented in chapter 4.

First, the raw housing transactions are filtered to keep only regular residential market deals with realistic property features. The cleaning process excludes non-market transactions and auctions,

Water boundaries used for coastal classification



Figure 3.7: The coastline and water layers used to compute coastal distance and classify water-related geography.

entries missing important data (such as building period, urbanisation grade, and living area), extreme price-per-square-metre outliers by excluding the top and bottom 1%, and some physically implausible records (for instance, cases where the living area divided by the number of rooms is less than five). The sample size was initially 1,922,124 and dropped to 1,859,714 property sales from 2010 to 2023.

Second, wind turbines from the WindStats dataset are classified by turbine type, indicating whether they are onshore, offshore at sea, or offshore inland, using the sea and inland-water polygons described in subsection 3.1.4, with manual corrections for known edge cases.³ For each transaction, distances to turbines are calculated in the Dutch RD New coordinate system using Euclidean distance. A transaction is considered to be currently treated, such that `currently_treated = 1`, when at least one turbine of the relevant type within a specified treatment radius is operational on the transaction date. The pipeline also records ever-treated status, treatment timing, event time (measured in years relative to treatment for a specific property), and the number of operational turbines. For heterogeneity analyses later on, each property also has cohort-based characteristics, defined as turbine characteristics during its initial exposure year.

Third, Socioeconomic Status (SES) variables are added to each observation; these are available at the neighbourhood level from 2014 to 2023 and recorded annually. Because neighbourhood definitions change over time and thus cannot be merged with the housing transaction data, the cross-sectional SES variables are aggregated into 4-digit postal code (PC4) regions for each year using household weights, thereby creating a panel structure. Every transaction is then matched to its SES data by merging them at the PC4 region level in the corresponding sale year. Consequently, SES controls are available only for the 2014-2023 subsample.

The final R preparation step changes the data types accordingly for the analysis and filters out additional observations. It excludes transactions in which a property was previously exposed to a turbine and then had that exposure removed due to turbine decommissioning within the treatment range. This process removes an additional 30,931 observations to ensure that the final sample aligns with methodological assumptions. The full number of observations after each step can be seen in Table 3.1.

Table 3.1

The number of observations of the analysis sample after each step of data merging and cleaning.

Step	Observations	Dropped from previous
Raw NVM transactions	1,922,124	
After NVM transaction filters	1,859,714	62,410
Removing previously treated observations	1,828,783	30,931
SES observations after steps 1-3	1,450,230	378,553

3.3. Exploration of the Working Dataset

3.3.1. Treatment Exposure

After combining the data, it becomes clear that turbine exposure is limited to a minority of all transactions. Approximately 8.8% of transactions are classified as ever-treated, indicating that they will at some point be exposed to at least one turbine. In contrast, around 7.2% are treated at the

³As seen in Figure 3.7, the inland water bodies extend into the North Sea, meaning that some turbines were initially misclassified. Based on visual inspection and manual turbine type assignment by coordinates, these are reclassified as offshore turbines at sea. Turbines out of range of these water bodies were also classified in the same manner.

moment of sale, aligning with expectations from previous Dutch studies. The vast majority of transactions remain unexposed to onshore turbines, reflecting their uneven distribution across the Netherlands (see Figure 3.4).

The transaction prices in the resulting dataset vary widely, with a median of approximately 265,000 euros and a mean of around 318,905 euros, indicating a right-skewed distribution typical of housing markets. This is caused by a few outliers, which are houses with very large transaction prices, driving the average up while the median price remains largely unaffected. A breakdown of housing characteristics in relation to turbine exposure will be provided in Table 4.1.

3

Onshore treatment exposure within 2 km




Figure 3.8: A heatmap showcasing the distribution of the share in properties exposed to onshore turbines.

When the turbine distributions were discussed in subsection 3.1.2, it was already clear that turbines are not distributed evenly across the Netherlands. Figure 3.8 maps the share of treated transactions in the Netherlands, highlighting that turbine exposure is localised to specific regions, with significant differences in the percentage of treated transactions. This irregular spatial distribution

of exposure is essential to the analysis: it shapes both the identifying variation used for causal inference in subsection 3.3.2, and impacts the socioeconomic composition of exposed areas discussed in subsection 3.3.3, which will be discussed next.

3.3.2. Identifying Variation

The Difference-in-Differences (DiD) specification relies on comparing pre- and post-treatment units within the temporal window and requires sufficient observations to accurately estimate the treatment effect. This section determines the identifying variation that is available with the data sources along 3 lines: (i) determining the evolution of houses with turbine exposure over the years, (ii) determining whether newly constructed onshore turbines lead to new exposure within areas, and (iii) the distribution between the number of pre-treatment and post-treatment observations across 4-digit postal code (PC4) areas.

The share of property transactions within 2 kilometres of an onshore operating wind turbine is plotted in Figure 3.9. While the share fluctuates modestly, only a slight upward trend is evident. This contrasts the findings by [Eichholtz et al. \(2023\)](#), who found that the share rose rapidly from 0% to 11% from 1987-2010, after which it remained roughly constant with no further change.⁴ This implies that prior Dutch literature has data throughout the initial adoption phase, where a lot of transactions went from unexposed to exposed to turbines, which all predates our sample window regarding time, meaning the nearly constant results were to be expected.




Figure 3.9: The percentage of housing transactions that have 1 or more turbines within 2 kilometres over the years.

The near-constant share of sold properties with exposure is consistent with two scenarios: (i) there are few newly constructed turbines from 2010 to 2023, or (ii) many turbines are constructed, but add very little exposure to sold properties. Figure 3.10 highlights that during 2010 to 2023, while 1,435 onshore turbines were constructed, this led to only 312 PC4 areas in which an observation first experienced turbine exposure. This means that, on average, each newly constructed turbine leads to approximately 0.22 newly exposed PC4 locations. The near constant share of sold properties with turbine exposure, therefore, lies in the added exposure per turbine, rather than the number of

⁴Note that the higher share in 2010 is caused by a higher radius definition of 2.5 kilometres ([Eichholtz et al., 2023](#)).

turbines constructed between 2010 and 2023.




Figure 3.10: The number of onshore turbine constructions compared to the number of PC4 areas, newly exposed to onshore turbines.

The low exposure per turbine can be caused by two mechanisms. Firstly, new turbines may be placed in regions already exposed to older turbines, increasing turbine density rather than exposing new properties to turbines. Secondly, turbines could be placed in regions with relatively low populations to minimise any adverse effects. Figure 3.11 showcases the turbine placement per construction year over the Netherlands, and it is consistent with both mechanisms. Many newly constructed turbines are placed in regions where older turbines already exist. Simultaneously, new clusters of turbines are also placed in less densely populated areas, such as in the provinces of Groningen, Drenthe, Limburg, and Gelderland.

Across all PC4 regions with observations that ever experienced turbine exposure in the sample, the distributions of pre-construction and post-construction observations vary widely. Figure 3.12 highlight that quite a large number of PC4 regions do not contain any sales before the relevant turbines were constructed, indicating that these were constructed before the housing transactions sample starts in 2010. Furthermore, Figure 3.12 highlights that, even when there are pre- or post-construction observations within PC4 regions, there are either many or very few. If one looks at the minority side of the PC4 regions with both pre- and post-construction observations, then the median number of observations is 19, the 25th percentile lies at 4 observations, and 38.4% of all PC4 regions have fewer than 10 observations.⁵

Figure 3.13 decomposes the total merged sample per PC4 region into categories, highlighting that not all PC4 areas contribute equally for identification. Of the 4,002 PC4 regions, 2,941 have no observations that ever experience turbine exposure, so these will be the control group. Among the 1,061 ever-treated PC4 areas, 195 were already fully exposed to turbines and thus have no pre-construction observations, as all turbines were constructed before 2010, while 36 areas were only recently exposed to turbines and thus have no post-construction observations. These remaining 830 PC4s have observations from both before and after exposure to nearby turbines, comprising 422,136 observations out of the initial 1.86 million. Note that of these, 72.6% are in the control group and 27.4% are in the treated group, corresponding to 306,370 and 115,766 observations, respectively.

⁵With location fixed effects at the postal-code level, estimating the treatment effect requires observations from both before and after a turbine becomes operational within the same location. The minority side is the smaller group between the pre- and post-treatment observations within that location. Since identification depends on comparing these two periods, the size of the smaller group limits the amount of variation that the area can provide.




Figure 3.11: Onshore turbines per construction period scattered over the Netherlands.

The identifying sample is the subset which allows causal inference. Whether 830 PC4 areas with 422,136 housing transactions provide adequate identification depends not only on the number of observations, but on the within-PC4 distribution of observations across pre-construction and post-construction groups, which is examined next.

Taken together, the descriptive diagnostics showcase both the size and the structure of the identifying variation within the merged dataset. Of the 4,002 PC4 areas, there are 830 regions that contain both the pre- and post-construction observations. While these still account for over 422,000 observations, the minority side within PC4 regions has a median value of 19 observations, while a substantial percentage has fewer than 10. The combination of an already saturated wind turbine landscape across the Netherlands and a relatively short temporal range limits the effectiveness of




Figure 3.12: Distributions of the number of pre- and post-turbine construction observations.




Figure 3.13: Decompositions of the merged sample regarding the effective identifying sample for causal inference.

this data for inferring small, precise causal effects.

3.3.3. Socioeconomic Composition of Dataset

After merging the Socioeconomic Status (SES) data with the housing transactions data by the Nederlandse Vereniging van Makelaars (NVM), where SES variables are added to housing transactions at the PC4 level for the year of the housing transaction between 2014 and 2023, it can be observed that SES are not homogeneous or independent and identically distributed (i.i.d.). Rather, Figure 3.15 showcases that the overall composite SES distribution, along with the financial, educational, and labour-history scores, shows signs of clustering in larger cities. Interestingly, while the educational scores are higher in these larger cities, their financial wealth and labour-history scores are lower, which seems counterintuitive. This could highlight that the spread, or inequality within these PC4 regions, is high, which is confirmed in Figure B.5 of Appendix B. This not only means that SES variables differ spatially over the Netherlands, but there can also be large differences in SES variables even within specific localised areas.




Figure 3.14: A comparison between the number of observations before, after, and on the minority side of the turbine exposure for identifying PC4 regions.

Table 3.2

Distribution of pre-treatment, post-treatment, and minority side observations within identifying PC4 areas.

	Pre-treatment observations	Post-treatment observations	Minority side
Mean	369.1	139.5	78.4
10 th percentile	6.0	2.0	2.0
25 th percentile	37.0	7.0	4.0
Median	174.0	38.0	19.0
75 th percentile	556.5	152.0	95.8
90 th percentile	996.5	421.0	244.2
percentage ≥ 5	91.1%	80.5%	73.6%
percentage ≥ 10	87.8%	69.9%	61.6%
percentage ≥ 20	81.6%	60.4%	49.5%

While SES data differ spatially, it should also be mentioned that socioeconomic conditions develop over time. Figure 3.16 highlights trends of transaction-weighted and PC4-year weighted SES scores. The composite score shows an S-curve, initially declining marginally before increasing and smoothing out over time. It is also noteworthy that the educational and labour-history sub-scores show a similar trend, while the financial sub-score is roughly constant. Furthermore, across all panels, the transaction-weighted trend consistently lies above the PC4-year-weighted trend, indicating that housing transactions frequently occur in high SES neighbourhoods. However, both weight-




Figure 3.15: The distribution of socioeconomic status across the Netherlands.

ing schemes demonstrate nearly identical trend patterns, indicating that SES development is robust across weighting schemes.




Figure 3.16: The average socioeconomic status in the Netherlands demonstrates an s-curve over time.

Figure 3.17 compares the median transaction price per square metre across composite SES quintiles. A clear socioeconomic gradient is evident: Transactions in PC4 regions with higher composite SES tend to be more expensive when controlling for the living area, and this relation holds over time as well. This confirms that SES is correlated with housing prices, even after conditioning on living area. This relationship is relevant because earlier work (e.g., see [Dröes & Koster, 2016](#)) shows that wind turbines tend to be sited in areas with lower housing prices, such as rural areas. If properties with turbine exposure also follow different socioeconomic trajectories over time, and SES itself is correlated with property prices, then price trends may differ systematically between treated and untreated properties for reasons unrelated to turbine exposure.




Figure 3.17: The socioeconomic score quintiles highlight that a higher score is linked to a higher transaction price, conditioned on living area.

Figure 3.18 demonstrates the difference between ever-treated and never-treated properties in the median price index, and showcases these per SES quartile. While ever-treated and never-treated properties within the same quartile mostly demonstrate parallel trends (with the exception of the

lowest SES quartile from 2020 onwards), it should be noted that different SES quartiles have different price growth rates. The lower SES quartile exhibits faster price growth per PC4 region than the higher SES-PC4s. This shows that SES values are not only associated with property prices, but also with price growth over time.




Figure 3.18: The median price index of socioeconomic score quantiles mostly highlights similar trends for ever- and never-treated properties, but price growth is more pronounced in neighbourhoods with lower socioeconomic scores.

Figure 3.19 highlights the standardised differences across all SES variables between treated and untreated properties, showcasing large differences between treated and untreated PC4 regions, while there are also large differences in SES compositions between onshore and offshore exposed areas. Areas with onshore turbine exposure show little difference in wealth variables but have a noticeably higher share in middle education at the expense of higher education, and are more continuously employed. Regions with offshore turbine exposure show large differences between untreated groups in wealth, income, and assets, and are also more highly educated. This results in having a very high SES value, whereas areas with onshore turbine exposure do not exhibit these large differences between treated and untreated areas. It should also be noted that both onshore and offshore exposed units score low on SES spread, meaning that these PC4 areas with treated units are more homogeneous than their untreated counterparts.

Figure 3.20 exhibits that over time and after adjusting for national trends, the composite SES and sub-scores diverge between ever-treated and never-treated observations. Across the composite score and all sub-scores, it is evident that treated and never-treated observations follow different trends regarding socioeconomic development over time.

The same patterns remain visible in Figure 3.21, where the SES variables have been adjusted for urbanisation-related effects per year. Together with Figure 3.20, this could imply that treated and never-treated properties have different socioeconomic development over time. The divergence is not limited to the aggregate SES measure, but is also visible across several underlying SES dimensions. This observation is particularly relevant given the earlier finding that SES is strongly associated with both housing prices and housing price appreciation. Combined with the substantial baseline SES differences between treated and untreated properties, the divergent SES trajectories




Figure 3.19: The standardised differences between treated and non-treated properties exhibit that differences in socioeconomic status exist.




Figure 3.20: The socioeconomic score, adjusted for national trends, shows diverging socioeconomic trends over time.

imply that the parallel-trends assumption in housing prices may not automatically hold in the absence of additional controls and fixed effects.




Figure 3.21: The socioeconomic score, adjusted for urbanisation trends, shows diverging socioeconomic trends over time.

Research Design and Methodology

This chapter introduces the causal specifications used for the onshore analysis and the descriptive strategy used for the offshore analysis. It presents the DiD identification strategy and the baseline onshore specification, explains the robustness and heterogeneity analyses, describes how socioeconomic controls are incorporated, and shows the descriptive approach taken for offshore turbines. The methodological choices, such as the use of an estimator suitable for staggered adoption and fine-grained fixed effects, will determine the credibility of the subsequent results.

4.1. Identification Strategy

To estimate the effects of onshore wind turbines on housing prices, a Difference-in-Differences (DiD) identification strategy will be employed. DiD compares the changes in property prices for treated properties, properties that at some time t get exposed to turbines, with changes in prices for comparable non-treated properties. The core assumption is that, in the absence of treatment, the treated group would have followed the same trend as the non-treated group. Under this *parallel trends* assumption, the non-treated or control group provides a credible counterfactual for how treated properties would have evolved if no turbine had been installed.

The main idea behind this identification strategy is illustrated in Figure 4.1. Before the turbine exposure, the treated and control groups exhibit identical trends. After the turbine is installed, the treated group diverges from the counterfactual trend approximated by the control group because of turbine exposure. One can now determine the differences in outcomes between pre-treatment and post-treatment for both the treated and control groups. The treatment effect equals the difference between the treated and control groups before and after treatment; hence the name "Difference-in-Differences".

In the Dutch context of wind turbine adoption, treatment does not occur simultaneously for all units at a single point in time. Instead, wind turbines are constructed across multiple periods and locations throughout the Netherlands, implying that properties are treated at different times. While this reflects the real-world timing of wind turbine development and decision-making, the staggered adoption setting can bias Two-Way Fixed Effects (TWFE) estimates. The reason is that TWFE specifications pool all observations into a single regression, which relies on comparisons that are not suitable for causal inference. The TWFE specifications may rely on 'forbidden comparisons' in which already-treated units are implicitly used as controls for units that have not yet been treated, even though the already-treated units no longer provide a valid counterfactual ([Goodman-Bacon, 2021](#)). He has shown that the coefficient of interest could be subjected to negative weighting, or even a sign reversal of the true treatment effect. Given the staggered nature of turbine placement, relying on TWFE specifications would risk biasing the average treatment effect of the treated due to forbidden comparisons between treated and untreated units across different time points.

To address the limitations imposed by the standard TWFE specification under staggered adoption, modern DiD estimators designed for staggered treatment timing will be used. The estimator de-




Figure 4.1: An illustration of the treatment effect of the Difference-in-Differences identification strategy.

veloped by [Sun & Abraham \(2021\)](#) will be used in the main specifications. This method avoids forbidden comparisons by explicitly determining an event time based on a cohort definition and the treatment time, thereby avoiding TWFE biases. A treatment effect will then be estimated for each cohort, and these estimates can be aggregated into an Average Treatment Effect of the Treated (ATT). The same estimator will also be utilised in heterogeneity analyses, where multiple ATTs will be estimated for different sub-samples and then compared. This reduces the number of observations per sample split but allows assessment of whether different sample splits affect turbine effects on housing prices.

To test whether the differences between these sample splits are statistically significant, a traditional TWFE specification will be used where covariates are multiplied by the treatment indicator. This differs from the [Sun & Abraham \(2021\)](#) estimator, which multiplies cohorts by relative time dummies and aggregates them into effects, rather than estimating a single treatment coefficient under TWFE specifications.¹

4.2. Baseline Specification

To estimate the effects of turbines on property prices, multiple Difference-in-Differences (DiD) specifications will be utilised with the [Sun & Abraham \(2021\)](#) estimator. A relatively simple and classic specification is shown in Equation 4.1. It regresses the logarithm of the transaction price P_{it} on the treatment indicator D_{it} , which equals 1 if there is an operational turbine within d kilometres of property i at time t . Following [Dröes & Koster \(2016\)](#), the distance d for a property to be treated is defined as 2 kilometres.² This threshold is empirically motivated: the authors argue that the main

¹The estimator by [Callaway & Sant'Anna \(2021\)](#) was also considered for baseline and heterogeneity analyses of turbine effects. Unlike the approach in [Sun & Abraham \(2021\)](#), it identifies group-time ATTs under a conditional parallel trends assumption and allows estimation via outcome regression, inverse probability weighting, or doubly robust methods ([Callaway & Sant'Anna, 2021](#)). Ultimately, the [Callaway & Sant'Anna \(2021\)](#) estimator was not used, as in the Dutch setting, the conditional parallel trends assumption was difficult to defend: once covariates were added, the pre-treatment design became singular, leaving only living space available for adjustment, which would risk omitted variable bias.

²Do note that [Eichholtz et al. \(2023\)](#) use a different specification in which they define distance d as 2.5 kilometres. Such a specification will also be tested in section 5.1. However, rather than being explicitly defined for turbine treatment, it is also used for other energy sources, such as coal, gas, and biomass plants. For this reason, it was not chosen to be the main treatment distance.

local externalities of wind turbines, including shadow flicker, noise, and visual disamenity, are concentrated within a limited range, and their distance-profile estimates show that the house-price effect decreases with distance and is no longer statistically significant beyond approximately 2 kilometres. As newer turbines are generally larger and may have different impacts, alternative distance definitions will be assessed in section 4.3. Additionally, because the observations in the sample span multiple time periods, these location-time effects will be accounted for using location-time fixed effects γ_{mt} , which allow each municipality to follow its own local housing market trajectory.³ This accounts for the heterogeneity of the Netherlands, where local unobserved time shocks may evolve very differently across municipalities. For instance, a very fast-growing company like ASML could (temporarily) distort housing markets at the municipal level in ways that are not representative of the broader housing market. By including municipality-year location-time effects, these unobserved local time trends are absorbed, reducing the risk that the estimated Average Treatment Effect of the Treated (ATT) is driven by unobserved localised temporal effects, rather than the turbine exposure itself. Finally, ε_{it} denotes the error term.

$$\ln P_{it} = \beta D_{it} + \gamma_{mt} + \varepsilon_{it} \quad (4.1)$$

While Equation 4.1 provides an intuitive starting point for estimating the ATT, it does not yet account for possible differences in observed housing characteristics between treated and control properties. To address this, X_{it}^k is included, which is a vector of housing characteristics such as living area, maintenance quality, and house type, where k denotes the number of housing characteristics. Each of these housing characteristic variables will receive its own coefficient, such that the estimated treatment effect is not confounded by differences in observed housing attributes. The full list of housing characteristics is reported in Table E.1, while Table 4.1 compares these characteristics between the treated and untreated units.

$$\ln P_{it} = \beta D_{it} + \delta^k X_{it}^k + \gamma_{mt} + \varepsilon_{it} \quad (4.2)$$

While controlling for hedonic housing attributes accounts for observable differences in housing quality, it does not fully resolve concerns about unobserved spatial sorting. Properties with identical housing attributes may sell for systematically different prices across locations because some areas are more desirable, more urban, or consistently more constrained by housing supply than others. For instance, a house in Amsterdam is likely to sell for a higher price than the same property in a rural municipality, not because of structural differences but because of local amenities, land scarcity, and other housing market conditions.

To address this issue, location fixed effects will be introduced in Equation 4.3 and Equation 4.4. In these equations, α_j and α_k captures all time-invariant characteristics of regions j and k , where j and k are defined at the 4-digit postal code (PC4) and the 6-digit postal code (PC6) level, respectively. Identification will then come from properties within the same postal code region, so that houses are compared only to other houses in the same local market. This means that any time-invariant effect at the regional level is absorbed by the location fixed effects, strengthening the causal interpretation of the estimated ATT. These location fixed effects are similar to the location-time fixed effects, but operate at a very fine-grained level, whereas the location-time fixed effects are more coarse but are allowed to vary over time.

$$\ln P_{it} = \beta D_{it} + \delta^k X_{it}^k + \gamma_{mt} + \alpha_j + \varepsilon_{it} \quad (4.3)$$

$$\ln P_{it} = \beta D_{it} + \delta^k X_{it}^k + \gamma_{mt} + \alpha_k + \varepsilon_{it} \quad (4.4)$$

³Personal communication with Dröes and Koster (May 2026) confirmed the importance of municipality-year fixed effects to absorb local time-varying shocks, which was adopted in the main DiD specifications.

Another concern in hedonic pricing models is that unobserved locational characteristics aren't fully captured by using location fixed effects. [Gong \(2017\)](#) notes that within a neighbourhood, the value of properties with similar housing characteristics may still differ significantly, depending on the exact location. While including fixed effects at the fine-grained postal code level will reduce the likelihood of biased estimates, [Gong \(2017\)](#) states that these cannot fully capture locational effects, noting that some omitted variable bias and spatial autocorrelation will likely remain. To combat these concerns, standard errors are clustered at the neighbourhood level rather than assumed to be independent and identically distributed (i.i.d.). This means that within a neighbourhood, correlation between observations is allowed, further strengthening the causal interpretation.⁴

The specifications shown so far all rely on comparisons between treated units and units not exposed to turbines. Table 4.1 shows the housing characteristics of the treated and control units. It is clear that properties with turbine exposure are, on average, cheaper and smaller.⁵ This confirms that wind turbines are not randomly placed across the Netherlands, which further highlights the need for rich controls and fine-grained fixed effects of the baseline DiD specifications. Note that these controls and fixed effects do not explicitly test or ensure parallel trends. Instead, they mitigate concerns regarding endogenous turbine placement and observed compositional differences between treated and control properties. By absorbing permanent PC6 differences and municipality-year-specific housing market shocks, they help make the parallel trends assumption more credible.

Table 4.1

Descriptive statistics: treated and control transactions.

	<2 km of a turbine				>2 km of a turbine			
	Mean	SD	Min	Max	Mean	SD	Min	Max
Transaction price (EUR)	298,376	169,971	36,200	7,200,000	320,490	210,868	33,750	10,825,000
Usable living area (sqm)	117.577	46.196	25.000	1,250.000	119.320	50.879	10.000	2,288.000
Number of rooms	4.623	1.422	1.000	29.000	4.636	1.605	1.000	99.000
Maintenance inside (1-9)	6.973	1.038	1.000	9.000	6.964	1.052	1.000	9.000
Maintenance outside (1-9)	7.013	0.941	1.000	9.000	7.017	0.912	1.000	9.000
Listed monument	0.010	0.097	0.000	1.000	0.010	0.102	0.000	1.000
Built before 1945	0.227	0.419	0.000	1.000	0.226	0.418	0.000	1.000
Built 1945-1970	0.194	0.395	0.000	1.000	0.222	0.416	0.000	1.000
Built 1971-1990	0.255	0.436	0.000	1.000	0.276	0.447	0.000	1.000
Built 1991-2010	0.274	0.446	0.000	1.000	0.239	0.427	0.000	1.000
Built after 2010	0.051	0.220	0.000	1.000	0.037	0.188	0.000	1.000
Apartment	0.173	0.378	0.000	1.000	0.255	0.436	0.000	1.000
Bungalow	0.019	0.138	0.000	1.000	0.022	0.147	0.000	1.000

Continued on next page.

⁴As a robustness check, standard errors were also clustered at the PC4 level, which follows the example of [Dröes & Koster \(2016\)](#). The effects of the main specification remain unchanged. This suggests that the results are not sensitive to the level at which standard errors are clustered.

⁵The full list of housing characteristics are shown in Table E.1.

	<2 km of a turbine				>2 km of a turbine			
	Mean	SD	Min	Max	Mean	SD	Min	Max
Manor, farmhouse, or country home	0.062	0.241	0.000	1.000	0.065	0.246	0.000	1.000
Penthouse	0.004	0.061	0.000	1.000	0.004	0.067	0.000	1.000
Single-family home	0.693	0.461	0.000	1.000	0.595	0.491	0.000	1.000
Villa	0.027	0.163	0.000	1.000	0.031	0.174	0.000	1.000
Other residential	0.022	0.146	0.000	1.000	0.027	0.162	0.000	1.000
Observations	131,398				1,697,673			

4.3. Robustness Analyses

To assess the robustness of the baseline result of chapter 5, three sets of robustness analyses examine how the estimated effect responds to alternative specification choices. This will provide further insight into the strengths and the confidence of the evidence on Dutch turbine effects using recent turbine and housing transaction data.

1. **Anticipation effects:** It may be possible that prices for treated and non-treated units diverge prior to a unit's treatment. To assess whether the results are affected by anticipation or pre-treatment deviations, the baseline specification is re-estimated after excluding selected pre-treatment event times. Two exclusions are considered: (i) three years prior to treatment, following [Dröes & Koster \(2016\)](#), where they found significant pre-trends, and (ii) earlier pre-treatment years that show statistically significant deviations in the event-study plot.
2. **Alternative control groups:** The control group is restricted to properties located closer to turbines. Instead of including all properties beyond two kilometres as controls, alternative specifications use control groups located within 2-3, 2-5, and 3-5 kilometres from the nearest turbine. These specifications determine whether the baseline estimate is robust across all control groups.
3. **Alternative treatment distances:** To test whether turbine effects are highly localised, the same specification is used, but now with a treatment distance d of 2.5 kilometres. This aligns with the treatment distance used by [Eichholtz et al. \(2023\)](#), who found statistically significant results with those treatment definitions. These specifications test whether the turbine effects are localised in this dataset and whether adding a buffer layer affects the estimated effect.

4.4. Heterogeneity Analyses

The baseline specification estimates an average treatment effect across all treated properties. To examine whether this average hides variation in turbine effects, four dimensions of heterogeneity are analysed: the number of turbines, turbine height, the urbanisation level of the surrounding area, and the distance to the nearest turbine.

Each heterogeneity analysis applies the 6-digit postal code (PC6) fixed effects specification, in which subsamples are defined by their characteristics at the initial treatment of a property. This means that the subsamples are not dependent on additional constructed turbines in later years, but only consider turbines within 2 kilometres of the initial treatment year. This time-invariant definition is important because it ensures that heterogeneity is attributed to differences during the initial treatment year, rather than to changes in exposure over time.⁶ In other words, the analysis com-

⁶It will be noted in subsection 5.1.3 if treatment effects may not reflect a clean estimate due to turbines constructed

pares properties that were first treated under different local turbine exposures, thereby allowing determination of the heterogeneity in turbine effects across initial exposure types.

The split-sample approach is necessary because the [Sun & Abraham \(2021\)](#) estimator does not support interaction effects as seen in Two-Way Fixed Effects (TWFE) specifications. While the TWFE specification would include an interaction term between the treatment indicator and the heterogeneity variable, the modern estimator requires that heterogeneity be examined through separate regressions on subsamples. Each subsample retains the full set of never-treated control properties, ensuring that the average treatment effects across heterogeneity groups are not attributable to differences in the control groups. TWFE specifications will be used with interaction effects to see whether differences between sub-groups are statistically significant.

4.5. Socioeconomic Score Extensions

Section 3.3.3 documents that socioeconomic conditions vary substantially across the Netherlands and are correlated with both housing prices and turbine exposure. This raises the question of whether part of the estimated turbine effects reflects underlying socioeconomic differences between treated and control areas rather than the turbines themselves. To assess this, the baseline specification is extended by adding (yearly) time-varying socioeconomic covariates at the 4-digit postal code (PC4) level. The new model is thus defined as

$$\ln P_{it} = \beta D_{it} + \delta X_{it} + \gamma_{mt} + \zeta S_{jt} + \alpha_k + \varepsilon_{it}, \quad (4.5)$$

where S_{jt} contains standardised socioeconomic score variables for PC4 area j in year t . Four Socioeconomic Status (SES) specifications are estimated: the aggregate socioeconomic score and three sub-scores capturing financial position, educational attainment, and labour market history. The aggregate SES score is treated as the preferred SES variable because it summarises the overall socioeconomic position of the PC4 area. All SES variables are standardised to have a mean value of zero and a standard deviation of one to improve interpretation between socioeconomic variables.⁷

The Central Bureau of Statistics (CBS) socioeconomic data are only available from 2014 onward, so the model extension is estimated on the 2014-2023 subsample. Additionally, the analysis is restricted to observations that can be matched to PC4-level SES data, which the CBS records only for neighbourhoods with a sufficient number of households. The SES extension is therefore best interpreted as a test of whether the baseline turbine effect remains after controlling for observed socioeconomic composition at the PC4 level.

4.6. Descriptive Offshore Analysis

So far, specifications have been introduced to estimate the effects of onshore turbines. However, the methodology developed for onshore turbines does not transfer directly to offshore wind farms. There are four characteristics of offshore wind turbines in the Netherlands that would warrant a descriptive analysis rather than a causal approach.

First, there is only limited identifying variation. Dutch offshore wind turbines were constructed only recently and in a small number of offshore clusters, as seen in Figure 3.3 and Figure 3.4, with turbines within each offshore wind park becoming operational at approximately the same time. The recent construction of offshore wind farms has led to few post-treatment observations,

at later points in time.

⁷It should be mentioned that the sub-scores, especially the financial and labour history SES values, have a high correlation ($r = 0.89$). This makes it difficult to identify the relative contribution when only a few SES controls are introduced. The three socioeconomic sub-score variables are therefore primarily entered separately, while the aggregate SES score will capture the overall status of a PC4 region.

and the clustered timing reduces the number of distinct treatment cohorts. After all, Figure 3.3 shows that offshore turbines have only been constructed in 4 out of the 14 years during the 2010-2023 sample. Together, it is hypothesised that these features leave little identifying variation for Difference-in-Differences (DiD) designs under staggered adoption, thereby limiting the ability to effectively estimate the effects of offshore turbines on property prices.

Second, to properly estimate offshore turbine effects, it is hypothesised that non-coastal observations should be excluded, resulting in a relatively small sample. Offshore turbines are visible only from coastal areas, but these deviate from the national average. In chapter 3, it is shown that coastal properties within 30 kilometres from the coastline exhibit different socioeconomic characteristics than onshore-treated properties, and that socioeconomic status correlates with both housing prices and price trajectories. To increase the credibility of any results, the offshore analysis would therefore need to restrict the sample to coastal communities only and thus compare treated coastal properties with untreated coastal properties, rather than with an unrestricted national control group. However, restricting the sample to only coastal communities would greatly reduce the sample size.

The two points mentioned above lead to the third reason: the combination of few treated coastal properties and a limited coastal control group makes precise estimation difficult under conventional DiD designs. Methods that synthetically construct counterfactuals to increase sample size, such as synthetic DiD specifications that operate without potential bias under staggered adoption, may be better suited to the offshore setting. However, specifications that take these points into account methodologically fall outside of the scope of this thesis and will be noted down as potential future research.

The fourth and final point is that the externalities, and thus the treatment definition, of offshore turbines differ from those of onshore turbines. The distance-based exposure measure suits onshore turbines because their externalities are highly localised, operating through noise, shadow flicker, and visibility, as noted in subsection 2.2.2. For offshore turbines, the dominant externality is visibility, as noise and shadow flicker are no longer relevant at very distant placements from the coastline. Some offshore studies find that visibility alone has statistically insignificant effects on prices (see section 2.3). If visibility is the primary channel, a credible specification would require data regarding turbine visibility in an offshore setting. A distance-based treatment definition, which has been applied in onshore settings, may not adequately capture the relevant offshore turbine exposure.

For these reasons, the offshore analysis is presented descriptively in section 5.2, documenting price trends in coastal communities relative to offshore wind farm development without claiming causal identification. Making any causal claim would require a methodology suited to offshore turbine exposure.

Results

This chapter presents the empirical findings regarding turbine effects. It first presents the baseline onshore effect and tests its robustness, then the heterogeneity of effects across distance, turbine number, height, and urbanisation. Additionally, it covers socioeconomic extensions and provides a descriptive offshore analysis. These results answer the first and third research questions and help establish what is known about Dutch offshore turbines. These results, particularly regarding heterogeneity and socioeconomic context, will fuel the subsequent discussion.

5.1. Onshore Effects

5.1.1. Treatment Effects Under the Baseline Specification

Table 5.1 presents baseline estimates of the effect of onshore wind turbines on property prices in the Netherlands. The specifications progress from a standard Difference-in-Differences (DiD) model in column one to richer models that add housing hedonic attributes and increasingly restrictive fixed effects in columns 3 and 4.¹ These columns correspond from Equation 4.1 through Equation 4.4, and all estimates are based on the full housing transaction sample.

The classic DiD specification, described in Equation 4.1, shows that properties located within 2 kilometres of an operational turbine experience no statistically significant change in prices. However, this estimate is difficult to interpret because this specification does not account for any housing attributes or spatial price effects. As shown earlier in Table 4.1, turbines tend to be placed in areas where properties are systematically cheaper, making this specification vulnerable to the omission of both housing characteristics and locational aspects.

When accounting for housing characteristics, a significant shift is observed, as the estimated effect is a statistically significant estimate of -2.5% . The transition from an insignificant to a significant estimate underscores that, after accounting for potential composition bias, wind turbines appear to have an adverse effect on housing prices. Do note that locational aspects are not reflected in column two, so even though the results are statistically significant, it could be that the estimate is biased by locational characteristics instead of being a pure turbine effect.

After adding 4-digit postal code (PC4) fixed effects, the turbine effect becomes statistically insignificant again, with the point estimate remaining only slightly negative. This suggests that once comparisons are restricted to properties within the same PC4 area, substantial spatial heterogeneity is absorbed. Identification now relies solely on within-PC4 variation, which is limited (see subsection 3.3.2). Based on the standard error, effects larger than approximately $\pm 1.6\%$ can be ruled out at an 80% confidence level, implying that any remaining turbine effects, if present, are small in magnitude.²

¹The full regression table can be observed in Table E.2, which shows the housing characteristics as well instead of indicating that these are included.

²The Minimum Detectable Effect (MDE) was calculated as $MDE = (Z_{1-\alpha} + Z_{1-\beta}) \cdot SE$. Under conventional significance and power levels, this suggests that effects above 1.57% are unlikely to be present, as they would have been

Table 5.1

The baseline onshore effects of wind turbines on housing prices.

	log price			
	Classic DiD	Housing characteristics	PC4 fixed effects	PC6 fixed effects
	(1)	(2)	(3)	(4)
Wind turbine ≤ 2 km	0.0138 (0.0174)	-0.0247*** (0.0075)	-0.0011 (0.0056)	-0.0081* (0.0041)
Observations	1,828,775	1,828,775	1,828,724	1,779,283
R ²	0.35430	0.84292	0.87975	0.94535
Within R ²	0.00278	0.75741	0.76350	0.62462
Year-Municipality fixed effects	✓	✓	✓	✓
Housing characteristics (50)		✓	✓	✓
PC4 fixed effects			✓	
PC6 fixed effects				✓

Notes: The table reports the Average Treatment Effect of the Treated (ATT) estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

In a more granular specification where PC4 is replaced by 6-digit postal code (PC6), the estimated effect becomes marginally significant at the 10% level, with a magnitude of roughly -0.8% . This change suggests that the estimated turbine effects are sensitive to the level at which location fixed effects are applied, as PC6 fixed effects will absorb more spatial variation. At the same time, the magnitude of the point estimate remains small, so the turbine effects are relatively modest if present.³

To determine whether the parallel trends assumption holds for the specified models, Figure 5.1 presents the pre-trends, in which event-time ATTs are estimated to measure how prices evolve in each year before and after the turbine becomes operational. In theory, the pre-trends should show no effect before the treatment, suggesting that price movements are driven by factors other than turbine operation. The results show that specifications without strict locational fixed effects exhibit substantial pre-trends, which indicates that within-location comparisons due to fixed effects are required for a specification to gain causal credibility. Models with PC4 and PC6 fixed effects show no pre-trends in the three years leading up to the turbine's operational start, but between 4 and 7 years in advance, there are statistically significant price effects. These effects range from -1% to -2.1% , depending on the year relative to the treatment and the fixed-effect level. This suggests that the parallel trends assumption is only partially supported, which could indicate that turbine announcements cause an initial price decline of nearby properties. However, this interpretation remains speculative, as no data on turbine announcement dates are available in the datasets used.

detected with 80% confidence.

³Similar Two-Way Fixed Effects (TWFE) specifications under PC4 fixed effects estimate effects of -0.78% (significant at 5%), which is similar to the baseline specification with PC6 fixed effects, which estimated an effect of -0.81% . Under PC6 fixed effects, the TWFE estimate is -0.44% and not significant. The full TWFE replication is reported in section C.1 of Appendix C.

Alternatively, the negative pre-treatment estimates in the -7 to -4 range may reflect cohort composition. In staggered adoption settings, the event-time coefficients at distant leads are often identified mainly by later-treated cohorts, especially since the temporal range of the housing transaction data is already limited. These later-treated cohorts may differ systematically from earlier-treated units, suggesting that pre-treatment estimates do not necessarily reflect anticipation effects.⁴




Figure 5.1: Event study of the baseline specifications with the Sun & Abraham (2021) estimator with $t = -1$ as the reference year.

The baseline specification suggests a small negative effect of onshore wind turbines on property prices. With PC6 and municipality-year fixed effects, the estimate is -0.8% and only marginally significant. This is somewhat smaller than the effects reported in prior Dutch studies, which typically find price effects in the range of about -1.2% to -2.4% in baseline specifications with the full sample (Dröes & Koster, 2016; Dröes & Koster, 2021; Eichholtz et al., 2023). At the same time, the pattern is consistent with the Dutch literature as the effects are negative and highly localised. However, unlike some earlier work, the event study does not show a clean pre-treatment decline beginning two years before the turbine becomes operational. Instead, it shows significant pre-trends that appear four to seven years before turbines become operational. This may reflect anticipation effects, although the absence of turbine announcement dates prevents testing this hypothesis.

5.1.2. Robustness Analyses

Anticipation Effects

Table 5.2 presents the baseline specification with event times -3 , -2 , and -1 excluded, following Dröes & Koster (2016), who found significant anticipation effects in this period prior to the turbines becoming operational. In contrast to their findings, the event study in Figure 5.1 indicates that the same event-time years do not exhibit significant pre-trends. Therefore, the same effects in which

⁴A sample split by sale year, reported in Table C.3, was attempted to determine if the estimated ATT would differ between the sub-groups, but the results are not fruitful, likely due to the lack of identifying variation.

the ATT becomes more negative are not expected, but remain interesting, as anticipation effects are logically most likely to occur just before treatment.

Under specifications with strict location fixed effects, the results confirm this expectation. The point estimate with PC4 fixed effects is nearly identical to the baseline (-0.13% versus -0.11%), and both are statistically insignificant. At the same time, the estimate with PC6 fixed effects weakens slightly to -0.63% and loses statistical significance. This pattern indicates that the marginally significant baseline effect with PC6 fixed effects is not biased downward by units about to be treated, which is consistent with the absence of significant pre-trends in those years.

The results in columns one and two without spatial fixed effects highlight a different pattern. The Classic DiD coefficient moves from 1.4% to 2.5% , while the housing characteristics specification moves from -2.47% to -1.84% . These shifts may indicate that specifications without locational fixed effects are more sensitive to event time restrictions. This sensitivity further reinforces the case that specifications require strict locational fixed effects.

Table 5.2

The onshore effects of wind turbines on housing prices, excluding potential anticipation effects up to three years prior to treatment.

	log price			
	Classic DiD	Housing characteristics	PC4 fixed effects	PC6 fixed effects
	(1)	(2)	(3)	(4)
Wind turbine ≤ 2 km	0.0252 (0.0222)	-0.0184** (0.0093)	-0.0013 (0.0068)	-0.0063 (0.0042)
Observations	1,816,965	1,816,965	1,816,914	1,767,168
R ²	0.35469	0.84303	0.87988	0.94546
Within R ²	0.00184	0.75720	0.76361	0.62482
Year-Municipality fixed effects	✓	✓	✓	✓
Housing characteristics (50)		✓	✓	✓
PC4 fixed effects			✓	
PC6 fixed effects				✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. Potential anticipation effects are mitigated by excluding the event times years -3 , -2 , and -1 relative to the treatment year, the period in which [Dröes & Koster \(2016\)](#) found potential anticipation effects. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

Because the event study in Figure 5.1 reveals significant pre-trends in event times from -7 to -4 years, rather than in the years closer to the treatment, Table 5.3 reports the baseline specification where these earlier event time observations are excluded. This means that, for these specifications, event times at which parallel trends are not plausible are excluded.

The estimates move in the opposite direction from the previous exclusion. The PC6 specification increases in effect from -0.81% (marginally significant at 10%) to -0.92% and becomes statisti-

cally significant at the 5% level. Likewise, the PC4 point estimate becomes modestly negative, but remains statistically insignificant. If the pre-trends from 7 to 4 years prior to treatment reflected anticipation effects, this increase in ATT is to be expected. After all, the difference in price between these observations before and after treatment would be smaller, which is why excluding these observations would eliminate a downward bias.

However, it is unlikely that these years reflect clean anticipation effects, especially because these effects also fade in the subsequent years prior to treatment. The temporal range is from 2010 through 2023, with staggered adoption. This would indicate that event times of -7 to -4 years are likely to be observations from mid- to late-treated cohorts, as it is impossible for an early-treated cohort, say 2011, to have observations with an event time of -4 years. It could therefore be that these effects are not anticipation effects, but rather differences between the average treated unit and mid- to late-treated units. Such changes are consistent with diverging socioeconomic trajectories highlighted in subsection 3.3.3.

5

Table 5.3

The onshore effects of wind turbines on housing prices, excluding event times with pre-trend deviations.

	log price			
	Classic DiD	Housing characteristics	PC4 fixed effects	PC6 fixed effects
	(1)	(2)	(3)	(4)
Wind turbine ≤ 2 km	0.0132 (0.0174)	-0.0253*** (0.0075)	-0.0025 (0.0057)	-0.0092** (0.0042)
Observations	1,818,010	1,818,010	1,817,960	1,768,231
R ²	0.35390	0.84287	0.87968	0.94537
Within R ²	0.00176	0.75722	0.76356	0.62476
Year-Municipality fixed effects	✓	✓	✓	✓
Housing characteristics (50)		✓	✓	✓
PC4 fixed effects			✓	
PC6 fixed effects				✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. Potential anticipation effects are mitigated by excluding the event times in years -7 , -6 , -5 , and -4 relative to the treatment year, which showed pre-trends in the event study. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

Together, these two analyses suggest that the baseline estimate is not driven by anticipation in the standard sense but rather by cohort composition effects arising from the limited sample window. The relatively stable baseline point estimate across both exclusions, compared to the original baseline, suggests that the result is not driven by these significant pre-treatment event times. The estimated effect remains small throughout, shifting from -0.81% to -0.92% under PC6 fixed effects. So, while alternative control groups showed that the baseline estimate is fragile, the magnitude and sign are generally stable when the significant pre-treatment years are removed.

Alternative Control Groups

To assess whether the baseline result depends on the definition of the control group, Table 5.4 re-estimates the PC6 specification using alternative control groups with excluded observations. Column one reproduces the baseline, in which all properties beyond 2 km serve as controls. Columns two and three restrict the control group to properties within 2 to 3 kilometres and 2 to 5 kilometres of a turbine, respectively, ensuring that treated and control properties are more similar in overall location characteristics. Column four uses a 3 to 5 kilometre control group, sometimes referred to as a doughnut, as it excludes properties 2 to 3 kilometres from a turbine that may themselves experience price effects from nearby turbines. Including these transactions within the control group may bias the ATT toward zero, because the comparison group is itself partially treated and experiencing price declines. Therefore, by excluding these observations within 2 to 3 kilometres, spillover effects are mitigated.

Table 5.4

The onshore effects of wind turbines on housing prices with alternative control group definitions.

	log price			
	Full sample (no restrictions)	2-3 km control group	2-5 km control group	3-5 km control group
	(1)	(2)	(3)	(4)
Wind turbine ≤ 2 km	-0.0081* (0.0041)	0.0032 (0.0063)	0.0013 (0.0059)	-0.0044 (0.0121)
Observations	1,779,283	299,828	641,886	556,444
R ²	0.94535	0.94491	0.94670	0.94739
Within R ²	0.62462	0.62145	0.63684	0.63983
Housing characteristics (50)	✓	✓	✓	✓
Year-Municipality fixed effects	✓	✓	✓	✓
PC6 fixed effects	✓	✓	✓	✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator with various control group definitions, using the logarithm of the transaction price as the dependent variable. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

The results reveal that the marginally significant baseline estimate of -0.81% fades after restricting the control groups in various specifications. Under the 2-3 km and 2-5 kilometre control group definitions (columns two and three), the point estimate switches signs to small positive values, neither of which is statistically significant. Under the doughnut specification (column four), the estimate remains negative at -0.44% but is not statistically significant at any conventional level and is roughly half the baseline magnitude.

The sign change in columns two and three aligns with two possible explanations that the data cannot distinguish: (i) either treated properties do not systematically differ from nearby controls, or (ii) the nearby controls are affected by spillover effects that eliminate the treated-control difference, and thus treatment effects. The second hypothesis is directly tested in column four using the doughnut specification, with nearby properties removed. The fact that the doughnut estimate remains

negative, while the control groups with possible spillover effects turn positive (but insignificant), is consistent with spatial spillovers biasing the latter toward zero. However, this cannot be distinguished within the data. Overall, these findings suggest that the baseline estimate of 0.81% is small, fragile, and not robust to the choice of control group. It switches signs in two of three restricted definitions and remains statistically indistinguishable from zero under all of them. This is different from prior literature, as it was previously found that the estimated turbine effects would increase in absolute terms by restricting the control group to only nearby units of ever-treated units (Dröes & Koster, 2016; Dröes & Koster, 2021). This does not invalidate these findings, but it could also imply that recent Nederlandse Vereniging van Makelaars (NVM) data allow only limited identification.

Alternative Treatment Definition

To test whether an alternative treatment distance definition would also find any turbine effects, Table 5.5 reports a similar specification to the baseline specification, but now d is defined as 2.5 kilometres, meaning that properties are now treated within 2,500 metres of an operational turbine.

Table 5.5

The onshore effects of wind turbines on housing prices if the treatment range is increased to 2.5 kilometres.

	log price			
	Classic DiD	Housing characteristics	PC4 fixed effects	PC6 fixed effects
	(1)	(2)	(3)	(4)
Wind turbine ≤ 2.5 km	-0.0049 (0.0163)	-0.0274*** (0.0074)	-0.0071 (0.0050)	-0.0002 (0.0034)
Observations	1,828,775	1,828,775	1,828,724	1,779,283
R ²	0.35418	0.84288	0.87976	0.94535
Within R ²	0.00260	0.75734	0.76353	0.62462
Year-Municipality fixed effects	✓	✓	✓	✓
Housing characteristics (50)		✓	✓	✓
PC4 fixed effects			✓	
PC6 fixed effects				✓

Notes: The table reports the ATT estimates with the Sun & Abraham (2021) estimator, using the logarithm of the transaction price as the dependent variable. A unit is now considered treated if there is an operational turbine within 2.5 kilometres of the property at the time of the transaction. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

Columns one and two are very similar to the baseline results in terms of statistical significance, but it is noteworthy that the point estimates in both columns have become more negative. Additionally, the sign changed in column one from positive to negative, although it remains statistically insignificant. The estimated effect under PC6 fixed effects falls to -0.02% and is statistically insignificant, compared to -0.81% at the two-kilometre baseline result. This pattern is consistent with strongly localised turbine effects, as properties between 2 and 2.5 kilometres experience little measurable

price impact; including them as treated would shift the average effect toward zero. This would be consistent with the findings of [Dröes & Koster \(2016\)](#), who also found that turbine effects on housing prices were localised.

5.1.3. Heterogeneity Analyses

Number of Turbines

To examine whether the average turbine effects mask heterogeneity in treatment intensity, Table 5.6 shows the results split by turbine exposure. It differentiates the effects between properties that received one turbine and those that received two or more turbines during their initial treatment period. Note that the cohort definition is determined by the number of turbines during the initial treatment and that both specifications use the non-restricted control group sample.

Table 5.6

The onshore effects of wind turbines on housing prices depending on the number of turbines during the treatment period.

	log price	
	1 Turbine (1)	2+ Turbines (2)
Wind turbine ≤ 2 km	-0.0172*** (0.0052)	0.0049 (0.0066)
Observations	1,730,738	1,671,151
R ²	0.94553	0.94564
Within R ²	0.62482	0.62629
Housing characteristics (50)	✓	✓
Year-Municipality fixed effects	✓	✓
PC6 fixed effects	✓	✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. It showcases the differences in trajectories between units with 1 turbine and those with 2 or more turbines. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

It is shown that properties initially exposed to a single turbine exhibit an ATT of -1.72% , statistically significant at the 1% level. Properties initially exposed to two or more turbines show an estimated ATT of $+0.49\%$, not statistically significant at any conventional level, but Table 5.7 shows that this group is statistically different from the single-turbine group, where measured effects may even be slightly positive. These results differ from the baseline specification, which found an overall effect of -0.81% . This suggests that the modest adverse effect reflects opposing patterns across initial turbine intensity during the initial treatment period.

The results from Table 5.6 are counterintuitive, as it was initially expected that more turbines during the initial treatment period would lead to higher adverse effects. The effect between the 2 groups is also statistically significant, which is shown in Table 5.7.⁵ The estimate for single-turbine exposure is -1.72% , which closely aligns with the findings of [Dröes & Koster \(2016\)](#), who reported

⁵Note that Table 5.7 uses TWFE specifications as [Sun & Abraham \(2021\)](#) does not support standard interaction terms,

−1.85%, also significant at the 1% level. They also found no statistically significant results in a joint significance test when adding turbines as an interaction term. However, this estimate is on top of a −1.85% adverse effect, meaning that additional turbines do not have a larger effect than single turbines. This differs from the finding in column two, which shows an insignificant effect for two or more turbines. This could imply two things: (i) there may be selection bias of locations where multiple turbines are installed, as it is confirmed in Table 5.7 that clusters are installed in regions where pre-existing price trends were already negative, and (ii) it could be possible that there are modest adverse effects, but that the data lacks the statistical power to properly identify these effects.⁶

It should be explicitly noted that these cohort-based specifications define treatment intensity as the exposure during initial treatment, but 27.2% of the post-treatment observations in the single-turbine group are exposed to additional turbines in later stages. The share remains below 7.5% in the first five years after initial treatment, but rises in extended time horizons. This means that the observed effect of −1.72% may partially capture effects of subsequent turbines, which should be taken into account.

5

Turbine Height

Heterogeneity by turbine height is analysed using cohort definitions based on the turbine's tip height at the initial turbine exposure, with categories below 50 metres, between 50 and 150 metres, and above 150 metres, following the categories as [Dröes & Koster \(2021\)](#). The composition of these groups differs widely: small turbines (<50 metres in height) were predominantly installed before 2010, with a median installation year of 1994, meaning that only 2 PC4s have any pre-treatment observations. The opposite is true for the largest turbine category (>150 metres): They have only been installed since 2012 and have a median installation year of 2021, whereas the housing transaction data only goes up to 2023. They likely provide extensive pre-treatment observations, but only limited post-treatment observations. Additionally, this is the smallest group in terms of the number of treated observations, as only 4,666 properties were ever treated. Note that only 138 and 2,354 PC4 and PC6 regions have both pre- and post-treatment observations, respectively, meaning that there are very few regions that contribute to identification. Only the medium-height turbines (50-150 metres) offer substantial pre- and post-treatment variation within postal code regions for identification.

Table 5.8 shows the effects of various turbine heights. The medium height turbines have an ATT of −1.09%, significant at the 5% level. The estimate is lower, almost half, than that reported by [Dröes & Koster \(2021\)](#), who found an adverse effect of −2% for turbines of the same height, also significant at the 5% level. Compared to the baseline estimate, it is slightly more negative than the reported −0.81%. The small turbine category has an estimate of −4.73% with a large standard error, meaning that this estimate is not significant at any conventional level. This reflects the lack of pre-treatment observations noted earlier, and the difference is not statistically different from that of the medium-height turbine group, as shown in Table 5.7. The largest turbine category exhibits a counterintuitive result, with an estimate of +0.9%, significant at the 5% level, and statistically

which is why subgroups were used in the first place. However, because the overall magnitude and sign are generally consistent under strict fixed effects, this approach can be used as an approximation to determine whether differences between groups are statistically significant. These specifications are shown because, while the sub-group DiD analyses showed varying estimates for the treated sub-group units compared to the full control group, they do not indicate whether the differences between subgroups are statistically significant.

⁶Out of 161,460 treated properties, 69% of these properties were initially exposed to a single turbine, rather than multiple turbines. The lower number of observations with two or more turbines during initial exposure, combined with a narrow temporal range, could mean that there are not sufficient pre- and post-treatment observations within PC6 regions to properly identify any modest effect. If one applies PC4 instead of PC6 fixed effects, like in Table C.2, the sign flips for the multiple turbine category, although it remains statistically insignificant at any conventional level.

Table 5.7

The interaction effects of all the presented subgroups in heterogeneity analyses using a TWFE specification.

	log price				
	(1) Baseline TWFE	(2) Number of turbines	(3) turbine height	(4) Urban effects	(5) Distance effects
Wind turbine ≤ 2 km	-0.0044 (0.0037)	-0.0099** (0.0047)	-0.0163*** (0.0045)	-0.0023 (0.0037)	-0.0006 (0.0090)
2+ turbines		-0.0174*** (0.0063)			
Wind turbine ≤ 2 km $\times$ 2+ turbines		0.0184*** (0.0056)			
height < 50 m			-0.0026 (0.0477)		
height > 150 m			-0.0127** (0.0064)		
Wind turbine ≤ 2 km $\times$ height < 50 m			0.0240 (0.0460)		
Wind turbine ≤ 2 km $\times$ height > 150 m			0.0411*** (0.0063)		
Wind turbine ≤ 2 km $\times$ urban				-0.0040 (0.0069)	
distance 1-2 km					0.0008 (0.0047)
Wind turbine ≤ 2 km $\times$ distance 1-2 km					-0.0043 (0.0085)
Observations	1,779,284	1,779,284	1,779,284	1,779,284	1,779,284
R ²	0.94533	0.94533	0.94533	0.94533	0.94533
Within R ²	0.62445	0.62446	0.62449	0.62445	0.62445
Housing characteristics (50)	✓	✓	✓	✓	✓
PC6 fixed effects	✓	✓	✓	✓	✓
Year-Municipality fixed effects	✓	✓	✓	✓	✓

Notes: The table reports the ATT estimates utilising TWFE specifications, using the logarithm of the transaction price as the dependent variable. Interactions were utilised to assess whether the differences between subgroups used to estimate heterogeneity effects are statistically significant. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

different from the medium height turbine category shown in Table 5.7. As previously discussed, this could be an artefact of the limited number of pre- and post-treatment observations per postal code region. Table C.5 shows the same specification but now with PC4 fixed effects rather than at the PC6 level, and reports an effect of -1.69% for the largest turbines, while it reports statistically insignificant results for other turbine height categories. This shows that these specifications remain highly sensitive to the definition of fixed effects.

Table 5.8

The onshore effects of wind turbines on housing prices grouped by turbine height groups.

	log price		
	<50 metres	50-150 metres	>150 metres
	(1)	(2)	(3)
Wind turbine ≤ 2 km	-0.0473 (0.0493)	-0.0109** (0.0049)	0.0090* (0.0051)
Observations	1,659,481	1,718,800	1,647,047
R ²	0.94561	0.94555	0.94575
Within R ²	0.62594	0.62554	0.62609
Housing characteristics (50)	✓	✓	✓
Year-Municipality fixed effects	✓	✓	✓
PC6 fixed effects	✓	✓	✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. It highlights the distinct treatment effects of turbines across height categories. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

These results thus do not show the similar patterns reported by [Dröes & Koster \(2021\)](#), in which prices decrease with turbine size. The difference likely stems from the sample period rather than from a change in perceptions of turbines of various heights. The small number of observations within postal code regions for both small and large turbines precludes reliable estimates of turbine effects. Small turbines were predominantly built well before 2010, while the already very few larger turbines lack post-treatment observations. To properly identify the effects of various turbine heights, a longer sample period is required that captures the initial adoption period of turbines. However, this will likely not allow identification of the tallest turbine group, as these were only recently constructed.

Urban Effects

To determine whether urban properties experience a larger decline in property value due to turbine exposure, Table 5.9 shows the treatment effects of urban properties and those that are not located in urbanised regions. Urban properties are defined as those with more than 1,500 addresses per square kilometre, while non-urban properties have fewer addresses per square kilometre. Note that the control groups in both specifications are not filtered by urbanisation category and thus represent the full non-treated sample.

Column one shows that urban properties, those with over 1,500 addresses per square kilometre, experience adverse turbine effects estimated at -1.02% , significant at the 5% level. This is slightly higher than the average baseline estimate of -0.81% , indicating that urban properties, on average, experience greater property depreciation due to turbine exposure. The subgroup specification in column two is -0.51% and not statistically significant at any conventional level. Remember that [Dröes & Koster \(2016\)](#) and [Eichholtz et al. \(2023\)](#) found that urban properties experience at least

Table 5.9

The onshore effects of wind turbines on housing prices grouped by urbanisation category.

	log price	
	Urban units (1)	Non-urban units (2)
Wind turbine ≤ 2 km	-0.0102** (0.0052)	-0.0051 (0.0062)
Observations	1,698,269	1,704,115
R ²	0.94608	0.94504
Within R ²	0.62588	0.62524
Housing characteristics (50)	✓	✓
Year-Municipality fixed effects	✓	✓
PC6 fixed effects	✓	✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. It highlights the different treatment effects by the urbanisation category of treated units, with urban properties defined as those with more than 1,500 addresses per square kilometre. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

1.2% additional adverse price effects due to urbanisation when using the same definition.⁷ While the difference between the subgroup estimates of columns one and two appears very large, column four of Table 5.7 shows that urban properties do not statistically experience larger adverse price effects at any conventional significance level, which would contrast existing literature. The point estimate remains negative, which may suggest but does not prove that urban properties experience larger price declines. While an urbanisation effect has not been proven, the greater number of properties in urbanised areas suggests that it may still be undesirable to construct turbines there rather than in more rural locations. This is further supported by column two, which shows that non-urban properties did not, on average, experience a statistically significant price effect, unlike urbanised properties.

Distance Effects

To determine if properties closer to a turbine experience more effects on average, Table 5.10 shows the effects of two treated property groups, where they are categorised by the closest distance to the first constructed turbine. Do note that, based on the defined cohort structure, only the distance during the initial treatment period is considered, meaning that any additional turbine constructed later has no influence on the group a property is sorted into. Additionally, for both specifications, the control group has no distance-based restrictions.

Column one shows the effects of turbines when the initial turbine is constructed within one kilome-

⁷Do note that the urbanisation definitions differ: [Dröes & Koster \(2016\)](#) define urbanised properties as properties with at least 5,000 persons per square kilometre, but [Eichholtz et al. \(2023\)](#) use a urbanisation measure of five distinct categories. It is assumed that this aligns with the definition used in Table 5.9 as they use the same data source for housing transactions.

Table 5.10

The onshore effects of wind turbines on housing prices depending on the distance to the closest turbine during the treatment period.

	log price	
	Treated within 0-1 km (1)	Treated within 1-2 km (2)
Wind turbine treatment	-0.0292** (0.0148)	-0.0067* (0.0040)
Observations	1,637,940	1,763,931
R ²	0.94571	0.94545
Within R ²	0.62629	0.62480
Housing characteristics (50)	✓	✓
Year-Municipality fixed effects	✓	✓
PC6 fixed effects	✓	✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. It showcases the ATTs of the distance subgroups within 1 kilometre and those between 1 and 2 kilometres of a turbine. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

tre, estimated at -2.92% , significant at the 5% level. The effect for turbines initially placed farther away, at one to two kilometres, is estimated at -0.67% , significant at the 10% level.⁸ Note that while the estimates of the closest distance group are 4.3 times as high in absolute terms as the more distant group, Table 5.7 shows that the difference between these groups is not statistically significant. These results do not align with prior findings by [Dröes & Koster \(2021\)](#), who document that turbine effects generally decrease as the distance to the turbine gets larger up to 2,250 metres, after which they are indistinguishable from zero. However, it does align with the finding of [Eichholtz et al. \(2023\)](#), who found that turbine effects are roughly constant regardless of a turbine's distance from a property when opening.

5.1.4. Socioeconomic Score Extensions

To determine whether socioeconomic elements are partially driving the baseline results, Table 5.11 shows how the baseline results change after controlling for various standardised socioeconomic score variables, including an aggregated overall measure. Note that because Socioeconomic Status (SES) variables have only been recorded since 2014, all observations from the first four years of the sample period will be dropped. The baseline specification is re-estimated with an identical subset to match the specifications with SES data. Additionally, it should be noted that not all observations have SES data, most likely because the Central Bureau of Statistics (CBS)⁹ only documents data from neighbourhoods with sufficient households, which should be kept in mind when interpreting

⁸More detailed distance bands within the 0-1 km and 1-2 km ranges are not shown because the resulting subsamples lack the statistical power to detect marginal effect sizes.

⁹The CBS is the Netherlands' independent national statistical office, responsible for collecting, processing, and publishing official statistical information about the country.

the results.

Table 5.11

The baseline onshore effects of wind turbines on housing prices when controlling for SES variables.

	log price				
	(1) Baseline Result	(2) SES Financial	(3) SES Education	(4) SES Labour History	(5) SES Aggregated
Wind turbine ≤ 2 km	-0.0126** (0.0050)	-0.0109* (0.0062)	-0.0102* (0.0062)	-0.0107* (0.0062)	-0.0107* (0.0062)
Standardised financial SES		0.0138*** (0.0050)			
Standardised educational SES			0.0717*** (0.0056)		
Standardised labour history SES				0.0220*** (0.0034)	
Standardised composite SES					0.0340*** (0.0047)
Observations	1,394,320	1,394,321	1,394,321	1,394,321	1,394,321
R ²	0.94780	0.94781	0.94789	0.94782	0.94783
Within R ²	0.63208	0.63215	0.63272	0.63224	0.63230
Housing characteristics (50)	✓	✓	✓	✓	✓
PC4 fixed effects	✓	✓	✓	✓	✓
Year-Municipality fixed effects	✓	✓	✓	✓	✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator, using the log-arithmetic of the transaction price as the dependent variable. Standardised SES variables were included as controls to determine whether the base estimate would change. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

Table 5.11 shows that adding socioeconomic controls produces a few findings. Firstly, restricting the baseline specification to observations between 2014 and 2023 with SES data shifts the estimate from -0.81% to -1.26% , significant at the five per cent level. The additional SES control variables have a slightly lower estimate with -1.07% for the aggregated SES measure. The estimate thus becomes lower in absolute terms once socioeconomic variables at the PC4 level are taken into account. Furthermore, the direction of the control variable estimates is economically logical: an increase in a SES will increase a property's transactional value. This suggests that SES variables absorb variation that fine-grained fixed effects did not fully capture. Secondly, when observing the ATTs across columns two to five, the shift in turbine effect seems to be rather consistent, no matter which SES variable is used. The financial, educational, labour history, and overall aggregate scores all yield nearly identical treatment estimates when included alone. No matter which SES metric

is included, the signs and estimated effects all point in the same direction.¹⁰ However, it does not highlight whether turbine effects differ based on socioeconomic factors.¹¹

Urban-SES Heterogeneity

In subsection 5.1.3, a heterogeneity analysis looked into whether turbine effects were more impactful in urbanised areas. To examine whether the urbanisation pattern reflects underlying differences in socioeconomic composition, Table 5.12 presents the urbanisation split with and without standardised aggregate SES control variable for the 2014-2023 subsample.¹²

Table 5.12

The onshore effects of wind turbines on housing prices grouped by urbanisation category, with and without the standardised aggregated SES as a control variable.

	log price			
	Urban		Non-urban	
	without SES controls (1)	with SES controls (2)	without SES controls (3)	with SES controls (4)
Wind turbine ≤ 2 km	-0.0198*** (0.0075)	-0.0182** (0.0076)	-0.0056 (0.0060)	-0.0064 (0.0060)
Standardised composite SES		0.0341*** (0.0048)		0.0350*** (0.0048)
Observations	1,331,479	1,331,479	1,336,673	1,336,673
R ²	0.94846	0.94849	0.94753	0.94755
Within R ²	0.63341	0.63359	0.63280	0.63299
Year-Municipality fixed effects	✓	✓	✓	✓
Housing characteristics (50)	✓	✓	✓	✓
PC6 fixed effects	✓	✓	✓	✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. It compares urban and non-urban properties and highlights whether using standardised, aggregated SES results yields a meaningful difference. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

After restricting the sample to observations with non-NA SES variables, the base estimate changes

¹⁰Some of the partial SES scores are highly correlated. The correlation between financial SES and labour history SES is 0.89, indicating they measure nearly identical constructs. The correlations between educational SES and financial SES and labour history SES are 0.50 and 0.36, respectively. See Figure B.6.

¹¹An analysis splitting the sample by SES quartile was also conducted to test for heterogeneity in turbine effects across socioeconomic score quartiles. The estimates were statistically insignificant across nearly all quartiles, likely due to limited identifying variation in most of them; splitting the sample by quartile leaves too few observations per quartile for reliable estimation. This means that it cannot be determined whether lower or higher SES scores lead to varying turbine effects. This analysis can be seen in Table C.8. A similar analysis was performed by splitting the sample in two, but this also did not yield significant results.

¹²For all specifications, observations without an aggregate SES measure were dropped to ensure that the results were not driven by a sample differences. Again, keep in mind that the CBS only documents data from neighbourhoods with sufficient households, so the following table does not reflect these observations, as those are also dropped.

from -1.02% to -1.98% , a substantial difference. By removing all observations between 2010 and 2013, later-treated cohorts receive more weight in estimating the effect of turbines, and these cohorts have taller turbines, yielding larger treatment effects.¹³ These differences are not observed for the non-urban specifications, as both have a similar, non-statistically significant, point estimate.

When controlling for the standardised aggregate SES variable, the urban ATT decreases from -1.98% to -1.82% , a shift of 16 basis points, or approximately 8% of the effect magnitude. While the SES variable in urban regions thus marginally lowered the turbine effect in absolute terms, the gap between the urban effects and the non-urban effects (which are not statistically significant) remains present. This indicates that the urban effects of turbines are not fully explained by the socioeconomic backgrounds of treated PC4 regions, but rather reflect characteristics of urban housing markets themselves.

5.2. Descriptive Offshore Analysis

In recent years, wind turbines have also been constructed at sea at a fast pace. Figure 5.2 highlights the locations of the wind parks along the Dutch coast and also shows how many offshore turbines became operational in what year. It further illustrates which coastal properties are ever-exposed to offshore turbines during the temporal range of the housing transaction data, which would be required for any inference. For this descriptive analysis, properties were assumed to be exposed to turbines if they are within 30 kilometres of any offshore turbine. This range was chosen because [Sullivan et al. \(2013\)](#) found that offshore turbines are noticeable to observers at approximately 29 kilometres, but remain visible with concentrated or extended viewing at distances beyond 40 kilometres. This likely remains a conservative range, as offshore turbines have become increasingly larger in recent years (see Figure 3.5). A subsample with a cut-off distance of 10 kilometres from the coastline was chosen because the relevant offshore externalities are expected to be highly localised through turbine visibility, considering that visibility is often already constrained by dunes and other landscape features. Using this cut-off distance will therefore focus on the areas in which these effects are most likely to matter.

Table 5.13 shows the descriptive statistics of the coastal communities, defined as properties within 10 kilometres from the coastline. It highlights the number of treated, ever treated, never treated, and relevant percentages. It should be noted that the definitions used in the descriptive analysis are based solely on distances to the nearest turbine, but offshore turbine exposure may be more nuanced in Dutch coastal settings.

Table 5.13

Descriptive statistics regarding coastal transactions with turbine exposure.

Sale year	Currently treated	Not currently treated	Ever treated	Never treated	Total	Percentage currently treated	Percentage ever treated
2010	4,496	20,052	12,705	11,843	24,548	18.3	51.8
2011	4,122	18,397	11,401	11,118	22,519	18.3	50.6
2012	4,221	18,214	10,991	11,444	22,435	18.8	49.0

Continued on next page.

¹³Figure 3.5 shows that the onshore turbines have gotten taller, on average, between 2010 and 2013. Although subsection 5.1.3 did not establish that larger turbines generate systematically larger price effects, it is important to note that the average turbine height in this period frequently falls in the middle-height category (50-150 metres). Additionally, the largest turbines constructed near the end of the sample period have very few post-treatment observations. Therefore, the findings by [Dröes & Koster \(2021\)](#), who documented that turbine heights negatively affected property prices, seem more trustworthy than the results in this thesis.

Sale year	Currently treated	Not currently treated	Ever treated	Never treated	Total	Percentage currently treated	Percentage ever treated
2013	4,165	18,246	10,872	11,539	22,411	18.6	48.5
2014	6,204	25,357	15,195	16,366	31,561	19.7	48.1
2015	8,548	29,216	18,405	19,359	37,764	22.6	48.7
2016	9,904	31,196	20,866	20,234	41,100	24.1	50.8
2017	9,315	29,223	19,448	19,090	38,538	24.2	50.5
2018	8,444	25,971	17,528	16,887	34,415	24.5	50.9
2019	8,921	26,369	17,993	17,297	35,290	25.3	51.0
2020	9,614	27,069	19,096	17,587	36,683	26.2	52.1
2021	7,659	21,414	15,690	13,383	29,073	26.3	54.0
2022	7,294	18,511	14,572	11,233	25,805	28.3	56.5
2023	12,212	15,388	15,544	12,056	27,600	44.2	56.3



To compare price trends in the two coastal groups while accounting for differences in location and timing, the logarithm of the transaction price per square metre is compared after removing the average price levels associated with 6-digit postal code (PC6) regions and the sale year. These effects are removed in an iterative process, alternating between demeaning within postcodes and within years until the residuals stabilise, meaning that the remaining residuals cannot be further explained by time and locational properties. This process occurs for each transaction year, with 2014 used as the baseline year, while still using the 10 kilometre coastal subset and the 30 kilometre exposure range of offshore turbines.

Figure 5.3 reports the differences between the coastal groups. The results are somewhat surprising as the differences are pronounced and not constant over time. The difference is positive early on, starting around +4%, then declines at a constant rate to -5% around 2016, only to recover from 2020 onward to positive differences. It should be noted that the declines begin well before the majority of offshore turbines were constructed, which means the declining differences occur when most units are unlikely to be exposed. Additionally, it is noteworthy that the effects do not stabilise at a lower level as more turbines are constructed; rather, the differences between ever-treated and control units increase in magnitude. This suggests that these results may be driven by external factors not accounted for in this analysis.

The observed pattern is not a result of the definition of coastal communities. Figure 5.3 reproduces the difference using coastal zones of 5 and 15 kilometres, next to the 10 kilometre baseline. The V-shape, magnitude and timing for the minimum and maximum differences are all similar. Furthermore, replacing the PC6 location effects with 4-digit postal code (PC4) also yields similar findings (See Figure D.2 in Appendix D), indicating that these findings are robust across various coastal definitions.

While subsection 3.3.3 showed that a lower aggregate socioeconomic score is associated with greater onshore turbine exposure, the opposite pattern holds in the offshore setting. Figure 5.4 reports the mean aggregate Socioeconomic Status (SES) value of ever-treated and never-treated properties over time. Treated coastal properties, those ever within 30 kilometres from an offshore turbine, score higher on average than control properties throughout the sample, and the gap between the two groups is persistent although it drops slightly over the years. One possible explanation is that households with higher aggregate SES values tend to choose to live closer to the sea, meaning



Restricted Coastal Sample and Offshore Sea Turbines



(a) The placement of offshore turbines and the coastal properties within a 30-kilometre range.

(b) The construction of offshore wind turbines over the years in the Netherlands.

Figure 5.2: The turbine placement and construction numbers over the years in the Netherlands.



(a) The differences in price levels between ever-treated and control units after adjusting for PC6 and year effects.

(b) The robustness of the price differences for alternative coastal distance definitions.

Figure 5.3: The price differences between ever-treated and control units and their robustness.

that they are more likely to be treated under a distance-based definition.

This compositional difference in socioeconomic scores helps explain the positive price differences between treated and control properties. If higher SES households are over-represented near the coast, and if higher SES is associated with stronger price growth, then the treated group may appreciate faster for reasons unrelated to offshore turbines. It does not, however, explain the dip in the differences around 2016. Additionally, the differences in Figure 5.3 increase after 2016, as many offshore turbines became operational, which might seem counterintuitive.




Figure 5.4: The aggregate socioeconomic score over the years for Dutch coastal communities.

The difference between treated and control units can also be examined relative to the timing of offshore construction. Figure 5.5 reports the same differences in the logarithm of the transaction price between treated and control properties by event time, where event time is measured in years relative to the earliest turbine within 30 kilometres becoming operational, where $t = -1$ is the reference year. Note that year and locational effects are still removed, like in Figure 5.3.

The pre-treatment leads do not lie flat around zero. The difference rises from about -2% at $t = -5$ to roughly $+1\%$ at $t = -2$, one year prior to the reference year. A parallel trends assumption is therefore, at first glance, not supported descriptively, as the two groups are already moving relative to one another before any treatment occurs.¹⁴

After the turbine becomes operational, the difference drops sharply to approximately -8% in year $t = 1$, and remains highly negative afterwards. At first glance, this would suggest that offshore turbines have large effects, although this is descriptive rather than causal. However, the lower panel of Figure 5.5 shows that the number of observations drops to fewer than 2,000, while leads consisted of up to 10,000 observations. This could imply that there is little support, diminishing the power of these findings. The reason for the drop in observations is that turbines were only recently

¹⁴Keep in mind that anticipation effects can occur as found in prior literature for onshore turbines, but cannot be tested as this turbine data is not available.

constructed, meaning there are only a few years' worth of transactions to determine any descriptive post-treatment effect.




Figure 5.5: The descriptive differences in prices for coastal communities between treated and control properties relative to their treatment year, where $t = -1$ is the reference year. It also highlights the total observations per event time.

It should be kept in mind that turbine adoption is staggered but occurs only in a few years because offshore parks are not constructed every year (see Figure 5.2). Event times can therefore be dominated by one or two specific offshore parks, such that the effects for a specific event time may reflect the effects of a particular wind farm rather than a more general response to offshore turbine exposure. Given the thin and unbalanced post-treatment support and the small number of cohorts, the conditions are not ideal for causal inference, and no firm conclusions should be drawn from the event-time pattern.

The extent of this limitation is shown in Figure 5.6. Among the PC6 areas with treated observations, 39.7% have no pre-treatment observations and 34.4% have no post-treatment observations. Additionally, 62.0% of all PC6 areas have at most two post-treatment years. The descriptive event time shown in Figure 5.5 is thus built on very little support, further highlighting that no hard conclusions should be drawn from these graphs.

As a final step, the global year effects with municipality-by-year effects, as seen in the main specification of the onshore effects, and recompute the treated-control difference. Figure 5.7 compares the two assumptions. Under global year-time effects, the difference has the distinct V-shape documented earlier, whereas with year-municipality time effects, the differences appear to be constantly zero across all years.

It is hypothesised that these differences could be due to two reasons. First, the differences observed under global year effects were to some extent location effects, such that the differences highlight




Figure 5.6: The bars report the share of PC6 areas with treated observations with no pre-treatment observations, no post-treatment observations, and often no more than 2 post-treatment years.




Figure 5.7: The differences between treated and control units after accounting for different time effects, comparing global versus municipal time effects.

local housing-market trajectories rather than genuine offshore-related price patterns. Second, since treatment is based on large distances for all coastal communities, the treated and control properties seldom share the same municipality-year group. Therefore, eliminating municipality-by-year effects also eliminates the variation that differentiates the two groups. Any offshore effects that might be present would be absorbed into regional trends and would no longer be detectable.

The two readings are not mutually exclusive, but the overlap suggests that the second reading is dominant. Only 8.8% of municipality-year cells include both groups, covering 8.2% of all observations. This near-zero difference may be due to absorption rather than indicating a precisely estimated statistically insignificant effect. Consequently, there is no model that simultaneously accounts for regional trends and maintains a treated-control contrast when using distance-based treatment definitions of turbine exposure, because the same time control eliminates both confounding factors and the comparison group. Given this and the limited, unbalanced event-time data, offshore results are presented descriptively.

Discussion

This chapter discusses the results and places them within the broader literature. It explains why the estimated effects are less pronounced than in previous Dutch studies, explores the policy implications of the heterogeneity results, considers the challenge of balancing overall costs with the fair distribution of wind energy externalities, and puts the estimated effects in the context of the Dutch housing market. It aims to translate empirical results into findings that can improve the public debate on wind turbine siting strategies.

6.1. Main Findings in Context


The baseline analysis, estimating turbine effects with a staggered Difference-in-Differences (DiD) design, finds a small negative effect of approximately -0.8% when using 6-digit postal code (PC6) and year-municipality fixed effects, marginally significant at the 10% level. It is important to note, however, that this estimate is sensitive to specification choices. Under 4-digit postal code (PC4) fixed effects, the point estimate falls to near zero and loses significance at any conventional level. Additionally, under alternative control group definitions, the effect similarly shifts towards zero, sometimes even changing signs, and remains statistically insignificant. Taken together, these specifications suggest that the -0.8% effect is small, fragile, and not robust to the choice of control groups, and is statistically indistinguishable from zero under most alternative specifications. However, the baseline effect is robust to controlling for neighbourhood socioeconomic composition. Adding the standardised aggregate Socioeconomic Status (SES) shifts the estimate only marginally from -1.27% to -1.07% , and no single SES sub-score drives this change. This indicates that the estimated turbine effect is not fully attributable to socioeconomic differences between treated and control areas, but persists even after conditioning on the socioeconomic composition of PC4 areas.

The estimated effect is substantially smaller than the -1.4% to -2.7% range reported in the prior Dutch literature by [Dröes & Koster \(2016\)](#); [Dröes & Koster \(2021\)](#) and [Eichholtz et al. \(2023\)](#). However, the reported results are consistent with the broader international literature, which generally finds small and localised wind turbine effects ([Schütt, 2024](#)). The localised nature of the effect is supported directly by the findings in chapter 5: the estimated effect is only statistically significant with the 2 kilometre treatment definition, becomes insignificant if it is increased to 2.5 kilometres, and it is implied that effects within one kilometre are larger (-2.92%) than between one and two kilometres (-0.67%), although the differences between these groups are not statistically significant.

The causal interpretation of the baseline is limited by pre-trends. The event study in Figure 5.1 reveals notable price differences between treated and control properties from 7 to 4 years before treatment, rather than the flat pre-trend suggested by a strict parallel-trends assumption. As discussed in subsection 5.1.1, these pre-trends are most likely due to cohort composition, in which distant pre-treatment event times are mainly observed in later-treated cohorts that systematically

differ from the average treated unit. Although this explanation aligns better with the data than the notion that anticipation effects occur so far in advance, the presence of imperfect pre-trends weakens any causal claims. Excluding the event times that show pre-trends slightly shifts the estimate from -0.81% to -0.92% , suggesting that the pre-trends bias the baseline toward zero, as pre-trends already displayed a negative price level. While the causal interpretation should be approached with caution, the pre-trends do not appear to drive the estimated effect.

The economic magnitude of the estimated effect is modest. For a property at the national average transaction price, a -0.8% effect corresponds to a price reduction of roughly 3,880 euros, using the national average price of 485,000 euros in the first quarter of 2026 (NVM, 2026). Furthermore, Figure 6.1 presents the aggregated effects across the affected housing stock and the effects of single turbines per year. The total loss between 2010 and 2023 amounts to 315 million euros. While this number may appear large, it is spread over 14 years and across all housing transactions affected by turbines, and represents only a small percentage of the total housing market. Expressed per turbine, each turbine resulted in approximately 16,000 euros in lower transaction costs annually, which is modest compared to the turbines' costs.



(a) The aggregate transaction value loss due to onshore turbines over time.

(b) The transaction value loss measured per onshore turbine over time.

Figure 6.1: Time losses caused by onshore turbines, with price effects of -0.8% , graphed over time.

The estimated magnitude is relevant to the Dutch planning compensation policy, as the impact of onshore turbines on housing prices is limited. Under the 'Omgevingswet', property owners are entitled to compensation if the depreciation exceeds the 'societal risk'. However, this societal risk under the Omgevingswet is four per cent, and Minister Hermans of the previous cabinet found this threshold sufficient (Gastel, 2025). Throughout this thesis, none of the estimated effects across the baseline and all heterogeneity specifications exceeds the threshold. This means that the price effects documented in chapter 5 would rarely, if ever, warrant any compensation, even if the price depreciation is economically relevant and statistically significant, not because the Omgevingswet is poorly designed, but because the measured effect is genuinely small.

6.2. Differences From Prior Dutch Studies

The previously mentioned estimated effects are consistently smaller than those reported in the prior Dutch literature. There are four plausible reasons for these differences, with the sample period likely the most impactful.

The first factor regards the limited temporal range of the housing transaction data. Earlier studies used transaction data from the same source, dating back to 1985, capturing the early adoption of turbines in the Netherlands during the 1990s and early 2000s. The data used in this thesis begins only in 2010, by which point much of this adoption had already occurred. This restricts the identifying variation required for causal inference: in many areas with treated observations, the sample contains

few pre-treatment observations, few post-treatment observations, or both. As shown in chapter 3, the proportion of treated properties rises only slightly throughout the sample period. This suggests that the initial adoption phase, which offered considerable variation in prior Dutch studies, largely falls outside the temporal range of the currently available data (see Figure 3.9). This limitation is, however, structural and thus cannot be resolved through specification choices.

A second factor relates to the data itself. Although the data source is excellent, personal communication with Dröes and Koster (May 2026) revealed that the Nederlandse Vereniging van Makelaars (NVM) data for export has become more restricted in recent years.¹ A first-order estimate indicates that the number of transactions in the dataset is not restricted, taking into account that NVM realtors hold roughly a 70% market share. However, it is less clear whether the control variables have been maintained at the same level of detail as in prior publications. If the granularity of housing characteristics has decreased, directly comparing current estimates with prior estimates becomes more challenging, since the granularity of the control variables can influence the estimated treatment effect of turbines.

A third factor is the estimator used to determine effects. This thesis uses an estimator developed by Sun & Abraham (2021), whereas the prior Dutch literature relies on conventional Two-Way Fixed Effects (TWFE) Difference-in-Differences (DiD) specifications (Dröes & Koster, 2016; Dröes & Koster, 2021; Eichholtz et al., 2023). Methodologically speaking, as reported in subsection 2.2.4, the TWFE estimates can be biased under staggered adoption due to forbidden comparisons between treated and control units, a concern addressed by the Sun & Abraham (2021) estimator. section C.1 reports a TWFE replication of the baseline specification, and both estimators yield small negative effects that are marginally significant or insignificant depending on the fixed effects levels, indicating that the results are similar once strict fixed effects are applied. This indicates that the choice of estimator is not the primary driver of the divergence from the literature. The remaining differences may reflect TWFE bias due to staggered adoption or the different aggregation method used by the Sun & Abraham (2021) estimator, which estimates effects per event time and can then be aggregated to an Average Treatment Effect of the Treated (ATT).

A fourth factor concerns how heterogeneity analyses are conducted. Prior literature typically examines heterogeneity through interaction terms with TWFE DiD specifications. The approach taken in subsection 5.1.3 estimates heterogeneity using subsamples, defines treated groups by time-invariant characteristics measured at the initial exposure, and compares each against the full control group. This approach yields clean treatment effects per subgroup that are consistent with the staggered-adoption structure of the Sun & Abraham (2021) estimator. The trade-off, however, is reduced statistical power: splitting the sample into subgroups leaves less identifying variation per subsample, making any effects more difficult to detect. The subgroup estimates are therefore more easily interpretable than the pooled interaction coefficients of the TWFE specifications, where the estimated effect depends on how the interaction is defined, but they may also be less precise.

These factors thus suggest that the smaller estimates reported in chapter 5 reflect the characteristics of recent Dutch data and the methodological choices regarding staggered turbine adoption, rather than contradictory findings relative to prior literature. The sample period is most likely the most important factor, as no alternative design choice can recover identifying variation that falls outside the temporal range of the housing transaction data.

¹These are the authors who first reported the effects of wind turbines on property prices in the Dutch context in 2016, and reported a subsequent analysis in 2021 (Dröes & Koster, 2016; Dröes & Koster, 2021).

6.3. Heterogeneity and Policy Implications

While the baseline estimates an average treatment effect, heterogeneity analyses reveal that this average masks different effects across turbine characteristics and contexts. These heterogeneity analyses tell a different story and provide insights into the trade-offs politicians face when deciding where and how turbines should be sited.

The clearest and most policy-relevant finding concerns the number of turbines. Properties first exposed to a single turbine show an effect of -1.72% , while those first exposed to two or more turbines show a statistically insignificant effect of $+0.49\%$, and the interaction test confirms the difference between these two groups is statistically significant at the 1% level. Therefore, additional turbines do not generate additional price depreciations. If anything, evidence points toward positive price effects. This result is counterintuitive, yet several explanations are possible for this insignificant effect. First, multi-turbine clusters may be sited at greater average distances from properties, thereby decreasing the estimated effects. Second, multi-turbine areas may offer little identifying variation in this data: although roughly 31% of treated observations were initially treated with 2+ turbines, many may have been exposed before the sample period begins, reducing the number of pre-treatment observations to detect effects under strict fixed effects. Finally, it could be possible that the clusters may also consist of smaller turbines that likely generate smaller effects, though these hypotheses were not tested. The absence of a detectable multi-turbine effect, combined with the finding documented by [Dröes & Koster \(2021\)](#) that additional turbines do not generate marginal statistically significant effects, supports that turbines should be clustered to minimise effects. This would concentrate turbine externalities rather than exposing more properties to potentially higher turbine effects.

The turbine distance results are suggestive but not formally established. The subgroup estimates for initial distances up to one kilometre and within one to two kilometres suggest, though they do not confirm, a larger effect closer to turbines, echoing the localised finding of [Dröes & Koster \(2021\)](#). However, the interaction test indicates that the differences between these subgroups are not statistically significant, meaning that an 'urban effect' is not established. It can be established that the effects are spatially confined; Table 5.5 shows that no significant turbine effects are observed when the treatment distance is increased from 2 to 2.5 kilometres. An implication is, therefore, that turbine effects are spatially localised within two kilometres and do not show a statistically significant decay within those two kilometres. To minimise turbine effects, these turbines should be sited far from residential properties, ideally more than 2 kilometres away.

A similar interpretation should be used when determining the effects of urbanisation. The subgroups show that turbine effects are significant in urban areas (-1.02%) but insignificant in non-urbanised areas (-0.51%), and that these differences persist after controlling for the socioeconomic composition of 4-digit postal code (PC4) areas. The interaction test, however, does not establish that the urban and non-urban effects are statistically distinguishable. The honest reading is therefore that an effect is detectable in urban areas but not in non-urban areas, while the difference between the two cannot be formally confirmed in this data. While urban effects are therefore not proven, the higher density of affected properties per square kilometre in urban areas remains a consideration regardless of whether urban effects are statistically distinguishable. The urban-rural gap is therefore not explained by socioeconomic differences between treated areas, although the underlying reason why urban regions show larger effects cannot be determined from this data. Considering the higher density of affected properties per square kilometre in urban areas, these observations suggest placing turbines away from densely populated regions, where the impact on individual properties and the total number of affected properties is greater.

The height analysis provides the least definitive results. Only the medium-height category (50-150

metres) offers a credible estimate of -1.09% . The small-turbine category shows no significant effect, likely due to the lack of pre-treatment observations since these turbines were mostly built before the sample period. The largest-turbine category yields estimates that change sign across fixed-effects specifications, suggesting that the data cannot reliably estimate an effect for this group, likely due to the lack of post-treatment observations. The interaction test indicates that both the smallest and largest categories differ from the medium-height reference group, but given the unreliability of those two estimates, this should be regarded only as weak identification rather than as evidence of genuine heterogeneity across heights. Unlike Dröes & Koster (2021), it therefore cannot be concluded that taller turbines generate larger effects. Do note that these different results are likely due to the data's temporal range. Their data contains more identifying variation on turbine height needed to estimate height effects, and the estimator by Sun & Abraham (2021) requires more observations for larger turbines, as it estimates an effect for every event time, making it more difficult to estimate effects for the largest turbines.

Overall, the heterogeneity analyses confirm one strong siting principle and highlight some weaker factors. The key finding is that adding more turbines is not associated with more negative effects, indicating that clustering turbines localises fewer capitalised externalities on fewer properties. Other observed patterns, such as effects being limited to about two kilometres and more pronounced in urban areas than in rural ones, suggest avoiding turbine placement near residential or densely populated zones. However, these subgroup differences are not statistically significant and should therefore be interpreted with caution. Meanwhile, turbine height does not yield any practical recommendations due to data limitations. These recommendations reduce overall property value losses, but they do not consider how these losses are distributed across communities.

6.4. Distributive Implications

The heterogeneity findings highlight a key siting consideration with implications for property distribution. The most definitive result shows that clustering turbines does not further increase localised price effects, meaning that concentrating turbines affects fewer properties overall with similar price effects. The weaker patterns cautiously suggest that effects are constant within two kilometres of a turbine and are observable mainly in urban rather than non-urban areas. Overall, a strategy to reduce total property devaluation would cluster turbines in less densely populated zones. However, this approach focuses only on minimising aggregate property losses and does not address the uneven distribution of those losses across the Netherlands.

This distinction matters because subsection 3.3.3 and prior literature document that onshore turbines in the Netherlands are already disproportionately located in neighbourhoods with lower prices, which are associated with lower Socioeconomic Status (SES). A strategy that concentrates further wind development in rural areas would, in practice, reinforce this pattern by imposing additional burdens on communities already disproportionately exposed to turbines. From the perspective of distributive justice, which concerns the equitable allocation of environmental benefits and ills across groups (Jenkins et al., 2016; Vallor et al., 2022), this element of fairness is not captured by aggregate efficiency. Note that this explicitly concerns exposure to the turbines themselves, not the price effects experienced. After all, that is another question that could not be answered, as mentioned in section 5.1.4.

The tension is similar to a known limitation of cost-benefit reasoning, as cost-benefit analyses are also criticised for interpersonal aggregation (Hansson, 2010), essentially ignoring distributive aspects. Because such analyses aggregate monetised losses, they can favour outcomes that minimise the total costs while concentrating burdens on specific populations, particularly where those populations have lower property values and therefore incur smaller absolute losses. A siting strategy optimised purely to minimise overall costs would systematically direct turbine clusters toward

lower-value areas, associated with lower SES. Whether such a siting strategy is acceptable depends on how aggregate costs and benefits are weighed against the distribution of burdens, a normative question that this thesis cannot resolve through an empirical, data-driven approach.

It is important to consider the strength of this distributive concern alongside the magnitude and fragility of the estimated effects. The documented price effects are minor, and the subgroup differences that would justify a strict efficiency-based siting strategy are mostly not statistically significant. Thus, the distributive tension does exist but is modest in magnitude: siting decisions tend to impose small, capitalised losses on already disadvantaged communities, rather than large communities. While this does not eliminate concerns about fairness, it does place them in the proper context.

While the evidence shows where price effects are concentrated and who is currently most exposed, the decision on where to site turbines is ultimately a political one that balances the trade-off between minimising aggregate devaluations and distributive fairness. While this thesis does not offer the so-called best solution for navigating these trade-offs or for managing other forms of fairness, such as procedural justice or recognition justice, its contribution is to make the trade-off explicit. A siting strategy based on the heterogeneity results from subsection 5.1.3, applied without regard to distributional effects, would further increase existing inequality in the spatial allocation of wind energy, even if the average losses per property are small.

6.5. The Effects in Perspective

The earlier sections demonstrated that onshore wind turbines have a small and fragile effect on nearby housing prices, and examined how this effect and its associated burdens are distributed. This part contextualises these results within the wider Dutch housing market and spatial planning, exploring their policy implications more broadly.

The estimated impact should be considered in the context of the overall Dutch housing market. Between 2010 and 2023, Dutch house prices rose dramatically as visible in Figure 6.2, with the national median increasing by well over 80% across the period.² Considering year-on-year changes of this scale, a localised, fragile price effect of about 0.8% has only limited economic relevance as a market-wide factor. Wind turbines are therefore not a meaningful driver of housing market dynamics in the Netherlands. The main pressures on the Dutch housing market lie elsewhere: a persistent housing shortage of 396,000 properties, approximately 4.8% of the entire housing stock, with construction facing problems to reduce this number: regulations causing permitting constraints, nitrogen-related restrictions on construction, a shortage of building locations, population growth combined with changing family compositions, high costs associated with construction, and labour shortages in the construction sector ([Ministerie van Algemene Zaken, 2023](#); [Ministerie van Binnenlandse Zaken en Koninkrijksrelaties, 2024](#); [van den Beukel & Molleman, 2025](#)). These reasons contributed to very high property prices. At the same time, [De Nederlandsche Bank \(2020\)](#) shows that the maximum amount people can borrow for a mortgage contributes more towards high property prices than the property shortage itself. Relative to these problems in the Dutch housing market, the price impact of turbine siting is only marginal.

One could hypothesise that a decline in property prices is not necessarily undesirable from a societal perspective. In a housing market with significant shortages and high prices, a slight decrease near turbines could help prospective buyers, while only slightly reducing the wealth of current homeowners. Therefore, the effect is a modest transfer from existing homeowners to future buyers rather than a welfare loss. While this doesn't eliminate the distributive concern raised earlier, since the

²Also see Figure 3.17 for the price per square metre, and Figure B.7 for the graph highlighting this trend over time including the 80% claim.




Figure 6.2: The year-on-year growth of housing prices across the Netherlands over time.

6

transfer occurs unevenly across communities, it contextualises the magnitude of the harm relative to its impact. Whether such a transfer is ultimately desirable is yet another political choice. However, as the housing market has recently started to 'cool down', there is a risk of underwater mortgages, in which homeowners owe more on their mortgages than their properties' current market values (De Nederlandsche Bank, 2026). Moreover, this could slow the broader economy, as aggregate demand falls when mortgages are upside down, consistent with Keynesian macroeconomic theory. Although the effects are only modest, such cascading effects should not be overlooked.

Another factor to consider is that the smaller effect observed here, compared to earlier Dutch studies, may not be solely due to the limited identifying variation in recent data. It is also possible that societal attitudes toward wind turbines in general have evolved as they have become a more familiar feature of the Dutch landscape, potentially weakening the price effects. However, this explanation appears statistically unlikely, as research by Dröes & Koster (2021) found that the willingness to pay for living near turbines remained steady throughout the study period, suggesting perceptions haven't shifted significantly. Moreover, Dröes & Koster (2019) reports that effects are generally stable across different time periods, except between 2000 and 2009, further supporting the idea that perceptions have not evolved. The smaller effect is therefore most plausibly a combination of reduced identifying variation and the more granular fixed-effects structure and socioeconomic controls applied here, with changing attitudes a possible but unconfirmed contributor.

Housing prices were selected as the primary outcome among various externalities. While turbines influence landscapes, pristine nature, and bird migrations – many of which are challenging to quantify or compare – housing prices offer a clear, observable indicator through market transactions. They summarise local disamenities, including noise, visual impacts, and shadow flicker, into a single financial measure. Additionally, housing prices are directly linked to Dutch policy via compensation schemes. Studying them thus provides concrete, policy-relevant insights into the local costs associated with turbines. Nevertheless, the main finding that this effect is small suggests that if the housing-price externality remains modest, debates over turbine siting might focus more on less quantifiable externalities, such as landscape and ecological impacts, or on safety and health concerns, than on the housing market. For instance, Diederik Gommers³ mentioned concerns regarding

³Diederik Gommers, a medically trained expert in intensive care, became nationally renowned during the COVID-19

sleeping disorders due to low-frequency sound and its vibrations (Rensen, 2025). His argument is partly based on research by Godono et al. (2023), who found that people exposed to higher power levels are associated with more sleep disorders. At the same time, these effects are estimated to decrease with distance. He further advocates placing turbines at a safe distance of over 2 kilometres, at which point fewer than 25% of the population will experience sleep disorders. Overall, these points indicate that the policy discussion on turbine siting shouldn't focus only on housing-price impacts. If the direct economic externality is small, the key issues might lie in less easily measured areas, such as landscape, ecological importance, and health and safety, where societal significance could exceed what housing prices alone indicate.

6.6. General Limitations

Several limitations have been mentioned thus far in the document, some specific to particular analyses and others applying across the board. These are all listed below.

The most consequential limitation is the data's temporal range. The 2010-2023 window begins after the main phase of Dutch turbine adoption, which limits the available variation for identification and makes it more difficult to estimate turbine effects. As already mentioned, this limitation is structural rather than methodological and therefore cannot be solved by alternative design choices. In addition, the baseline effect is sensitive to the choice of control group, attenuating towards zero and switching sign under restricted control-group specifications. This reinforces the view that the estimate is fragile and statistically indistinguishable from zero under most alternatives, rather than indicating a robust effect.

A related data limitation concerns the Nederlandse Vereniging van Makelaars (NVM) export data. As noted, the granularity of available control variables may allegedly have declined in recent years, and the absence of turbine announcement dates prevents a direct test of anticipation effects. The pre-trends observed in the event study shown in Figure 5.1 cannot be conclusively linked to changes in cohort composition or actual anticipation effects. Additionally, the NVM data cover approximately 70% of all existing-home sales in the Netherlands. Contact with the NVM indicated that this coverage is higher in the Randstad, the north, and the centre of the country, and somewhat lower in Zeeland and Limburg, with no apparent bias by property type or price. The transaction sample is therefore broadly representative of the Netherlands. Still, it is important to note that the findings in chapter 5 are based not on all housing transactions but on a subsample. This is particularly relevant to the descriptive offshore analysis, as the relatively small sample size makes it more difficult to detect effects.

A further limitation arises from the assumption that a treated unit stays treated. Observations are excluded from the sample when a property reverts from treated to untreated status due to turbine decommissioning. Such properties remain excluded even if they are subsequently re-treated. This enforces a treatment structure consistent with the staggered-adoption estimator but also excludes transactions from regions where turbines are replaced over time. These regions are not assumed to be independent and identically distributed (i.i.d.) across the Netherlands, so removing about 30,000 observations may mean that the estimated Average Treatment Effect of the Treated (ATT) is not representative for all properties exposed to turbines.

Onshore analyses assess exposure based on the closest operational turbine at the time of sale. When a property is close to turbines of varying heights, isolating the impact of the nearest turbine from other nearby turbines is challenging. For instance, a 50-metre turbine next to a 150-metre turbine does not provide a clear measure of the smaller turbine's effect. The baseline estimate is extended with various heterogeneity analyses in subsection 5.1.3, but it should still be regarded as an average

pandemic through interviews and media appearances on how to slow the virus's spread.

effect that includes all types of turbine exposures.

The socioeconomic analysis is constrained in two ways. The Socioeconomic Status (SES) data are available only from 2014 onward at a yearly level and had to be aggregated from the neighbourhood level to the 4-digit postal code (PC4) level, rather than being available for each property. Furthermore, the data are recorded only at the neighbourhood level if there are sufficient households, meaning that very small neighbourhoods are not included in subsection 5.1.4. The high correlation among SES sub-scores further limits the identification of the exact dimension driving the observed associations, and the quartile-split analysis lacks sufficient identifying variation within strata to detect heterogeneity reliably. Secondly, socioeconomic data on household composition are not included in the analysis, so analyses examining heterogeneity by household size, migration status, and other factors cannot be conducted.

The offshore analysis is descriptive rather than causal. The limited number of offshore parks, their recent clustered construction, and the absence of a visibility measure preclude the use of the same identification strategy applied to onshore turbines, as discussed in section 4.6. The section, therefore, describes why estimating the effects of offshore turbines is difficult, but does not present a causal estimate of those effects in the Dutch context.

6

The heterogeneity analyses estimate effects for treatment trajectories defined at the initial exposure rather than for continuous exposure to a fixed characteristic. As explained in section 4.4, properties may experience changes in turbine exposure after initial treatment, so subgroups with relatively long lags should be interpreted with this in mind.

Finally, the hedonic approach captures the total capitalised effect of turbine proximity but cannot conclude what elements actually cause these effects. The analyses in chapter 5 do not determine whether price declines are driven by visibility, noise, shadow flicker, or other turbine-associated mechanisms, because distance-based exposure was used and these elements were not observed in the data.

Conclusion

This chapter concludes by answering each research question and synthesising the findings, while highlighting the normative trade-offs in turbine siting strategies. It reflects on the research process and the methodological decisions made, and it suggests directions for future research.

7.1. Answers to Research Questions

This thesis set out to answer the question of how onshore and offshore wind turbines affect housing prices in the Netherlands, and how these effects vary across space and socioeconomic context. The findings are synthesised here in relation to the research questions.

The main finding is that onshore wind turbines slightly reduce nearby housing prices by about 0.8% within a two-kilometre radius under the preferred specification. This effect is highly localised, disappearing beyond two kilometres, and remains when controlling for Socioeconomic Status (SES) at the 4-digit postal code (PC4) level.¹ Additionally, the effect is fragile: it decreases toward zero under most alternative specifications that use coarser location fixed effects and stricter control group definitions, becoming insignificant and sometimes even changing sign. The causal interpretation is further limited by imperfect pre-trends, but robustness checks suggest these do not drive the estimated effect. This suggests that, under current Dutch conditions, the impact of onshore turbines on housing markets is much smaller than previous Dutch research indicated, likely because recent data exclude the initial adoption phase and offer less identifying variation. Overall, the estimated losses total approximately 315 million euros between 2010 and 2023. This represents only a small portion of the Dutch housing market and remains well below the per-property compensation threshold.

For offshore wind, the available data do not support a causal estimate of effects on coastal housing prices. The small number of offshore parks, their recent, clustered construction, their confinement to relatively few coastal communities, and the absence of a visibility measure collectively preclude an identification strategy from determining any causal effect of offshore turbines. A descriptive analysis shows that price differences between coastal properties within and outside 30 kilometres of offshore parks already exist before turbines become operational, and that the apparent post-construction divergence is based on relatively few observations, making it difficult to draw any conclusion from these analyses. Rather than force an estimate the data cannot support, the analysis, therefore, documents why offshore identification is difficult and what future work would require, rather than providing a causal estimate.

The heterogeneity analysis yields one robust result and several weaker patterns. Formal testing shows that the effect does not increase with more turbines; additional turbines do not lead to larger price declines. This indicates that turbine clustering only exposes fewer properties to price effects. Although the effects appear concentrated near turbines and in urban areas, this has not been statistically confirmed. This distinction is important for policymaking, as the clustering effect has

¹On the matched sample, the effect attenuates from -1.26% to -1.07% after adding aggregate SES controls.

strong support, whereas the evidence for placing turbines away from residents and cities appears less convincing than subgroup estimates suggest.

These findings all address the main question: onshore turbines cause small, localised, and fragile price effects on nearby properties; offshore impacts remain uncertain; and among the various factors that might influence the magnitude of the effects, only the number of turbines is robustly established. The broader implications are twofold. First, since onshore turbines are often installed in lower-value neighbourhoods associated with lower SES, these minor impacts are unevenly distributed, potentially worsening existing inequalities if turbine siting strategy aims to minimise costs. Ultimately, the placement of turbines remains a political trade-off, balancing overall property devaluations with the need to ensure an equitable distribution of burdens. This trade-off is highlighted, but not resolved. Second, the modest effect size suggests that the housing market might not be the primary outcome of interest, inviting attention to externalities such as ecological and health impacts that are harder to measure.

7.2. Reflection

Having just completed an elective course called *statistical methods for causal inference and prediction* in the economics and finance package, I was interested in applying this toolkit, and analysing the effect of wind turbines on housing prices seemed like an appealing application, especially once it became clear that I could access excellent data sources. Looking back, I started somewhat naively, assuming that once the specifications were defined, the subsequent analysis would be relatively straightforward. However, I found the process more challenging than expected throughout the entire thesis period; many of them are listed below.

The first and most persistent surprise throughout the project was the difficulty and time required to prepare the data. I had assumed that most of my time would be spent modelling in R, but merging the WindStats turbine dataset with the Nederlandse Vereniging van Makelaars (NVM) transaction data took weeks, which felt more like a data science project than the expected econometric analysis. This process required multiple revisions. The initial pipeline was designed around a two-way fixed effects study structure, following the approach of earlier literature, but the [Sun & Abraham \(2021\)](#) estimator required different inputs given its cohort-based structure. Socioeconomic variables, which became relevant only halfway through the thesis, were incorporated at later stages. Some detailed data sources, such as neighbourhood household composition ('Kerncijfers Wijken en Buurten' to be exact), proved unusable because they could not be reliably made time-varying due to changes in how the data were collected over the years. Even after attempts to merge them correctly, they showed unexplainable fluctuations over the years and were eventually dropped altogether. The available Socioeconomic Status (SES) data, however, were only accessible from 2014, which reduced the already limited sample by four years. Each transition between R and Python for these iterations took more time than expected and planned.

The first simulations had genuinely confusing results. Using essentially the same data sources as [Dröes & Koster \(2016\)](#) and attempting to recreate their results, I found no significant effects despite using near-identical specifications. As this was my first time conducting such analyses, my initial thought was that I was at fault. It took a while to consider that the data might not allow identification because its shorter temporal range may have lacked the identifying variation required. This insight led me to examine more closely what was being identified at the local level, a move that proved quite valuable. Using a modern estimator and municipality-by-year fixed effects, marginally significant effects emerged, though they were not robust across control-group definitions. Looking back, I recognise that I was initially searching for statistical significance rather than being prepared to report the absence of an effect, mainly because all the previous studies found economically relevant and statistically significant effects. It took a while to settle on the idea that a statistically insignifi-

cant result can be insightful, rather than aiming only for statistically significant results.

During the project, several ambitions ultimately proved infeasible in the Dutch wind turbine setting. For instance, it was intended to use the [Callaway & Sant'Anna \(2021\)](#) estimator for the heterogeneity analyses, which assumes conditional parallel trends and works under staggered adoption. However, it was not feasible because the pre-treatment matrix became singular when more than one covariate was included, namely the living area. With variables such as housing type, maintenance status, and locational characteristics excluded, conditional parallel trends were implausible, as indicated by the *did* package itself.

A more ambitious idea was to replace the binary treatment with a continuous exposure index that combines turbine height and distance and accounts for multiple nearby turbines, such that all nearby turbines contribute to the overall turbine exposure.² However, there were two problems. First, the [Sun & Abraham \(2021\)](#) estimator cannot accommodate continuous treatment, as a property is either treated or not rather than partially exposed. A Difference-in-Differences (DiD) estimator that allows continuous treatment and staggered adoption could address this, but its implementation was not feasible in the time remaining, leaving Two-Way Fixed Effects (TWFE) specifications as the only option. Second, the estimates were positive, suggesting that increased turbine exposure raised prices. This is economically implausible and contradicts the baseline and heterogeneity results. The index also became vanishingly small for higher distance exponents, as the distance always exceeded the height. Additionally, the interpretation was more difficult, since a coefficient could no longer be read as the effect of a turbine within a certain range, but was now an abstract combination combining height, distance, and the number of turbines. For these reasons, the idea was set aside and documented here rather than presented as a research question and result.

The offshore turbine analysis followed a similar trajectory. I had initially assumed that once onshore effects were established, offshore effects could be estimated in the same way. It became clear that identifying onshore effects was already difficult, and that identifying offshore effects was even more challenging. The distance-based exposure is far less plausible at sea, where only visibility aspects remain; only a few offshore turbines were ever constructed, and all of them recently; the analysis is confined to coastal communities with fewer observations, and a credible definition of treatment would require some detailed measure of visibility that the data currently lacks. The offshore analysis therefore became descriptive, documenting why estimation is difficult and what future work would need to address, rather than offering a causal estimate.

Overall, these experiences showed me that the challenge in applied econometrics is usually not in the estimation method itself, but in whether the data can support the research question. The most valuable skill I developed was not running a particular model, but judging what a given dataset can and cannot identify, and reporting that honestly even when the answer is less clean than hoped.

7.3. Future Research

An important aspect for future research is to replicate this analysis using a repeated cross-section from the initial turbine adoption period. If a modern staggered-adoption estimator applied to such data recovers effects closer to those of earlier studies, this would confirm that the smaller estimates are due to the restricted temporal range rather than the estimator choice itself. If the estimates remain smaller, it would suggest that the choice of estimator, rather than the sample period, accounts

²The treatment dummy D_{it} would be replaced by E_{it} , the continuous exposure index, where $E_{it} = \sum_{k \in \tau_t} 1[d_{ik} \leq \bar{d}] \cdot \frac{H_k}{d_{ik}^p}$. In this specification, all operational turbines in set τ_t would contribute to the overall exposure if the turbine k was within the cut-off distance $\bar{d}$, and H_k is the tip height, d_{ik} is the distance, and p highlights the rate of decay with distance.

for part of the difference with the existing literature.

A second direction is future causal identification of offshore turbine effects, which this thesis could only address descriptively. Given the small number of offshore parks and treated properties, methods that construct synthetic counterfactuals, such as synthetic Difference-in-Differences (DiD), may help overcome the limited sample. Such work would also need to incorporate a measure of visibility rather than distance alone, since visibility is presumably the dominant externality at sea. The causal effects of offshore turbines are especially important, as in 2024, 40.5% of all wind energy was provided by offshore wind parks ([Centraal Bureau voor de Statistiek, 2025b](#)). This capacity continues to increase with the development of large offshore zones such as the Hollandse Kust (West) and IJmuiden Ver areas ([Rijksoverheid, n.d.-b](#)).

A third direction concerns what actually drives the property's depreciation. This thesis observes greater effects in urban areas that are not attenuated by aggregate Socioeconomic Status (SES) controls, but it does not identify the specific causes. The analyses in chapter 5 also cannot distinguish whether these price declines are due to visibility, noise, shadow flicker, or something else, since treatment was defined solely by distance. Future research could aim to separate these influences, perhaps using a viewshed analysis or noise modelling, to identify the actual drivers of these effects. This could guide targeted mitigation efforts, potentially reducing the disproportionate impacts observed in densely populated regions, leading to more equitable siting strategies.

A fourth extension would be to revisit the continuous exposure index, for instance by using the continuous treatment DiD estimator by [Callaway et al. \(2024\)](#), which was not implemented here due to time constraints.³ Such a design could aim to uncover the relationship between empirical definitions of turbine exposure and effects on housing prices, and may yield more interpretable estimates than the Two-Way Fixed Effects (TWFE) specification attempted in this thesis.

Finally, the socioeconomic analysis was constrained by the use of socioeconomic variables (*Kerncijfers wijken en buurten*), such as household composition, that could not be reliably made time-varying. Future work could draw on alternative time-varying socioeconomic data to examine dimensions of heterogeneity that this thesis could not, including household size and migration background. Additionally, future work could examine whether treatment effects are heterogeneous across groups with different socioeconomic backgrounds, a question this thesis could not address due to a lack of identifying variation.

³Due to time constraints, not a lot of research has been conducted regarding the assumptions of this estimator, for instance the *strong parallel trends assumption*, which should be studied carefully. Therefore, this estimator should be regarded only as a possible option, not the correct way forward.

References

- Brunner, E., Hoen, B., Rand, J., & Schwegman, D. (2024). Commercial wind turbines and residential home values: New evidence from the universe of land-based wind projects in the United States. *Energy Policy*, 185. <https://doi.org/10.1016/j.enpol.2023.113837>
- Callaway, B., Goodman-Bacon, A., & Sant'Anna, P. H. (2024, February). *Difference-in-differences with a Continuous Treatment* (tech. rep. No. w32117). National Bureau of Economic Research. Cambridge, MA. <https://doi.org/10.3386/w32117>
- Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230. <https://doi.org/10.1016/j.jeconom.2020.12.001>
- Centraal Bureau voor de Statistiek. (n.d.). Hoe interpreteer je de SES-WOA-scores en hoe zijn deze bepaald? [Last Modified: 11-11-2025T11:05:18]. Retrieved July 7, 2026, from <https://www.cbs.nl/nl-nl/faq/infoservice/hoe-interpreteer-je-de-ses-woa-scores-en-hoe-zijn-deze-bepaald->
- Centraal Bureau voor de Statistiek. (2024). CBS Open data StatLine. Retrieved June 20, 2026, from https://opendata.cbs.nl/portal.html?_la=nl&_catalog=CBS&tableId=85900NED&_theme=378
- Centraal Bureau voor de Statistiek. (2025a). 4. Windenergie [Last Modified: 2025-09-08T06:30:00+02:00]. Retrieved November 23, 2025, from <https://www.cbs.nl/nl-nl/longread/rapportages/2025/hernieuwbare-energie-in-nederland-2024/4-windenergie>
- Centraal Bureau voor de Statistiek. (2025b). StatLine - Hernieuwbare elektriciteit; productie en vermogen. Retrieved February 9, 2026, from <https://opendata.cbs.nl/#/CBS/nl/dataset/82610NED/table?ts=1633440075419>
- De Nederlandsche Bank. (2020). Huizenprijs hangt meer samen met financieringsruimte koper dan met woningtekort. Retrieved June 29, 2026, from <https://www.dnb.nl/algemeen-nieuws/nieuws-2020/huizenprijs-hangt-meer-samen-met-financieringsruimte-koper-dan-met-woningtekort/>
- De Nederlandsche Bank. (2026). De woningmarkt. Retrieved June 29, 2026, from <https://www.dnb.nl/actuele-economische-vraagstukken/woningmarkt/>
- Devine-Wright, P. (2005). Beyond NIMBYism: Towards an integrated framework for understanding public perceptions of wind energy. *Wind Energy*, 8(2), 125–139. <https://doi.org/10.1002/we.124>
- Dong, L., Gaur, V., & Lang, C. (2023). Property value impacts of onshore wind energy in New England: The importance of spatial heterogeneity and temporal dynamics. *Energy Policy*, 179. <https://doi.org/10.1016/j.enpol.2023.113643>
- Dong, L., & Lang, C. (2022). Do views of offshore wind energy detract? A hedonic price analysis of the Block Island wind farm in Rhode Island. *Energy Policy*, 167. <https://doi.org/10.1016/j.enpol.2022.113060>
- Dong, L., Lang, C., & Parent, J. (2024). Focusing the view: Improved methods for assessing viewshed impacts of onshore wind turbines. *Journal of Environmental Economics and Management*, 128. <https://doi.org/10.1016/j.jeem.2024.103068>
- Dröes, M., & Koster, H. (2016). Renewable energy and negative externalities: The effect of wind turbines on house prices. *Journal of Urban Economics*, 96, 121–141. <https://doi.org/10.1016/j.jue.2016.09.001>

- Dröes, M., & Koster, H. (2019). *Windturbines, zonneparken en woningprijzen* (tech. rep.). Rijksdienst voor Ondernemend Nederland. Retrieved June 29, 2026, from <https://www.rvo.nl/sites/default/files/2020/01/windturbines-zonneparken-en%20woningprijzen-2019.pdf>
- Dröes, M., & Koster, H. (2021). Wind turbines, solar farms, and house prices. *Energy Policy*, 155. <https://doi.org/10.1016/j.enpol.2021.112327>
- Eichholtz, P., Kok, N., Langen, M., & van Vulpen, D. (2023). Clean Electricity, Dirty Electricity: The Effect on Local House Prices. *Journal of Real Estate Finance and Economics*, 66(4), 743–777. <https://doi.org/10.1007/s11146-021-09878-6>
- European Commission. (2025). In focus: Wind energy powering the clean transition - European Commission. Retrieved December 4, 2025, from https://energy.ec.europa.eu/news/focus-wind-energy-powering-clean-transition-2025-03-17_en
- Gastel, E. v. (2025). Minister: Huidige compensatie burgers voor schade windmolens voldoende. Retrieved June 16, 2026, from <https://solarmagazine.nl/nieuws-zonne-energie/i41535/minister-huidige-compensatie-burgers-voor-schade-windmolens-voldoende>
- Godono, A., Ciocan, C., Clari, M., Mansour, I., Curoso, G., Franceschi, A., Carena, E., De Pasquale, V., Dimonte, V., Pira, E., Dallapiccola, B., Normanno, N., & Boffetta, P. (2023). Association between exposure to wind turbines and sleep disorders: A systematic review and meta-analysis. *International Journal of Hygiene and Environmental Health*, 254, 114273. <https://doi.org/10.1016/j.ijheh.2023.114273>
- Gong, Y. (2017). The Spatial Dimension of House Prices. *A+BE | Architecture and the Built Environment*, 7(4), 1–186. <https://doi.org/10.7480/abe.2017.4.1747>
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Themed Issue: Treatment Effect 1*, 225(2), 254–277. <https://doi.org/10.1016/j.jeconom.2021.03.014>
- Hansson, S. O. (2010). CostBenefit Analysis: Philosophical Issues. In Palgrave Macmillan (Ed.), *The New Palgrave Dictionary of Economics* (pp. 1–4). Palgrave Macmillan UK. https://doi.org/10.1057/978-1-349-95121-5_2918-1
- Heintzelman, M., & Tuttle, C. (2012). Values in the wind: A hedonic analysis of wind power facilities. *Land Economics*, 88(3), 571–588. <https://doi.org/10.3368/le.88.3.571>
- Hoen, B., Brunner, E., & Schwegman, D. (2025). Uplifting winds: The surprisingly positive community-wide impact of wind energy installations on property values. *Energy Research and Social Science*, 127. <https://doi.org/10.1016/j.erss.2025.104331>
- Hoen, B., Wiser, R., Cappers, P., Thayer, M., & Sethi, G. (2009). The Impact of Wind Power Projects on Residential Property Values in the United States: A Multi-Site Hedonic Analysis. *Journal of Real Estate Research*, 33. <https://doi.org/10.2172/978870>
- Informatiehuys Water. (2025). Dataset: KRW Oppervlaktewaterlichamen (INSPIRE geharmoniseerd). Retrieved February 17, 2026, from <https://www.pdok.nl/introductie/-/article/krw-oppervlaktewaterlichamen-inspire-geharmoniseerd->
- Jenkins, K., McCauley, D., Heffron, R., Stephan, H., & Rehner, R. (2016). Energy justice: A conceptual review. *Energy Research & Social Science*, 11, 174–182. <https://doi.org/10.1016/j.erss.2015.10.004>
- Jensen, C., Panduro, T., Lundhede, T., Nielsen, A., Dalsgaard, M., & Thorsen, B. (2018). The impact of on-shore and off-shore wind turbine farms on property prices. *Energy Policy*, 116, 50–59. <https://doi.org/10.1016/j.enpol.2018.01.046>
- Kadaster. (2025). Dataset: Zeegebieden (INSPIRE geharmoniseerd). Retrieved February 17, 2026, from <https://www.nationaalgeoregister.nl/geonetwork/srv/dut/catalog.search#/metadata/101304f0-2cf2-4ce0-aa23-cd7e4000fbba>
- Lang, C., Opaluch, J., & Sfinarolakis, G. (2014). The windy city: Property value impacts of wind turbines in an urban setting. *Energy Economics*, 44, 413–421. <https://doi.org/10.1016/j.eneco.2014.05.010>

- Mattmann, M., Logar, I., & Brouwer, R. (2016). Wind power externalities: A meta-analysis. *Ecological Economics*, 127, 23–36. <https://doi.org/10.1016/j.ecolecon.2016.04.005>
- McCarthy, I., & Balli, H. (2014). Windfarms and residential property values. *International Journal of Strategic Property Management*, 18(2), 116–124. <https://doi.org/10.3846/1648715X.2014.889770>
- Ministerie van Algemene Zaken. (2023). 900.000 nieuwe woningen om aan groeiende vraag te voldoen | Rijksoverheid.nl [Last Modified: 2023-07-24T09:08:32.299Z]. Retrieved June 29, 2026, from <https://www.rijksoverheid.nl/themas/bouwen-en-wonen/volkshuisvesting/nieuwe-woningen>
- Ministerie van Binnenlandse Zaken en Koninkrijksrelaties. (2024). Het statistisch woningtekort uitgelegd | Home | Volkshuisvesting Nederland [Last Modified: 2026-04-20T15:16:56.603Z]. Retrieved June 29, 2026, from <https://www.volkshuisvestingnederland.nl/onderwerpen/aanpak-woningnood/berekening-woningbouwopgave>
- NVM. (2024). *De Nederlandse markt voor recreatiewoningen* (tech. rep.). Retrieved March 23, 2026, from <https://www.nvm.nl/media/ql0ppi0t/2024-06-18-nvm-rapport-recreatiewoningmarkt-2023.pdf>
- NVM. (2026). Marktinformatie Wonen | NVM. Retrieved June 16, 2026, from <https://www.nvm.nl/wonen/marktinformatie-wonen/>
- Onakpoya, I. J., O'Sullivan, J., Thompson, M. J., & Heneghan, C. J. (2015). The effect of wind turbine noise on sleep and quality of life: A systematic review and meta-analysis of observational studies. *Environment International*, 82, 1–9. <https://doi.org/10.1016/j.envint.2015.04.014>
- Rensen, E. (2025). Hoe slecht is wonen naast windmolens eigenlijk? Diederik Gommers: Veel slaapstoornissen. Retrieved June 30, 2026, from <https://www.telegraaf.nl/binnenland/hoeslecht-is-wonen-naast-windmolens-eigenlijk-diederik-gommers-veel-slaapstoornissen/68032272.html>
- Rijksoverheid. (n.d.-a). Nadeelcompensatie: Wat is veranderd en overgangsrecht [Last Modified: 2024-11-25]. Retrieved February 18, 2026, from <https://iplo.nl/thema/ruimtelijke-ontwikkelingen/bijzondere-onderwerpen/nadeelcompensatie/veranderd-overgangsrecht/>
- Rijksoverheid. (n.d.-b). Waar staan en komen de windparken op zee? [Last Modified: 2025-08-11]. Retrieved December 4, 2025, from <https://windopzee.nl/onderwerpen/waar-staan-komen-windparken/>
- Rijksoverheid. (n.d.-c). Waarom windenergie op zee als energiebron? [Last Modified: 2025-01-16]. Retrieved December 4, 2025, from <https://windopzee.nl/onderwerpen/waarom-windenergie-zee/>
- Russell, A., Bingaman, S., & Garcia, H.-M. (2021). Threading a moving needle: The spatial dimensions characterizing US offshore wind policy drivers. *Energy Policy*, 157. <https://doi.org/10.1016/j.enpol.2021.112516>
- Schütt, M. (2024). Wind Turbines and Property Values: A Meta-Regression Analysis. *Environmental and Resource Economics*, 87(1), 1–43. <https://doi.org/10.1007/s10640-023-00809-y>
- Sims, S., Dent, P., & Oskrochi, G. (2008). Modelling the impact of wind farms on house prices in the UK. *International Journal of Strategic Property Management*, 12(4), 251–269. <https://doi.org/10.3846/1648-715X.2008.12.251-269>
- Skenteris, K., Mirasgedis, S., & Tourkolias, C. (2019). Implementing hedonic pricing models for valuing the visual impact of wind farms in Greece. *Economic Analysis and Policy*, 64, 248–258. <https://doi.org/10.1016/j.eap.2019.09.004>
- Sullivan, R. G., Kirchler, L. B., Cothren, J., & Winters, S. L. (2013). Research Articles: Offshore Wind Turbine Visibility and Visual Impact Threshold Distances. *Environmental Practice*, 15(1), 33–49. <https://doi.org/10.1017/S1466046612000464>

- Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Themed Issue: Treatment Effect 1*, 225(2), 175–199. <https://doi.org/10.1016/j.jeconom.2020.09.006>
- Sunak, Y., & Madlener, R. (2016). The impact of wind farm visibility on property values: A spatial difference-in-differences analysis. *Energy Economics*, 55, 79–91. <https://doi.org/10.1016/j.eneco.2015.12.025>
- Sunak, Y., & Madlener, R. (2017). The impact of wind farms on property values: A locally weighted hedonic pricing model. *Papers in Regional Science*, 96(2), 423–444. <https://doi.org/10.1111/pirs.12197>
- Torzewski, M. (2016). The impact of wind farms on the prices of nearby houses in Poland: A review and synthesis. *Real Estate Management and Valuation*, 24(2), 13–24. <https://doi.org/10.1515/remav-2016-0010>
- Vallor, Shannon, Raicu, I., & Green, B. (2022). *Ethics in Technology Practice*. Markkula Center for Applied Ethics at Santa Clara University. Retrieved February 9, 2026, from <https://www.scu.edu/ethics/>
- van Veelen, B., & Haggett, C. (2017). Uncommon Ground: The Role of Different Place Attachments in Explaining Community Renewable Energy Projects. *Sociologia Ruralis*, 57(S1), 533–554. <https://doi.org/10.1111/soru.12128>
- van den Beukel, L., & Molleman, S. (2025). *Bouw in beeld* (tech. rep.). Uitvoeringsinstituut Werknemersverzekeringen. Retrieved June 29, 2026, from <https://www.uwv.nl/assets-kai/files/521c9c1f-f4a3-42a1-a47c-d2b8c7cf0344/Bouw-in-beeld-oktober-2025.pdf>
- Vyn, R. (2018). Property value impacts of wind turbines and the influence of attitudes toward wind energy. *Land Economics*, 94(4), 496–516. <https://doi.org/10.3368/le.94.4.496>
- Vyn, R., & Mccullough, R. (2014). The Effects of wind turbines on property values in Ontario: Does public perception match empirical evidence? *Canadian Journal of Agricultural Economics*, 62(3), 365–392. <https://doi.org/10.1111/cjag.12030>
- Westlund, H., & Wilhelmsson, M. (2021). The socio-economic cost of wind turbines: A swedish case study. *Sustainability (Switzerland)*, 13(12). <https://doi.org/10.3390/su13126892>

AI Reflection

Several AI models were used throughout the thesis to improve the overall quality of the work. Codex (from OpenAI) and Perplexity Pro were used to assist with programming, correct coding strategies and logic, and debugging of the Python pipeline to merge all the Nederlandse Vereniging van Makelaars (NVM) housing transaction data, the WindStats wind turbine data, and SES data. All code was manually reviewed, tested, and validated before use. Perplexity Pro was also used to identify relevant sources that did not appear in the initial Scopus search. Grammarly Pro and Claude Pro (Anthropic) were used to improve the structure, clarity, and readability of individual sections. Do note that all the ideas presented, argumentation, and findings remain my own. Additionally, Claude Pro was used to actively challenge my own reasoning or disprove my hypotheses using economic reasoning.

Additional Descriptive and Socioeconomic Data Analysis

B.1. Socioeconomic Data

This appendix presents additional Socioeconomic Status (SES) diagnostics that complement the analysis shown in subsection 3.3.3. The figures depict the spatial distribution of financial and educational SES scores across the Netherlands, showing changes in SES scores from 2014 to 2023 and illustrating the spread and missing data at the 4-digit postal code (PC4) level. This highlights that (i) the SES scores vary significantly across regions and change unevenly over time, which highlights that SES should be time-variant, and (ii) that SES data is missing for the smallest neighbourhoods and observations without PC4 SES data, which is distributed all over the Netherlands but most frequently in the North. Additionally, a correlation matrix of all SES variables is provided, highlighting that some variables within the same category measure nearly identical constructs.




Figure B.1: The financial SES scores distributed across the Netherlands.

B




Figure B.2: The educational SES scores distributed across the Netherlands.




Figure B.3: The change in the partial and the aggregate SES scores between 2014 and 2023 shown across the Netherlands.

B




Figure B.4: The change in additional SES scores, related to finance, education, and the aggregate score, between 2014 and 2023 across the Netherlands.




Figure B.5: The spread of SES variables highlighting heterogeneous neighbourhoods and the PC4 areas that do not contain SES data.

B




Figure B.6: The correlation matrix of time-varying SES in the Netherlands between 2014 and 2023.


**B**

Figure B.7: The median transaction price growth over the years in the Netherlands with base year 2010.

Additional Specifications

This appendix reports additional regression specifications that supplement the main results. They serve as robustness checks and alternative specifications, mostly testing whether the conclusions of chapter 5 hold under different estimators, fixed-effects levels, and subgroup definitions.

C.1. Two-Way Fixed Effects Replication

This analysis aims to determine whether Two-Way Fixed Effects (TWFE) bias arising from staggered treatment timing yields results that differ from the main specification, using the [Sun & Abraham \(2021\)](#) estimator.

Table C.1

The TWFE replication of the baseline specification, highlighting the effects of wind turbines on housing prices with a similar method to prior literature.

	Classic DiD	log price		
		Housing characteristics	PC4 fixed effects	PC6 fixed effects
	(1)	(2)	(3)	(4)
Wind turbine ≤ 2 km	-0.0112 (0.0113)	-0.0327*** (0.0053)	-0.0078** (0.0034)	-0.0044 (0.0037)
Observations	1,828,776	1,828,776	1,828,725	1,779,284
R ²	0.35252	0.84263	0.87970	0.94533
Within R ²	3.34×10^{-5}	0.75696	0.76341	0.62445
Year-Municipality fixed effects	✓	✓	✓	✓
Housing characteristics (50)		✓	✓	✓
PC4 fixed effects			✓	
PC6 fixed effects				✓

Notes: The table reports the Average Treatment Effect of the Treated (ATT) estimates with classic TWFE models, using the logarithm of the transaction price as the dependent variable. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

C.2. Heterogeneity Analyses

This section reports the heterogeneity analyses under 4-digit postal code (PC4) fixed effects, rather than the 6-digit postal code (PC6) fixed effects levels presented in chapter 5. Note that the turbine height for tall turbines turns negative in Table C.5, which further illustrates that the heterogeneity results regarding turbine height are not robust across similar specifications.

Table C.2

The onshore effects of wind turbines on housing prices depending on the number of turbines during the treatment period with PC4 fixed effects.

	log price	
	1 Turbine	2+ Turbines
	(1)	(2)
Wind turbine ≤ 2 km	-0.0020 (0.0074)	-0.0021 (0.0079)
Observations	1,778,907	1,717,071
R ²	0.88003	0.88052
Within R ²	0.76405	0.76525
Housing characteristics (50)	✓	✓
Year-Municipality fixed effects	✓	✓
PC4 fixed effects	✓	✓

Notes: The table reports the Average Treatment Effect of the Treated (ATT) estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. It showcases the differences in trajectories between units with 1 turbine and those with 2 or more turbines. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** p<0.01, ** p<0.05, * p<0.10. The housing characteristics used as controls are listed in Table E.1.

Table C.3

The baseline estimate of onshore turbines divided by sale year.

	log price	
	Transactions from 2010-2017 (1)	Transactions from 2017-2023 (2)
Wind turbine ≤ 2 km	0.0015 (0.0050)	0.0049 (0.0048)
Observations	839,063	838,782
R ²	0.94289	0.94819
Within R ²	0.64304	0.64212
Housing characteristics (50)	✓	✓
Year-Municipality fixed effects	✓	✓
PC6 fixed effects	✓	✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. The analysis is divided by sale year, with column one containing the first 50% of transactions, and column two comprising the remaining half. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

Table C.4

The onshore effects of wind turbines on housing prices grouped by urbanisation category.

	log price				
	>2,500 addresses/km ²	1,500-2,499 addresses/km ²	1,000-1,499 addresses/km ²	500-999 addresses/km ²	<500 addresses/km ²
	(1)	(2)	(3)	(4)	(5)
Wind turbine ≤ 2 km	-0.0267** (0.0106)	-0.0069 (0.0057)	-0.0085 (0.0066)	-0.0028 (0.0094)	0.0030 (0.0110)
Observations	1,656,545	1,664,817	1,653,052	1,649,138	1,648,125
R ²	0.94590	0.94596	0.94573	0.94560	0.94524
Within R ²	0.62679	0.62562	0.62601	0.62606	0.62625
Housing characteristics (50)	✓	✓	✓	✓	✓
Year-Municipality fixed effects	✓	✓	✓	✓	✓
PC4 fixed effects	✓	✓	✓	✓	✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. It highlights the different treatment effects by the urbanisation category of treated units, highlighting the most detailed split by urbanisation grade possible with Nederlandse Vereniging van Makelaars (NVM) Brainbay housing transactions data. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** p<0.01, ** p<0.05, * p<0.10. The housing characteristics used as controls are listed in Table E.1.

Table C.5

The onshore effects of wind turbines on housing prices grouped by turbine height groups, with PC4 fixed effects.

	log price		
	<50 metres	50-150 metres	>150 metres
	(1)	(2)	(3)
Wind turbine ≤ 2 km	0.0022 (0.0354)	0.0027 (0.0067)	-0.0169* (0.0099)
Observations	1,705,564	1,765,949	1,691,714
R ²	0.88040	0.88018	0.88073
Within R ²	0.76492	0.76472	0.76544
Housing characteristics (50)	✓	✓	✓
Year-Municipality fixed effects	✓	✓	✓
PC4 fixed effects	✓	✓	✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. It highlights the distinct treatment effects of turbines across height categories. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** p<0.01, ** p<0.05, * p<0.10. The housing characteristics used as controls are listed in Table E.1.

Table C.6

The onshore effects of wind turbines on housing prices grouped by urbanisation category with Socioeconomic Status (SES) controls.

	log price		
	Baseline TWFE (1)	Baseline Urban (2)	Urban and SES (3)
Wind turbine ≤ 2 km	-0.0044 (0.0037)	-0.0023 (0.0037)	-0.0036 (0.0044)
Wind turbine ≤ 2 km $\times$ urban		-0.0040 (0.0069)	0.0054 (0.0083)
Standardised composite score			0.0344*** (0.0047)
Observations	1,779,284	1,779,284	1,394,321
R ²	0.94533	0.94533	0.94780
Within R ²	0.62445	0.62445	0.63213

Notes: The table reports the ATT estimates utilising Two-Way Fixed Effects (TWFE) specifications, using the logarithm of the transaction price as the dependent variable. It highlights the interaction effects of the treatment indicator and whether the property is located in an urban area, now also including the aggregate SES control variable. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

C.3. Socioeconomic Score Additional Specifications

This section shows whether the same specifications regarding Socioeconomic Status (SES) under 4-digit postal code (PC4) fixed effects still hold, compared to the original 6-digit postal code (PC6) fixed effect. The estimates in Table C.8 are statistically insignificant across nearly all quartiles, reflecting the limited identifying variation within most socioeconomic quartiles.

Table C.7

The effects of wind turbines on housing prices with SES as control variables, with year-municipality and PC4 fixed effects.

	log price				
	(1) Baseline Result	(2) SES Financial	(3) SES Education	(4) SES Labour History	(5) SES Aggregated
Wind turbine ≤ 2 km	-0.0011 (0.0056)	-0.0003 (0.0058)	-0.0002 (0.0059)	-0.0003 (0.0058)	-0.0003 (0.0058)
Standardised financial SES		0.0175*** (0.0050)			
Standardised educational SES			0.0715*** (0.0057)		
Standardised labour history SES				0.0167*** (0.0034)	
Standardised composite SES					0.0331*** (0.0047)
Observations	1,828,724	1,450,212	1,450,212	1,450,212	1,450,212
R ²	0.87975	0.88040	0.88050	0.88041	0.88042
Within R ²	0.76350	0.76706	0.76725	0.76707	0.76710
Housing characteristics (50)	✓	✓	✓	✓	✓
PC4 fixed effects	✓	✓	✓	✓	✓
Year-Municipality fixed effects	✓	✓	✓	✓	✓

Notes: The table reports the Average Treatment Effect of the Treated (ATT) estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. Standardised SES variables were included as controls to determine whether the base estimate would change. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.

Table C.8

The onshore effects of wind turbines split per aggregate SES quartile.

	log price			
	Q1 Lowest SES (1)	Q2 (2)	Q3 (3)	Q4 Highest SES (4)
Wind turbine ≤ 2 km	-0.0010 (0.0065)	-0.0138* (0.0078)	0.0041 (0.0089)	-0.0117 (0.0078)
Observations	442,132	440,019	437,533	440,790
R ²	0.94858	0.94043	0.93260	0.94154
Within R ²	0.62473	0.63465	0.62269	0.63338
Housing characteristics (50)	✓	✓	✓	✓
Year-Municipality fixed effects	✓	✓	✓	✓
PC6 fixed effects	✓	✓	✓	✓

Notes: The table reports the ATT estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. It shows the ATT by socioeconomic score quartile, with Q1 having the lowest and Q4 the highest aggregate socioeconomic score. The estimates are statistically insignificant across nearly all quartiles, reflecting limited identifying variation within most of them. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** p<0.01, ** p<0.05, * p<0.10. The housing characteristics used as controls are listed in Table E.1.

Additional Offshore Descriptive Elements

This appendix presents additional descriptive figures supporting the offshore analysis in section 5.2. The figures show the provincial distribution of offshore-exposed properties, illustrate the price differences between treated and control coastal units under PC4 location effects, and include further diagnostics of the identifying variation in the coastal sample. These figures reinforce the notion that causal inference in an offshore setting is currently very difficult.

D.1. Offshore Descriptives





Figure D.1: The number of properties ever exposed to offshore wind turbines in the Netherlands by province.



(a) The differences in price levels between ever-treated and control units after adjusting for 4-digit postal code (PC4) and year effects. (b) The robustness of the price differences for alternative coastal distance definitions.

Figure D.2: The price differences between ever-treated and control units and their robustness, now adjusted at the PC4 level.




Figure D.3: The descriptive differences in prices between treated and control properties at the PC4 level, relative to their treatment year, where $t = -1$ is the reference year. It also highlights the total observations per event time.




Figure D.4: The bars report the share of PC4 areas with treated observations with no pre-treatment observations, no post-treatment observations, and at most two post-treatment years.




Figure D.5: The differences between treated and control units after accounting for different time effects, comparing global versus municipal time effects at the PC4 level.

Housing Attributes

E.1. Housing Attributes of NVM Data

The Nederlandse Vereniging van Makelaars (NVM) dataset contains a wide variety of data on location, transaction prices, and housing characteristics. The housing characteristics included as controls are listed in Table E.1. These include structural characteristics, maintenance status, construction period, property type, and whether a property has monumental status. These characteristics allow the analysis to better isolate the turbine effect from differences in housing composition.

Table E.1

The types of housing characteristics used as controls in Difference-in-Differences (DiD) analyses.

Category	Variable	Clarifying comments
Structural	Living space (m^2)	
	Number of rooms	
Maintenance status (dummy)	Maintenance inside 1	The reference category. Values range from 1 to 9, with 1 being the worst and 9 being the best maintained.
	Maintenance inside 2	
	Maintenance inside 3	
	Maintenance inside 4	
	Maintenance inside 5	
	Maintenance inside 6	
	Maintenance inside 7	
	Maintenance inside 8	
	Maintenance inside 9	
	Maintenance outside 1	The reference category. Values range from 1 to 9, with 1 being the worst and 9 being the best maintained.
	Maintenance outside 2	
	Maintenance outside 3	
	Maintenance outside 4	
	Maintenance outside 5	
	Maintenance outside 6	
	Maintenance outside 7	
	Maintenance outside 8	
	Maintenance outside 9	
Construction period (dummy)	Building period: -1906	The reference category.
	Building period: 1906-1930	
	Building period: 1931-1944	
	Building period: 1945-1959	
	Building period: 1960-1970	

Continued on next page.

Category	Variable	Clarifying comments
Housing type (dummy)	Building period: 1971-1980	The reference category.
	Building period: 1991-2000	
	Building period: 2001-2010	
	Building period: 2011-2020	
	Building period: 2021-2030	
	Single family home	
	Mansion	
	Upstairs apartment	
	apartment with shared entrance	
	gallery apartment	
	Villa	
	Ground-floor apartment	
	Bungalow	
	Maisonette	
	Farmhouse	
	Manor	
	Stacked maisonette	
	Porch apartment	
	Penthouse	
	Canal house	
	Mezzanine	
	Estate	
	Student room	
	Double lower house	
Monument	Monumental status (yes/no)	

E.2. Baseline Specification

Table E.2 reports the baseline specification, but now presents the full set of housing attributes as well. A couple of things are noticeable. First, the (logarithm) living space is a strong predictor of transaction price. Under the specification with 6-digit postal code (PC6) fixed effects, a one per cent increase in living area is associated with an approximately 0.65 per cent higher transaction price, and the coefficient is even higher in other specifications. Next, maintenance also presents no surprises: as maintenance quality increases, whether inside or outside, the transaction price appears to rise as well. Both of these findings make economic sense, confirming that larger and better-maintained properties will lead to higher property prices.

The construction period of the properties also seems to matter. After applying location fixed effects, it becomes clear that recently constructed properties (from the 90s) have higher transaction prices, conditional on the other housing characteristics. This price premium for recently constructed properties appears to become larger the more recently a property was constructed. The property type itself also behaves as expected with general economic logic, relative to the single-family reference category. Detached and higher-status property types such as villas, manors, canal houses, and estates carry substantial price premiums, while apartment types such as porch apartments, maisonettes, student rooms and ground-floor apartments generally have a price reduction. It is interesting to note that the price premiums for more expensive properties decrease after increasingly strict fixed effects are applied. This implies that these more expensive properties are also located in neighbourhoods that, on average, are more expensive. For instance, the point estimate of villas drops from 0.332 to 0.305 after applying 4-digit postal code (PC4) fixed effects, and to 0.151 after applying PC6 fixed effects. Lastly, the monumental status of a property is associated with a small price premium.

Table E.2

The full regression table of the baseline results that include housing characteristics.

	Classic DiD	log price		
		Housing characteristics	PC4 fixed effects	PC6 fixed effects
Wind turbine ≤ 2 km	0.0138 (0.0174)	-0.0247*** (0.0075)	-0.0011 (0.0056)	-0.0081* (0.0041)
log living space		0.8424*** (0.0041)	0.7974*** (0.0026)	0.6476*** (0.0023)
Number of rooms		-0.0038*** (0.0005)	-0.0027*** (0.0003)	0.0038*** (0.0002)
Maintenance inside 2		0.0184* (0.0095)	0.0215*** (0.0082)	0.0170*** (0.0066)
Maintenance inside 3		0.0559*** (0.0071)	0.0543*** (0.0062)	0.0478*** (0.0049)
Maintenance inside 4		0.0732*** (0.0076)	0.0706*** (0.0066)	0.0646*** (0.0052)
Maintenance inside 5		0.0966*** (0.0072)	0.0983*** (0.0063)	0.0936*** (0.0049)
Maintenance inside 6		0.1299*** (0.0073)	0.1314*** (0.0063)	0.1266*** (0.0049)
Maintenance inside 7		0.1999*** (0.0073)	0.2011*** (0.0063)	0.1955*** (0.0049)
Maintenance inside 8		0.2532*** (0.0074)	0.2504*** (0.0065)	0.2474*** (0.0050)
Maintenance inside 9		0.2743*** (0.0075)	0.2711*** (0.0065)	0.2708*** (0.0051)
Maintenance outside 2		-0.0094 (0.0114)	-0.0125 (0.0104)	0.0071 (0.0086)
Maintenance outside 3		0.0186** (0.0089)	0.0166** (0.0081)	0.0355*** (0.0067)
Maintenance outside 4		0.0076 (0.0094)	0.0058 (0.0084)	0.0372*** (0.0070)
Maintenance outside 5		0.0189** (0.0090)	0.0192** (0.0080)	0.0621*** (0.0066)
Maintenance outside 6		0.0300*** (0.0091)	0.0272*** (0.0081)	0.0761*** (0.0066)
Maintenance outside 7		0.0525*** (0.0091)	0.0433*** (0.0081)	0.0956*** (0.0066)
Maintenance outside 8		0.0726*** (0.0091)	0.0591*** (0.0082)	0.0979*** (0.0067)
Maintenance outside 9		0.0756*** (0.0092)	0.0601*** (0.0082)	0.0934*** (0.0068)
Building period: 1906-1930		-0.0379*** (0.0057)	-0.0011 (0.0025)	-0.0087*** (0.0014)
Building period: 1931-1944		-0.0432*** (0.0078)	0.0210*** (0.0034)	-0.0021 (0.0018)
Building period: 1945-1959		-0.0945*** (0.0070)	-0.0316*** (0.0034)	-0.0065*** (0.0021)
Building period: 1960-1970		-0.1966***	-0.1045***	-0.0218***

Continued on next page.

	Classic DiD	log price		
		Housing characteristics	PC4 fixed effects	PC6 fixed effects
		(0.0066)	(0.0032)	(0.0021)
Building period: 1971-1980		-0.2068***	-0.1142***	-0.0188***
		(0.0063)	(0.0032)	(0.0022)
Building period: 1981-1990		-0.1401***	-0.0489***	0.0018
		(0.0066)	(0.0030)	(0.0024)
Building period: 1991-2000		-0.0626***	0.0323***	0.0523***
		(0.0064)	(0.0031)	(0.0024)
Building period: 2001-2010		-0.0525***	0.0561***	0.0778***
		(0.0070)	(0.0033)	(0.0025)
Building period: 2011-2020		-0.0308***	0.0645***	0.0848***
		(0.0067)	(0.0036)	(0.0032)
Building period: 2021-2030		0.0053	0.0964***	0.0992***
		(0.0088)	(0.0064)	(0.0076)
Mansion		0.1453***	0.1206***	0.0663***
		(0.0039)	(0.0023)	(0.0012)
Upstairs apartment		-0.0002	-0.0926***	-0.1246***
		(0.0059)	(0.0032)	(0.0028)
apartment (shared entrance)		-0.0249***	-0.0622***	-0.1064***
		(0.0048)	(0.0033)	(0.0031)
gallery apartment		-0.0655***	-0.0870***	-0.1074***
		(0.0050)	(0.0035)	(0.0029)
Villa		0.3320***	0.3054***	0.1509***
		(0.0038)	(0.0028)	(0.0018)
Ground-floor apartment		0.0378***	-0.0197***	-0.0534***
		(0.0052)	(0.0031)	(0.0033)
Bungalow		0.2596***	0.2345***	0.1008***
		(0.0032)	(0.0027)	(0.0020)
Maisonette		-0.1170***	-0.1417***	-0.1393***
		(0.0045)	(0.0029)	(0.0028)
Farmhouse		0.2009***	0.2341***	0.1204***
		(0.0047)	(0.0034)	(0.0030)
Manor		0.3561***	0.3430***	0.1839***
		(0.0045)	(0.0039)	(0.0034)
Stacked maisonette		0.0050	-0.0477***	-0.0459***
		(0.0103)	(0.0064)	(0.0045)
Porch apartment		-0.1055***	-0.1096***	-0.1270***
		(0.0108)	(0.0063)	(0.0055)
Penthouse		0.1942***	0.1396***	0.0618***
		(0.0049)	(0.0039)	(0.0040)
Canal house		0.1726***	0.1203***	0.0793***
		(0.0140)	(0.0084)	(0.0068)
Mezzanine		0.0667***	-0.0426***	-0.1091***
		(0.0080)	(0.0050)	(0.0039)
Estate		0.4337***	0.4582***	0.3438***
		(0.0253)	(0.0240)	(0.0283)
Student room		0.0890	0.0655	-0.0102
		(0.1607)	(0.1595)	(0.1683)

Continued on next page.

	Classic DiD	log price		
		Housing characteristics	PC4 fixed effects	PC6 fixed effects
Double lower house		0.1637*** (0.0155)	0.0347*** (0.0080)	0.0059 (0.0060)
Monument		0.0642*** (0.0056)	0.0338*** (0.0034)	0.0262*** (0.0022)
Year-Municipality fixed effects	✓	✓	✓	✓
Housing characteristics (50)		✓	✓	✓
PC4 fixed effects			✓	
PC6 fixed effects				✓

Notes: The table reports the Average Treatment Effect of the Treated (ATT) estimates with the [Sun & Abraham \(2021\)](#) estimator, using the logarithm of the transaction price as the dependent variable. The regression table is identical to the one presented in chapter 5, but now also includes the housing characteristics. The maintenance scores range from 1 to 9, with 1 being the lowest and 9 the highest. Standard errors are clustered at the neighbourhood level and reported in parentheses. Statistical significance is indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The housing characteristics used as controls are listed in Table E.1.