Wind Turbines and the Housing Market

Unfinished Draft. Please do not distribute or cite.

Carsten Andersen[1] and Timo Hener[2]

[1]Department of Economics and Business Economics, Aarhus University
[2]Department of Economics and Business Economics, Aarhus University, and CESifo

January 28, 2022

Abstract

We analyze the effect of wind turbines on house prices in proximity to new installations. Utilizing the universe of Danish house transactions since 1992 and data on all turbines ever established in Denmark, allows us to control for individual house fixed effects in a staggered difference-in-difference estimation. In the analysis, we distinguish between effects of proximity and shadow flicker from rotors partly covering the sun. Our results indicate a substantial negative effect on house prices from close turbines, which is stronger for tall turbines and houses affected by shadow flicker. Our results suggest a more nuanced view on the local externalities of wind turbines that heavily depend size and relative location.

1 Introduction

Wind power plays an important role in the policy mix aimed at slowing down climate change. Indeed, wind power is the second-fastest growing renewable energy source for electricity production in the world (IEA 2020)[1]. To stay on course with the 2030 Sustainable Development Goal (SDG7) of the United Nations to substantially increase the share of renewable energy in the global energy mix (IBRD 2020), wind power is expected to expand further. While wind power avoids the global externalities associated with conventional fossil fuels, it implies negative *local* externalities. Wind turbines emit a low-frequency noise, are visually unattractive, and cause a nuisance known as *shadow flickering*, which manifests as intermittent light behind spinning rotors. As these externalities intensify with proximity, expanding wind turbines to new places is becoming increasingly costly for local populations.

A large body of studies investigates effects of nearby wind turbines on, among others, suicide (Zou 2017), cardiovascular diseases (Poulsen and Raaschou-Nielsen 2018) and house prices (Jensen, Panduro, and Lundhede 2014; Dröes and Koster 2016).[2] The literature on wind turbine externalities in the housing market reaches very mixed conclusions. While some find significant, and economically important, negative external costs on house prices ranging between 2% and 16% (Jensen, Panduro, and Lundhede 2014; Gibbons 2015; Dröes and Koster 2016), others find no effects (Hoen et al. 2011; Lang, Opaluch, and Sfinarolakis 2014). In a back-of-the-envelope calculation, Dröes and Koster 2016 estimate that the external costs of a wind turbine is at least 16% of its construction costs. In general,

[1]Solar power had more net capacity addition since 2016.
[2]See also Zerrahn 2017 for a review of externalities of wind turbines.

the reasons for the varying results seem to be differences in the approaches and aggregations used, limited data availability on exact locations, timing, and characteristics of properties and turbines, and a focus on specific aspects of turbine nuisances.

In this paper, we examine the effect of wind turbines on the housing market and make three principal contributions, aiming at reconciling the vast variability of findings in the existing literature. First, we employ extraordinarily rich data from Denmark, one of the world's forerunners in wind energy production already reaching levels above 43% of total electricity production in 2017 (Energinet 2020). We observe geo-coded 6,000 wind turbine installments and the transactions of all 1.34 million traded properties, allowing for quasi-experimental identification of causal effects in generalized difference-in-differences estimations along with identification of important heterogeneities in treatment effects. Second, we are the first to identify the additional effect of shadow flickering, enabled by the use of variation in the relative angles between properties, turbines, and the position of the sun. The joint estimation of these effects allows us to reconcile findings on proximity (Dröes and Koster 2016) and visual pollution (Gibbons 2015). Third, we use not only newly built turbines but also turbines going out of service, allowing us to disentangle differential responses.

Evaluating the effects of wind turbine proximity on house prices is not straightforward. As wind turbines are not allocated randomly across geographic areas, a simple cross-sectional comparison of house prices in proximity to wind turbines with houses unaffected by wind turbines does not reveal the true extent of the externalities. For our estimation, we make use of data with rich variation both between areas and across time in turbine exposure to overcome selection and endogeneity issues common in hedonic pricing approaches. As the main specification we use a generalised difference-in-difference model with granular geographic, time and property controls, which address the issue of potentially unobserved pre-existing trends in both wind turbine exposure and housing market development. We employ detailed panel data on *every* housing transaction and *every* wind turbine in Denmark between 1992 and 2019. We use exact coordinates to calculate distances and angles between all Danish properties and 6,000 turbines. Information on wind turbines include height, diameter of rotors, and yearly electricity production, which we use to explore heterogeneity in effects of turbine exposure. Finally, many properties are traded multiple times in the period which allows the strictest econometric specifications to include property fixed effects. Hence, our empirical approach allows for arbitrary correlation between wind turbine exposure and time-constant unobservables of houses, and the vast number of transactions ensures that it is not at the cost of reasonable precision of the estimates. Similarly, we can show results net of location-specific time trends implying that unobserved time-varying trends are less likely to confound our estimations.

In our baseline estimation, we find that a turbine within 2km of a house decreases the property value by 2.3%, and no price effect is detectable after 2.5km. The negative price effect shows considerable heterogeneity for turbines of different height. While short turbines below 60m exhibit (insignificant) house price reduction of less than 2%, medium sized turbines of up to 120m show treatment effects above 3%. Dwarfing these impacts, modern tall turbines in excess of 120m cut property prices by 8%. Moreover, we find that houses affected by shadow flicker, net of turbine height, suffer an additional value loss of 5.2%.

The fact that our results indicate that tall turbines inflict the largest damages on house prices does not, however, imply a policy recommendation against tall turbines. The tallest turbines have at the same time the most powerful generator units and overcompensate the larger loss in property value by more than proportional savings in emissions. To make this comparison, we count the property damages from wind turbines against the positive environmental externalities of carbon dioxide savings using three scenarios for the social cost of carbon (SCC). Only the smallest turbines are not producing a net positive externality in the Danish case if we assume the lowest cost of SCC (€50). All other estimates imply that the turbines, even if reducing local property values, produce a net gain for society.

Compared to the literature, our proximity effects of turbines on house prices are larger on average than in Dröes and Koster 2016, and smaller than in Jensen, Panduro, and Lundhede 2014, and Gibbons 2015. We can reconcile these estimates with the effect heterogeneity across turbine heights and the additional impact from shadow flicker. To the best of our knowledge, this is the only study observing every property transaction and turbine in a national market. The only study coming close to this this wealth, and precision, of information is Dröes and Koster 2016, covering 70% of housing transactions in the Netherlands. Many studies are limited by only having information on a few selected turbine sites (e.g., Lang, Opaluch, and Sfinarolakis 2014; Jensen, Panduro, and Lundhede 2014 and Vyn and McCullough 2014), or vastly fewer traded properties (e.g., 2,100 property sales in Sunak and Madlener 2016 implying few repeat sales), or using data on the aggregated level of the postal code, (Gibbons 2015; Zou 2017). In comparison, our data holds information on *every* individual housing transaction and *every* individual wind turbine for an entire country over a period of 27 years.

The remainder is organized as follows. Section 2 introduces the data and estimation strategy, section 3 shows all estimation results and the comparison of externalities, and section 4 concludes.

2 Data

2.1 Data sources

In the analysis we combine three data sources.

1) Property trades

We select all property trades in Denmark 1992-2019 from publicly available transaction registers. This amounts to 2,810,039 transactions. After exclusion of vacation homes, transactions within the family and price outliers we have 2,104,742 transactions in the main analysis sample. These originate from 1,087,668 unique residential units leaving rich variation for specifications with house fixed effects. The data contains exact address information, selling prices, the exact date of sale, size and type of unit (e.g., apartment or house).

2) Wind Turbines

The Danish Energy Agency provides publicly available information on all wind turbines ever in operation beginning in 1977 (Danish Energy Agency 2021). The data includes information on the geographical coordinates, the date of commission and decommission, turbine height, capacity, and so on for 9,591 turbines. We exclude 686 turbines with missing geographical information and 383 turbines with missing information on height. We further exclude 1,069 private wind turbines for one-family homes and 956 offshore turbines. This leaves 6,880 onshore turbines for analysis. We define the total height of a turbine as the the axis height plus half the diameter of the rotor blades. In the analysis we split turbine height into three categories of low turbines (<60m), medium-sized (60-120m) and high (>120m). Figure 1 shows a map of the turbines in the analysis sample. Turbines are particularly prevalent along the western coast(s) where wind conditions are favourable.

3) Housing data

From the Agency for Data Supply and Efficiency we obtain housing information for all 3,794,380 million addresses ever existing in Denmark. Each address has geographical coordinates and information on use (e.g., residential or business). The application of this data is twofold. First, the data is merged with the property trades data set to supply the transactions with geographical coordinates. Second, to calculate the externalities from wind turbines we need to distinguish between housing unit types. With the latitude and longitude coordinates from the housing data we are then able to calculate haversine distances and angles from all addresses to all nearby turbines. In all specifications we focus on the characteristics of the *nearest* turbine in any given year.

Below, we use the terms “home”, “address”, “property”, “residential unit” and “house” interchangeably.

2.2 Descriptive statistics

Table 1 shows mean values for the variables we use, with standard deviations in parentheses for continuous variables. Key variables are prices and the main treatment indicator of being below 2 km from a wind turbine in a given year. Columns (1) and (2) are property transactions data over the period 1992-2019 while column (3) is a snapshot of all residential housing units in 2019. We mostly use column (1) in the empirical analysis, while column (3) is used for a back-of-the-envelope calculation of total externalities. In general differences between column (1) and (3) may be both due to transacted homes being different from non-transacted homes or general changes in 2019 vis-á-vis the period 1992-2019. Column (2) is the sub-sample we use for robustness analyses in the Appendix where we restrict the sample to transacted homes that are below 2 km from a wind turbine at some point in the data period (exploiting only the temporal variation in treatment).

Column (1) is the main sample of transactions used in the analysis. The average price of a home is €199,000 and 20 percent of all homes are within 2 km of a turbine at the time of sale. Hence, 20 percent of observations are treated according to our main definition of treatment and sample selection. In column (2), 76 percent of all transactions are below 2 km from a turbine at the same of sale. The average price of the sample in column (2) is also lower than the full sample in column (1), presumably both due to treatment effects and selecting for less urban homes. Similarly, these houses are larger. For the 2019 housing data we compute the value of each home (see section 4.5). The average value of all residential units in 2019 is €250,000, i.e., 25% higher than the value of the transacted homes (inferred by the prices). These figures are not directly comparable since average prices in 2019 is higher than prices 1992-2019 due to secular growth in house prices.

Returning to column (1) and the main sample we see mean values for interactions between the main treatment indicator and heights of the nearest turbine. More than half of the treated units originate from a turbine height of below 60 meters (13% in total in the full sample). Only one percent (0.002/0.20) of treated units have a high turbine (>120 meters) as the nearest turbine. Regulations stipulate that turbines have to placed at least four times total turbine height away from residential homes so it is unsurprising that high turbines are difficult to place close to many homes. In the full sample, only 0.2% of transacted homes are exposed to a high turbine below 2 km at the time of sale. This figure is identical to mean exposure to shadow flicker at the time of sale.

Moving further down the table we see variables that are used for heterogeneity analyses, allowing differential treatment effects both by distance and height. For example, 6.7% of sold homes are 1000-1500 meters from the nearest turbine while this is somewhat lower for the 2019 housing data (4.5%). However, the share of high turbines is higher for the housing data, probably reflecting that the share of high turbines is increasing over time.

Figure 2 shows trends in central variables over time. Figure 2 (a) shows growth in house prices for homes above and below 2 km from a wind turbine, respectively. For all years, homes close to a turbine is less expensive which probably reflects both treatment effects and selection effects. The aim in the first part of the analysis is to disentangle how much of the difference in prices can be attributed to wind turbine proximity itself. Apart from the financial crisis setback there is a secular price growth for both types but the difference is widening over time. In general, a widening gap can reflect both urban-rural price growth differentials, changes in treatment effects, or changes in the composition of homes sold.

Figure 2 (b) shows the number of active turbines in the sample by year. After 2002, there is a sudden drop due to a new center-right government in 2001 reducing subsidies for wind energy, consequently scrapping a substantial amount of medium-and low-sized turbines. When looking at dynamic effects we provide estimates both of commissioning turbines and decommissioning turbines, i.e., we use variation from homes subject to gains of treatment status and homes subject to loss of treatment status.

Figure 2 (c) shows that the average height of turbines in operation increase from 40 meters in 1992 to 80 meters in 2019. The growth is more pronounced for new turbines entering the sample where the average height towards the end exceeds 140 meters. Some years have very few turbines entering the sample which may lead to a noisy graph (e.g., 2007).

Figure 2 (d) shows the share of transacted homes that are close to a wind turbine is between 14% and 25% in 1992-2019. There is a sudden drop in 2002 mimicking the drop from (b). The subsequent downward trend may be explained both by compositional changes in transacted homes and a general trend towards decommissioning lower turbines and commissioning higher turbines.

3 Methods

Identifying the causal effect of turbines on house prices is challenging, because turbines are not randomly placed. Wind conditions, land values, and government regulation affect the decision where to install turbines. The sites are often close to the coast, where stable wind promises higher efficacy, and in areas with low land values that keep costs down. Governments may impose minimum distances to settlements and compensation for property owners. Any of the factors in the decision for turbine sites is a potential determinant of or correlated with property prices, yielding a bias in cross-sectional regressions.

Our identification strategy exploits variation in when and where turbines are running. We use information on the exact commissioning and decommissioning date for every turbine to identify whether a property is close to an operating turbine at any point in time. We exploit that turbines are installed and scrapped in the proximity of houses, while other properties either never or at a different point in time have a turbine close by. We essentially compare houses before and after a turbine was built or scrapped in its vicinity to houses in the same period that did not experience a turbine event.

We use two principal variables of interest, one that identifies the effect of a close turbine and one that elicits the impact of shadow flicker. We define an indicator variable $D_{i,t}$ that takes the value of one if house i in year t has a turbine active within 2 km, and zero otherwise. This dummy captures

Figure 1: Onshore turbines in Denmark

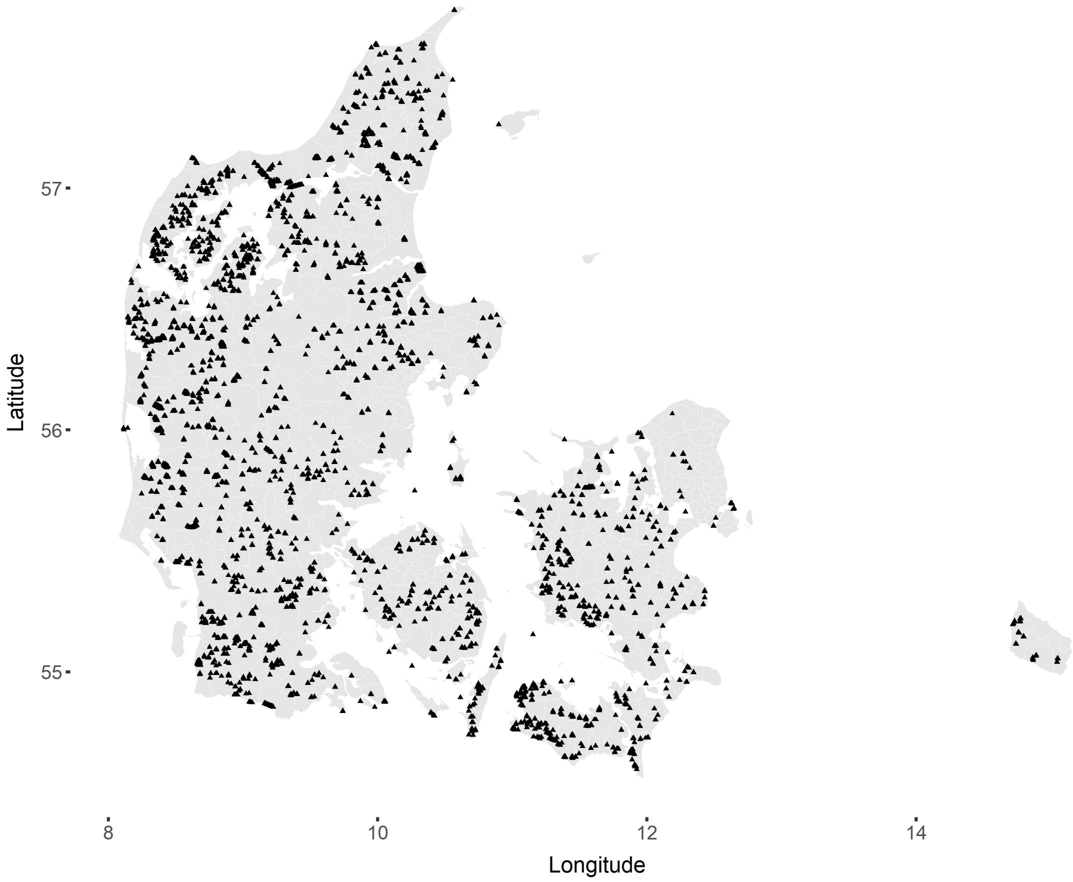

Notes: Figure 1 shows all onshore turbines in Denmark ever in operation in the period 1977-2019.

both the commissioning and decommissioning of turbines at close distance and we assume in the baseline that both events have the same impact magnitude with opposite signs. An active turbine may affect close houses through noise exposure (see Zou 2017) and visibility (see Gibbons 2015). As both impacts increase with proximity and decrease with blockages in direct sight, we regard noise and visibility effects as indistinguishable.

However, the impact of shadow flicker, our second variable of interest, is distinguishable as the exposure is only partially correlated with proximity. The nuisance of shadow flicker is the rhythmic change in light caused by the rotating turbine's blades partially blocking sunlight for a short moment. We use the coordinates of the house and the turbine, the tower height, and the diameter of the rotor to compute the orientation and minimum as well as maximum angle that the rotor appears from the house relative to the horizon. Then, we compute whether the sun ever passes the rotor as seen from the house. For the orientation towards the turbine, the azimuth, we compute the lowest possible altitude angle of the sun at winter solstice and the highest possible angle at summer solstice. For every house, this computation yields an angle range for the rotor and the sun as seen from the house. The shadow flicker indicator is one if the rotor and sun angle range overlap meaning that the sun at some point during the year passes behind the rotors, and zero otherwise.[3] We again define an indicator variable sf_{it} taking the value of one if the rotors of the closest turbine potentially cast a shadow at

[3]We discard sun altitude angle below three degrees, which is when shadows become very long and diffused such that shadow flicker is less noticeable.

Figure 2: Evolution in key variables 1992-2019

(a) Average prices by treatment status (2020 euros)

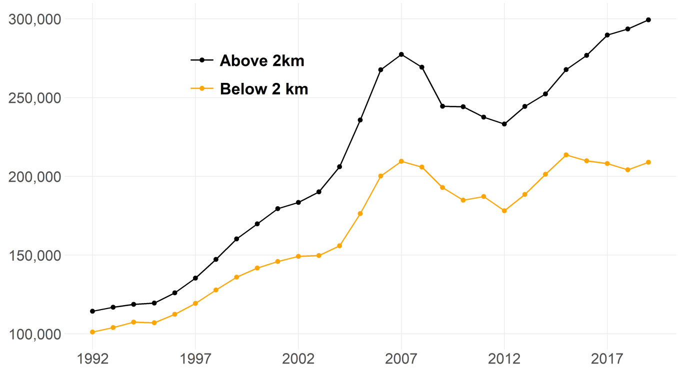

(b) Number of active turbines

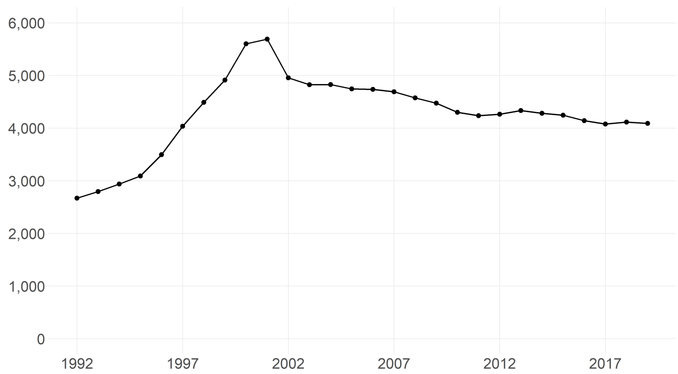

(c) Average height of wind turbines

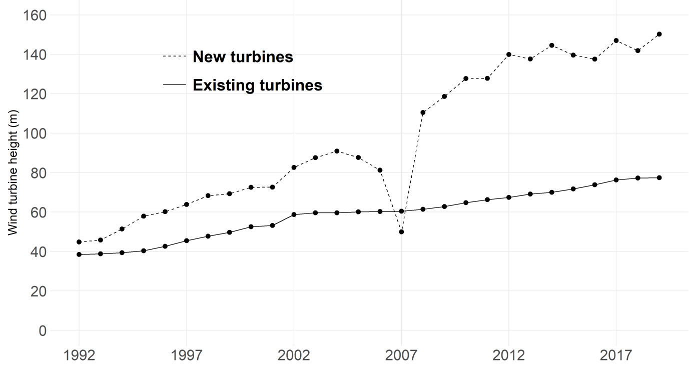

(d) Share of transacted homes below 2 km from a wind turbine at the time of sale

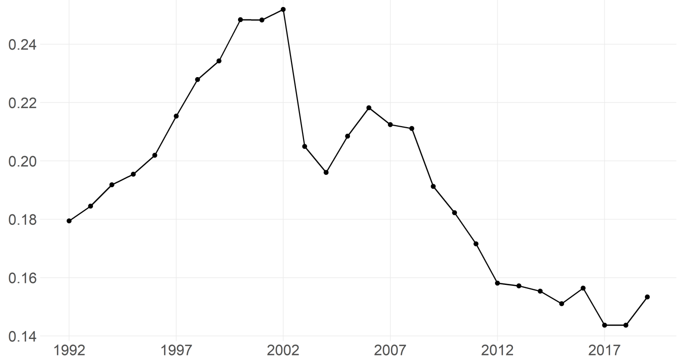

house i in period t. Because of the sun's changing altitude angle throughout the day, some houses are affected by shadow flicker, while others at the same distance from the turbine are not. Houses east and west of turbines have the highest likelihood of shadow flicker, north of turbines shadows are only present very close to turbines and they never appear south of turbines.

Our main outcome variable is the log price of a house. For the hedonic pricing regressions, we use a sample of all transactions. To investigate whether the probability of a sale is affected by wind turbines, we replace the outcome variable by an indicator that equals one if a house is sold, and zero otherwise, in a sample of all addresses. We describe our estimation for house prices, but the principle carries over to the estimation for sale probabilities.

Our identification strategy is twofold. We first use the indicator for turbine proximity in a flexible difference-in-differences type estimation, and second employ an event study to elicit anticipation and lagged impacts as well as differences between commissioning and decommissioning of turbines.

In the main estimation, we estimate the impact of turbine proximity on house prices at the address-year level in the following baseline equation

$$log(Y_{i,t}) = \alpha_i + \lambda_t + \gamma D_{i,t} + \varepsilon_{i,t}. \tag{1}$$

The dependent variable is the logarithm of the house price Y at address i in sales year t. As turbines are not randomly allocated, it is important to control for factors determining both the location of turbines and house prices. Typically, turbines would be placed in more rural areas with cheaper land. To exclude fixed differences in house prices between addresses, we include versions of location fixed effects in α_i. We vary the definition between specifications, with municipality fixed effects as the most aggregate regional measure, zip code fixed effects aggregating smaller local areas, and address fixed effects controlling for all time-constant house-specific unobservable characteristics. To capture temporal rises and falls in house prices that might be correlated with turbine expansions, we include fully flexible year fixed effects in λ_t.

A threat to identification are differential trends in house prices that correlate with turbine installments. As the above specifications only capture common price changes over time, the estimates are biased if turbines are placed in areas that are on the decline (or uprise) relative to the control areas. Therefore, we include specifications with interactions between α_i and λ_t. Specifically, we control for region-, municipality-, or zip code-specific year fixed effects that exclude flexibly deviations from common changes over time in increasingly smaller areas. These specifications are thus robust to turbine positioning that reacts to temporal price shifts in local areas, rendering the common trend assumption less demanding.

The indicator $D_{i,t}$ depicts whether a turbine is within a 2 km range when the house is sold. It can be zero if no turbine has been within 2 km or if a turbine within 2 km has been decommissioned in the past. The parameter γ identifies the effect of a turbine on house prices under the assumptions of being homogeneous across time and addresses, and of having the same magnitude with opposite signs for the commissioning and decommissioning of turbines. In a later specification, we include a second dummy for shadow flicker that identifies the additional impact on top of the effect of turbine proximity. In the heterogeneity analysis, we vary the treatment characteristics. First, we include flexibly distances to turbines in bins to identify at which distance the price impact vanishes. Second, we split the turbine indicator by the height of the turbine, allowing for differential impacts on prices. Third, we interact distance and height bins to identify differential distance and price impact by height of the turbine.

Another threat to identification appears when treatment effects are heterogeneous. The estimate of γ can be shown to represent a weighted average of all possible before-after comparisons with any other unit in the sample (De Chaisemartin and d'Haultfoeuille 2020; Goodman-Bacon 2018). Consequently, some of the control observations are already treated. If the treatment effects are not constant over time in the control units, i.e., prices increase or decrease because of the treatment years after commissioning, the comparison to the treated units become biased as too much or too little common time variation in the untreated potential outcome is subtracted. The weights attached to the individual comparisons are not in general proportional to the group size or even non-negative (De Chaisemartin and d'Haultfoeuille 2020). If the treatment design is staggered, such that an increasingly large fraction of the sample becomes treated over time, more observations towards the end of the sample receive negative weights, and observations in the middle of the period have larger treatment variance, giving them larger weights (Goodman-Bacon 2018). In our setting, treatment is only partially staggered, because the number of turbines grows in general, but the growth almost comes to a halt and many houses see their closest turbines removed. We test further for the robustness of the estimation to contaminated control groups by comparing results from samples that include all house sales and thus many never-treated control groups to results from estimations with houses that will at some point all be treated.

In the second approach, we estimate the dynamic treatment effects of turbines on house prices using the event study estimation equation

$$log(Y_{i,t}) = \alpha_i + \lambda_t + \sum_{\tau} \gamma_\tau \mathbf{1}[t - E_i = \tau] + \varepsilon_{i,t}. \tag{2}$$

We estimate only the effects of first turbines in a sample of houses that only change their treatment status once from zero to one in year E_i. The event indicator function $\mathbf{1}[.]$ equals one if the first turbine at house i is commissioned τ years before or after the sale in year t, where $\tau \in \{-7, -6, ..., 7, 8\}$, and zero otherwise. The endpoints at 8 years before and after treatment are trimmed, such that they include any event appearing before resp. after. This specification is equivalent to the standard case in Schmidheiny and Siegloch 2020. The left out category is 8 years and longer before the event, implying that the γ_τ are relative house prices 8 and more years before a turbine was installed. Impacts for $\tau < 0$ represent anticipation effects, which may well be present as the planning phase takes several years. Estimates for $\tau \geq 0$ represent immediate and lagged effects of the commissioning of turbines. In fully forward looking markets, the entire price reaction should be absorbed in the anticipation effect. Later impacts would point towards a salience effect that is based on experiencing the running turbine in proximity to the address.

To analyze the impact of decommissioning turbines, we add houses to the dynamic event study model that changed their treatment status from one to zero in year E_i'. In this model, we separately identify the dynamic treatment effects for commissioning and decommissioning of turbines in the extended estimation equation

$$log(Y_{i,t}) = \alpha_i + \lambda_t + \sum_{\tau} \gamma_\tau \mathbf{1}[t - E_i = \tau] + \sum_{\tau} \delta_\tau \mathbf{1}[t - E_i' = \tau] + \varepsilon_{i,t}. \tag{3}$$

Here, δ_τ identifies the dynamic effect of decommissioning turbines on house prices. Again, effect for $\tau < 0$ represent anticipation effects, which are to be expected of shorter duration, as a decommissioning may well have a shorter planning horizon. Immediate and lagged effects are seen for $\tau \geq 0$. After decommissioning of a turbine, the physical installation can stay intact for a long time until scrapping, which induces uncertainty about the permanence of the decommissioning for houses close by.

Table 1: Descriptives: House prices and housing data

	House prices		
	(1) Full sample	(2) Ever below 2 km	(3) Housing data 2019
Price	199,106 (141,048)	167,756 (123,764)	-
Value	-	-	249,684 (177,919)
Below 2 km	0.20	0.76	0.13
× Turbine height 0-60m	0.13	0.50	0.05
× Turbine height 60-120m	0.06	0.25	0.07
× Turbine height >120m	0.002	0.008	0.009
Below 3 km	0.37	0.85	0.29
Ever below 2 km	0.26	1.00	0.22
Ever below 3 km	0.46	1.00	0.42
Shadow Flicker	0.002	0.006	0.002
Size (m2)	129 (54)	144 (55)	111 (57)
Year	2005 (8)	2005 (8)	2019 (0)
Distance 0-500m	0.008	0.032	0.003
× Turbine height 0-60m	0.007	0.027	0.002
× Turbine height 60-120m	0.001	0.005	0.001
Distance 0-1000m	0.046	0.180	0.025
× Turbine height >120m	0.0003	0.0012	0.0011
Distance 500-1000m	0.038	0.147	0.022
× Turbine height 0-60m	0.026	0.103	0.009
× Turbine height 60-120m	0.011	0.043	0.011
× Turbine height >120m	0.0003	0.0012	0.0011
Distance 1000-1500m	0.067	0.260	0.045
× Turbine height 0-60m	0.044	0.171	0.016
× Turbine height 60-120m	0.022	0.086	0.026
× Turbine height >120m	0.0007	0.0028	0.0030
Distance 1500-2000m	0.083	0.322	0.064
× Turbine height 0-60m	0.052	0.204	0.024
× Turbine height 60-120m	0.029	0.114	0.036
× Turbine height >120m	0.0011	0.0042	0.0047
Observations	2,104,742	541,184	2,817,450

Notes: Table 1 shows mean values along with standard deviations in parentheses (only for continuous variables). The first column is the house price data used in the main analysis. The second column is a subset where we only exploit the temporal variation in the data (see results in Appendix). The third column is all residential units in Denmark as of 2019 used for the back-of-the-envelope calculation for total externalities. All prices and values are in 2020 euros.

4 Results

4.1 Static effects of close turbines on house prices

Our first set of results elicits the effect of a turbine within a 2 km diameter around the house on its sales price. The effect is identified from variation in the timing of turbines being put in and out of operation. We show results for increasingly demanding specifications of equation 1 in table 2, with the full sample results in panel A. Turbines are placed close to houses of significantly lower value than the average, as is evident from the large negative coefficients in column (1) where only year fixed effects are included. Much of this variation is explained by location. The negative effect reduces in magnitude when we include fixed effects for the municipality (column 2) and further with postal code fixed effects (column 3). None of these estimates is suitable to identify the causal effect of wind turbines on house prices. Only the specification with address fixed effects in column (4) can exclude endogenous location choices of turbines. Indeed, the estimate is again smaller in magnitude, but still negative and statistically significant at the 1% level. The estimate implies that placing a turbine within 2 km of a house decreases the sales price by 2.9%. Under the assumption that turbines are not placed close to houses that are expected to lose value relative to control houses, we can interpret the estimate as the causal impact of turbines.

In columns (5) and (6), we take a closer look at the timing assumption. It is conceivable that turbines are placed in areas that are on a decline relative to more prosperous control areas. Controlling for fully flexible price changes over time specific to the municipality (column 5) or postal code (column 6) reveals estimates of -2.3% for close turbines, both statistically significant. The estimates are thus robust to differential changes in prices over time in small areas around the houses down to the level of the postal code. Our most conservative estimate of the impact of a turbine within 2 km on house prices of -2.3% is considerably larger than the preferred estimate by Dröes and Koster 2016 of -1.4%.

We limit the sample for the estimation to houses that are within 2 km of a turbine at some point during our observation period and present results in panel B of table 2. As expected, the estimates are smaller in specifications without and with only coarse location fixed effects in columns (1)-(3), because the control houses are limited to more comparable close addresses. The preferred specifications with interacted fixed effect for year and municipality resp. year and postal code yield similar treatment effect estimates. The negative 2.2% effect in column (5) is very comparable to the estimate from the full sample and statistically significant. In column (6), the flexible specification is overly restrictive and leaves too little variation to identify a statistically significant effect. The within postal code times year fixed effect can only be identified by variation within those cells, but there is mechanically not much variation left when all control houses that are never treated are dropped from the sample. All further results shown use the full sample, the corresponding restricted sample results can be found in the appendix.

Table 2: Effect of wind turbine proxmity on house prices for different specifications

Dependent Variable:	ln(Price)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full sample						
Turbine $\leq$ 2km	-0.245***	-0.100***	-0.065***	-0.029***	-0.023***	-0.023**
	(0.022)	(0.011)	(0.011)	(0.011)	(0.008)	(0.009)
Observations	2,104,742	2,104,742	2,104,742	2,104,742	2,104,742	2,104,742
R^2	0.28	0.48	0.51	0.84	0.85	0.86
Panel B: Ever below 2 km (temporal variation only)						
Turbine $\leq$ 2km	-0.090***	-0.029***	-0.034***	-0.041***	-0.022***	-0.015
	(0.024)	(0.011)	(0.010)	(0.011)	(0.008)	(0.009)
Observations	541,184	541,184	541,184	541,184	541,184	541,184
R^2	0.22	0.37	0.40	0.80	0.81	0.82
Controls						
Year	Yes	Yes	Yes	Yes		
Municipality		Yes				
Postal Code			Yes			
Address				Yes	Yes	Yes
Municipality × Year					Yes	
Postal Code × Year						Yes

Standard-errors clustered at postal code in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

4.2 Heterogeneity to distance and height of turbines

In the baseline dummy specifications above, the relationship between house prices and turbines is grossly simplified. House prices are more realistically modelled as a function of distance to and size of turbines. We test with fewer restrictions on the functional form, how distance to the closest turbine, the height of a turbine, and the interaction of both measures affect house prices.

Figure 3: Effects on log(price) of nearest turbine in 500m bins

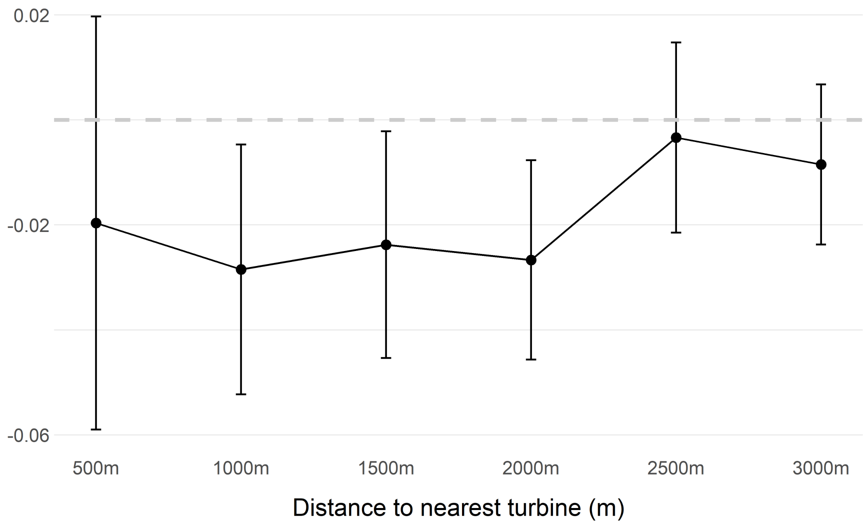

Figure 4: Effects on log(price) of nearest turbine in 1000m bins

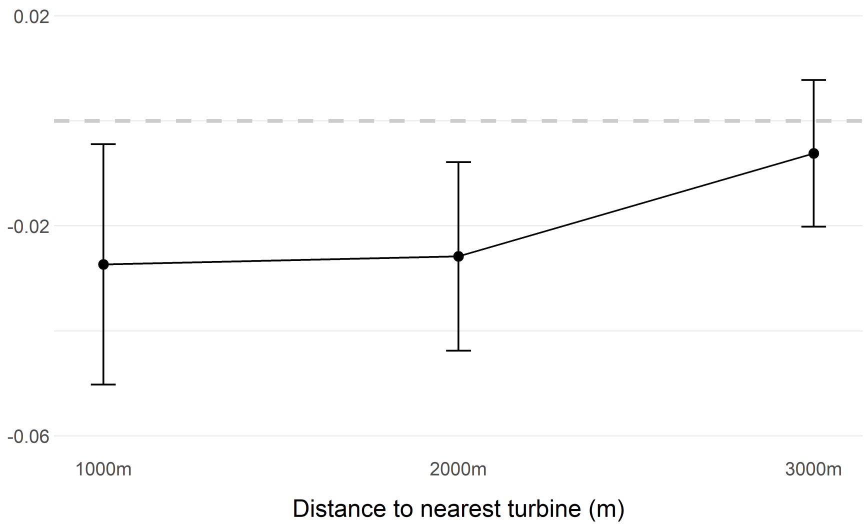

Figure 3 shows the impact of a turbine where the distance is binned in 500m intervals and the x-axis depicting the upper limit. Very few turbines are installed within 500m of a house due to minimum distance legislation and no significant effect is found. Above 500m distance to the turbine, house prices decrease by just over 2% and at conventional statistical significance. The effect drops sharply off above 2km distance, after which no impact of the turbine is detectable. If we bin the distances with 1 km intervals as in figure 4, both the 0-1000m and the 1000-2000m turbines have a significantly negative impact, while the 2000-3000m dummy is still not significant even with a larger number of observations. The distance specification in the baseline estimation is thus supported by the data. The threshold of 2 km is also consistent with Dröes and Koster 2016, while other studies suggest much farther reach (e.g., Gibbons 2015; Zou 2017).

Figure 5: Effects on log(price) of nearest turbine below 2 km - by height

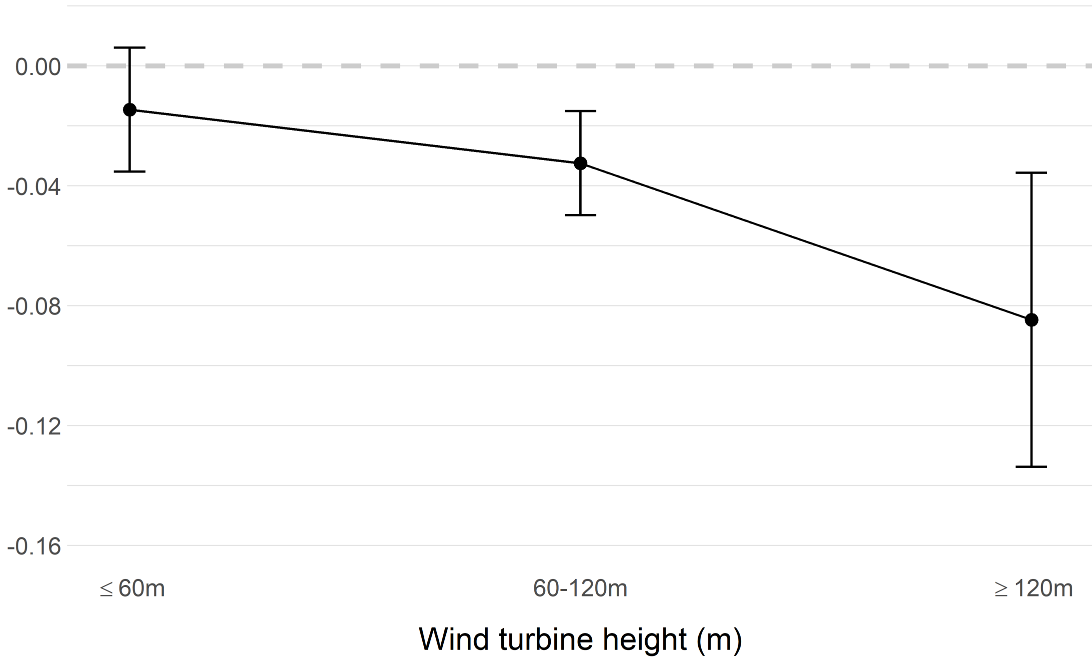

The impact on house prices is plausibly moderated by the height of turbines as well. Visual and noise nuisances increase with the size of the installation and consequently we expect the marginal effect on house prices to be of larger magnitude. Figure 5 depicts estimates separated by three height categories of turbines. The impact of the smallest turbines of at most 60m height is negative with a point estimate below 2% that turns out statistically insignificant. For the two larger size categories, estimates turn significant. The effect estimated for medium sized turbines with heights between 60m and 120m is around -3%. The tallest turbines of 120m and more decrease house prices by over 8%. These estimates for the tall turbines are considerably larger than the 3.7% price decrease that Dröes and Koster 2016 find for turbines above 100m. The discrepancy is testament to the increasing height of newer turbines that make up for a larger share in our more recent data and it highlights the importance of recognizing differential effects of turbine heights.

As much as turbine height contributes to larger house price effects, salience of the visual and noise nuisance is arguably a combination of distance and height. Taller turbines tend to be louder and thus spread noise further, and the perceived visual height is moderated by the distance to it. In Figure 6, we show results for separate turbine heights by distance bins. The largest impact is found for tall turbines at 500-1,000m distance (shorter distance to buildings not allowed) and estimated to be -12%. Between 1,000m and 2,000, the impact is -8%. For medium sized and tall turbines, the estimate reduces with distance, while all estimates are small and insignificant for the smallest turbines. The only turbine that still have a negative impact at 2,500m distance, although just not statistically significant, are of the tallest category. At 3,000m distance, estimates for all turbine heights are essentially zero.

Figure 6: Effects on log(price) of nearest turbine - by height and distance

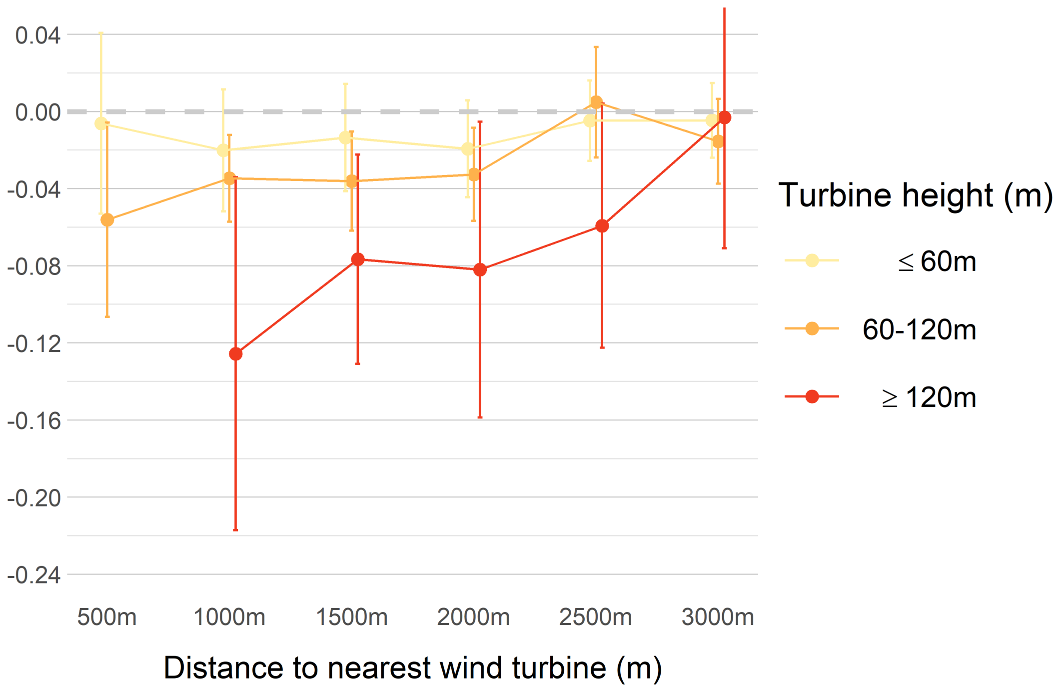

4.3 Effect of shadow flicker

Turbines are not only responsible for noise pollution and unpleasant views of the landscape. Adding to these nuisances, the periodic shadows from rotating wind turbine blades in particular are associated with annoyance (Voicescu et al. 2016) and seizures from photosensitive epilepsy (G. Harding, P. Harding, and Wilkins 2008). Table 3 presents estimates for the effect of the shadow flicker indicator on house prices. Results in column (1) do not account for turbine proximity and are thus likely to be overstated. When we include the dummy for a turbine within 2km in column (2), resembling the baseline specification from Table 2 panel A column (5), houses potentially affected by shadow flicker at some time during the year experience an additional price drop of 6.8% that is statistically significant.[4] Consistent with shadow flicker constituting a severe nuisance, the impact on house prices is larger than for having an average turbine within 2km.

As turbine size increases both the proximity effect on house prices and the probability of shadow flicker, we include the three turbine height specifications as controls in column (3) (resembling the Figure 5 specification). Here, we compare houses with the same height of the closest turbine, where the treated houses are in the shadow flicker area and the control houses are not. The estimate of the shadow flicker effect is slightly smaller, at a 5.2% price drop, and only statistically significant at the 10% level. When we include turbine height and distance interactions in column (4) the coefficient of shadow flicker stays the same, but the standard errors increase and the estimate turns statistically insignificant. The precision might be decreased due to a power issue when we compare fewer houses with and without shadow flicker within many small cells.

[4]Results are not directly comparable to others in the literature. Dröes and Koster 2016 do not find a significant effect, which they point out could be due to measurement error in the shadow flicker variable used. Indeed, they use a simplified assumption that only houses in the northern half of the 2km circle of a turbine can be affected by shadow flicker and if they are within a distance of 12 times the rotor diameter.

Table 3: Effect of shadow flicker on house prices for different specifications

Dependent Variable:	ln(Price)			
	(1)	(2)	(3)	(4)
Shadow Flicker	-0.078***	-0.068**	-0.052*	-0.051
	(0.030)	(0.030)	(0.030)	(0.035)
Controls				
Turbine ≤ 2 km		Yes		
Turbine ≤ 2 km interacted with height			Yes	
Turbine distance and height in 500m bins				Yes
Municipality × Year	Yes	Yes	Yes	
Address	Yes	Yes	Yes	Yes
Observations	2,104,742	2,104,742	2,104,742	2,104,742
R^2	0.85	0.85	0.85	0.85

Standard-errors clustered at postal code in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

4.4 Dynamic treatment effects

A wind turbine commissioning that we use as our treatment is the latest step of setting up a turbine. The process begins with a long planning procedure involving local authorities and public hearings and leads to the physical construction of the turbine before it is commissioned. House prices may well incorporate some or all of the turbine nuisance before the commissioning date. We investigate these anticipation effects in an event study design that describes the full dynamic treatment effect. We use only houses that change their treatment status from zero to one for this analysis, where the commissioning year is set to zero for all houses, following estimation equation 2.

Results of the dynamic event study estimation are depicted in Figure 7. It is apparent that none of the anticipation periods before the commissioning shows a significant effect on house prices. All estimates are relative to 8 years and more before the year of commission. Looking at the point estimates reveals that they start to turn negative at four years prior and remain around -2.5%. This is weak evidence of an anticipation effect and suggests that the house price effect of turbine commissioning understates the full impact of turbines.

The event study furthermore is used to investigate whether treatment effects are stable or heterogeneous over time. The first significant treatment effect in Figure 7 appears in the first year after commissioning with an almost 5% price drop. Point estimates thereafter hover around -4%, indicating that treatment effects materialize right after the turbines are commissioned and are permanent.

As sudden changes to the profitability of individual turbines, due to market conditions, policy changes, or mechanical fatigue and accidents, are not easily predicted, the dynamic treatment pattern for decommissioning turbines is expected to show less of an anticipation effect. For a separate identification in the dynamic model according to equation 3, we add houses that change treatment status from one to zero to the estimation sample. Results in Figure 8 show the dynamic treatment

Figure 7: Event study

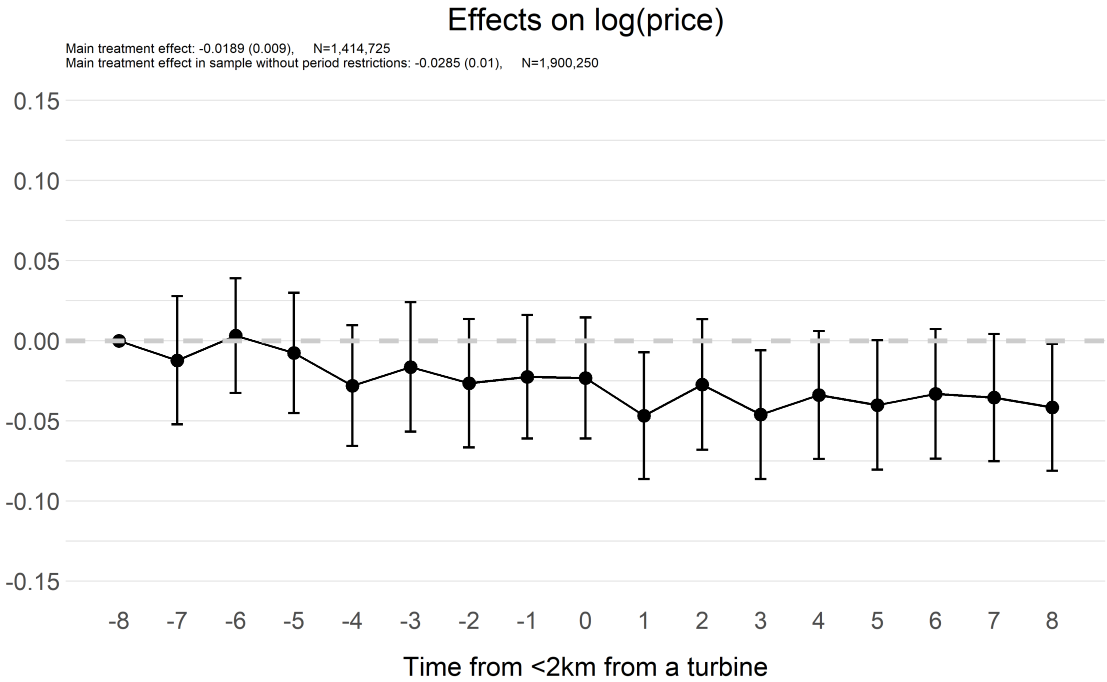

effects for both decommissioning and commissioning of turbines, whereas the latter results only change marginally due to the larger sample compared to Figure 7. The decommissioning treatment, which equivalently should show a positive price effect, indeed reveals little to none anticipation effect. However, there is a considerable lag in the price effect that continuously increases until 7 years after the treatment, when it reaches almost 10%. There are at least two plausible explanations for this pattern. First, several years may pass before the demolition of turbines, while house owners and buyers are still exposed to the view of the turbine and are confronted with uncertainty about re-commissioning. Second, a decommissioning may be followed by other turbines in the area shutting down or other local investments that improve the attractiveness of the area. Another possibility is that the actors overestimate temporarily the impact of turbines on the value of a house.

4.5 Total external costs and cost-benefit analysis

The above analyses gave us the relative price effects from wind turbine proximity. In this section we provide back-of-the-envelope calculations for the extent of the total house value losses for 2019, the last year of our data. The idea is to take the price effects we have estimated from the transactions data and apply these damage estimates to all Danish residential properties in 2019 with their associated 2019 turbine proximity status. To compute the total externalities in monetary terms we need a measure of the value of each house. Official property valuations are available for Denmark but they were last updated in 2011 and are widely considered unreliable.

Instead, we use the 2016-2019 average square meter price of the local area and multiply this with the size of the house to get the house value for all observations.[5] For the vast majority of observations (> 99%) we use the postal code and for the rest we use the municipality prices due to very few observed house transactions (< 50). We then multiply the relative price effect from our preferred specification

[5]Using this approach, the total housing value in Denmark in 2019 is €706.1 billion. Considering house price growth 2013-2019 this figure is very similar to Statistics Denmark's estimate for 2013 of €585.5 billion (Abildgren 2015).

Figure 8: Event study

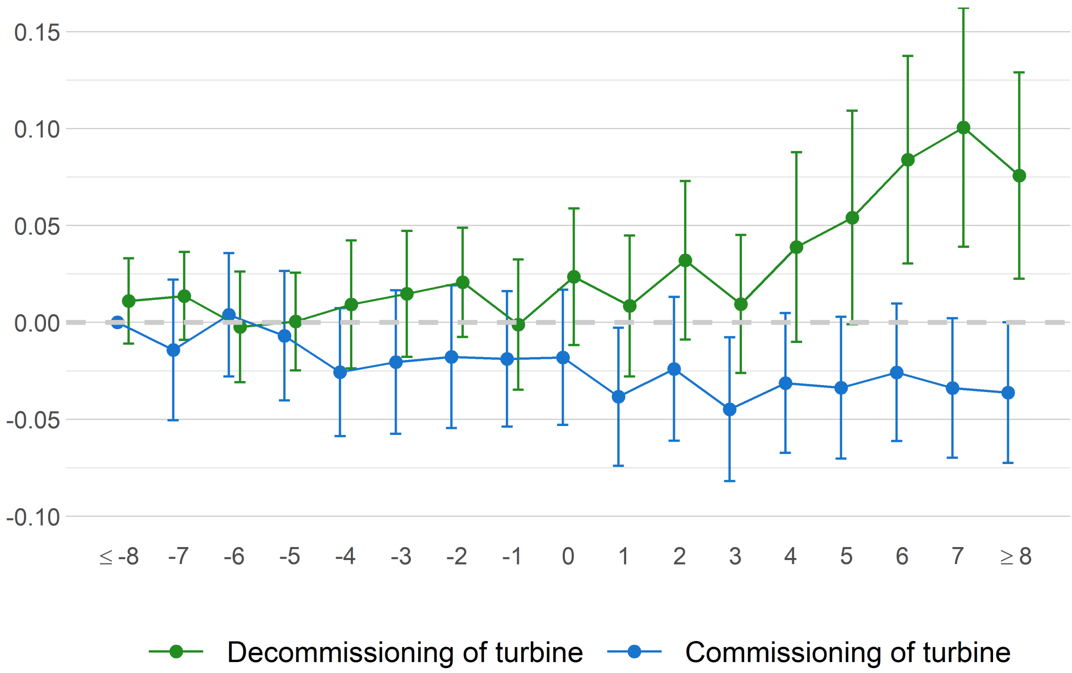

allowing heterogeneity in height of and distance to nearest turbine (Figure 6) with this house value to get the value that is lost due to wind turbine proximity.

Note that we also include the insignificant coefficients from Figure 6. Figure 5 showed a borderline significance of the low turbines so the insignificance from Figure 6 may be due to low power when allowing heterogeneity both in distance and height. Similar considerations apply to the estimate at 2500m for the high turbines (Figure 6) and the coefficient on shadow flicker (Table 3). Including all coefficients makes it more likely that we overestimate the total damages from wind turbines. On the other hand, we do not account for anticipation effects in the main estimations, which would imply an under-estimation of total effects and damages. The total damages found below would be roughly 40% lower if we excluded the insignificant coefficients (see Table 6 in the Appendix).

Table 4 panel A shows the total losses and various summary statistics. The total externality in 2019 is €2,716 million, which amounts to 0.38% of the total 2019 housing value (see footnote 5). The average loss for the 4,091 active onshore turbines in 2019 is €663,000. The average loss per MWh is €238. In both cases, the median loss is markedly lower than the average which indicates that relatively fewer turbines are responsible for substantial house value losses.

Similarly, the average loss per house is €976 while the median is 0. Hence, at least 50 percent of all residential houses suffer no loss from turbine proximity (in fact, the 90th percentile is also 0 so less than ten percent of all houses carry all losses).[6]

The average loss per turbine is relatively similar across turbine heights. Figure 5 and Figure 6 showed that the relative price effect from the highest turbines is more damaging than the medium-sized turbines so this result suggests that the highest turbines are rarely placed close to (high-value) residential areas (see also Table 1). Given the higher capacity of the turbines above 120 meters it is then clear that they are less damaging per MWh they produce: The average loss per high turbine is

[6]Figure 15 in the Appendix shows the spatial variation in average relative price effects, highlighting how damages from onshore wind turbines are distributed unevenly across areas. Areas around urban centers such as Copenhagen face zero losses while other more rural postal codes see their total 2019 housing value diminished by more than two percent.

€86 for the high turbines, €305 for the medium-sized and a staggering €876 for the low turbines below 60 meters.

In panel B we compare the positive environmental externality of reducing carbon dioxide emissions to the negative externality of lessening property values from producing electricity with wind turbines of different sizes. A full cost-benefit analysis would include the reduction in other pollutants and several other costs and benefits as well, but that is beyond the scope of this analysis. The exercise here is simply comparing the value of the displaced carbon dioxide to the destroyed property value. In panel B, we compute the gain per MWh which is given as

$$\text{Gain} = \text{SCC} \times \text{DEF} \times n - \text{Lost housing value}, \tag{4}$$

where SCC is the Social Cost of Carbon, n is the lifetime of the turbine, and DEF is the displacement emissions factor, which is the tonnes of carbon dioxide avoided per additional MWh of wind energy produced. We calculate the gain for three different values of SCC. The low value of €50 is, at the time of writing in May 2021, the current market price in the EU Emissions Trading System (EU ETS). The high cost of €200 is the current recommendation of the Danish Climate Council. Over time, the estimates of the SCC have been increasing. We set the DEF at 0.69 which is currently the best estimate for Denmark (Christensen, Datta Gupta, and Santucci de Magistris 2021). For lost housing values we use the average loss per MWh from panel A, and we set $n = 20$.

Table 4 panel B shows a positive average gain of producing one MWh with onshore wind energy for all configurations of the SCC. It is also seen that estimates differ by turbine height and are positive for all turbines except for the low configuration of the SCC, where it is negative for low turbines. Hence, in this simplified comparison, the benefits of producing electricity with onshore wind turbines exceeds the costs for all combinations of SCC and turbine heights except the one that has both low SCC and low turbines.

5 Conclusion

We have shown that wind turbines inflict significant damage on the value of local properties up to 2.5km away. The impact increases with the height of the turbine, such that more modern tall turbines suppress housing values more heavily. Houses within the area where turbines produce shadow flicker with their rotors suffer an additional drop in value. While the house price effects are significant both in a statistical and an economic sense, wind turbines, especially newer tall versions, overcompensate their damages with savings in carbon dioxide emissions when social costs of carbon are assumed at any conventional level.

For policy purposes, our results have a number of implications. First, to fully compensate property owners for their losses, at least three indicators, distance, turbine height, and shadow flicker, have to be taken into account. Second, turbines produce a considerable social net benefit. Thus, expanding wind farms is socially beneficial even if more houses than currently are affected from new turbines. Third, tall turbines with higher efficiency should be preferred even if their marginal effect on house prices is larger.

Table 4: Overall losses in house values and cost-benefit analysis (2020 euros)

	All turbines N=4,091	≤ 60m N=884	60-120m N=2,561	≥ 120m N=646
Panel A: Losses				
Total loss (millions)	2,716	594	1,670	452
Average loss per turbine	663,000	672,000	651,000	699,000
Median loss per turbine	119,000	128,000	121,000	101,000
Average loss per MWh	238	876	305	86
Median loss per MWh	58	177	59	13
Average loss per house	976			
Median loss per house	0			
Panel B: Cost-benefit analysis of producing one MWh with onshore turbines				
High SCC (200)	2522	1884	2455	2674
Medium SCC (125)	1487	849	1420	1639
Low SCC (50)	452	-186	385	604

Notes: Calculations are based on the specification in Figure 6. There are 4,091 active onshore turbines in 2019. We assume a capacity factor of 30% so a wind turbine of 1 MWh delivers $365 \times 24 \times 0.30 = 2,628$ MWh per year.

6 Appendix

6.1 Selected results with alternative sample

Figure 9: Effects on log(price) of nearest turbine in 500m bins

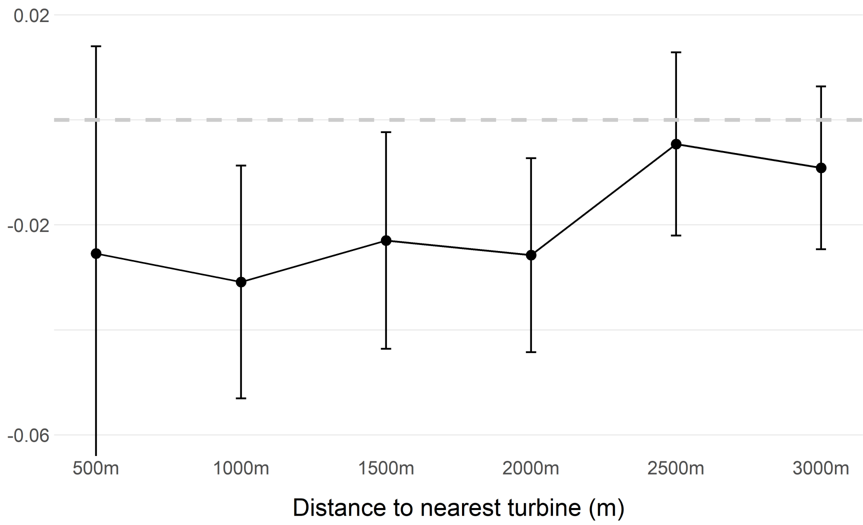

Figure 10: Effects on log(price) of nearest turbine in 1000m bins

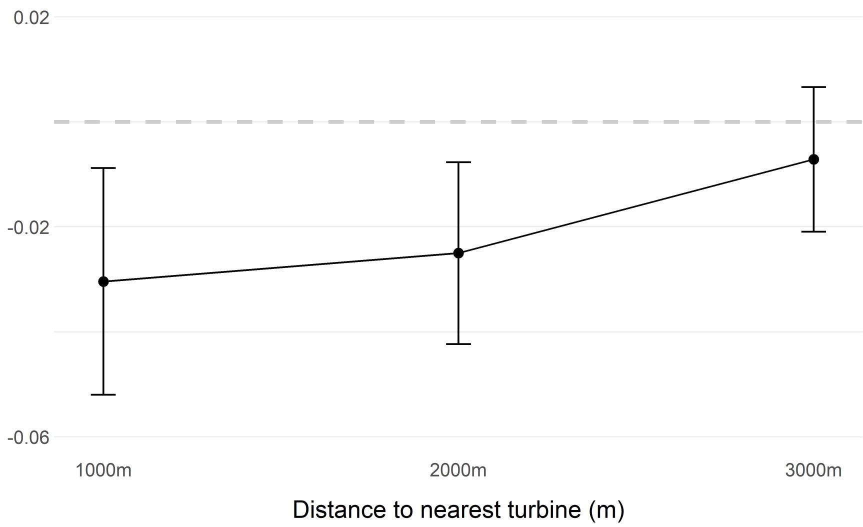

Figure 11: Effects on log(price) of nearest turbine - by height and distance

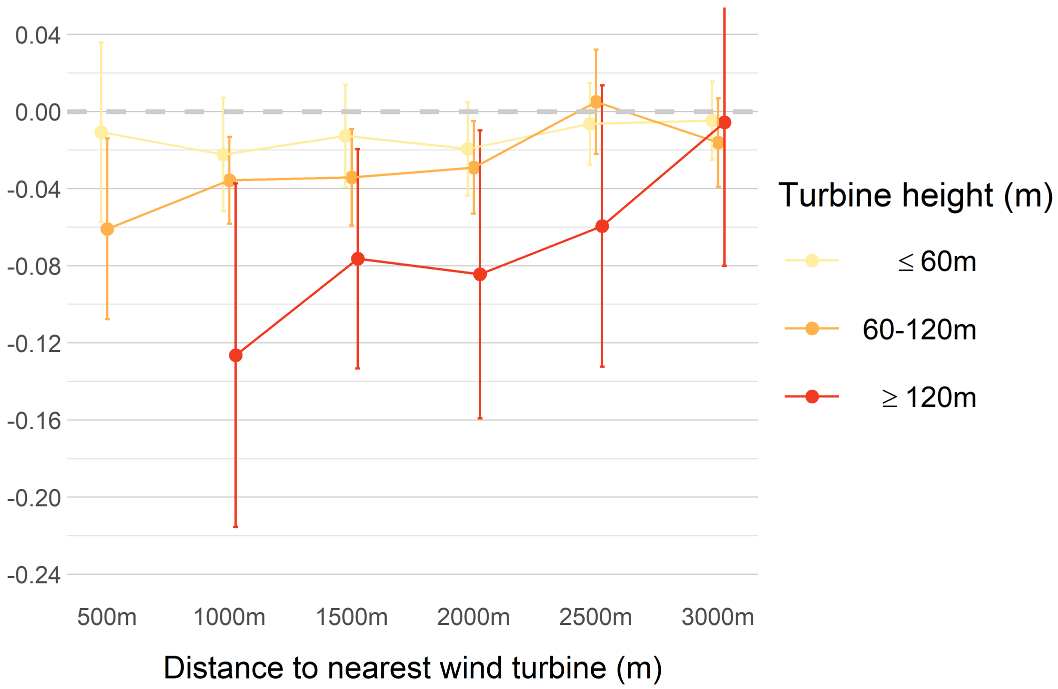

Figure 12: Effects on log(price) of nearest turbine below 2 km - by height

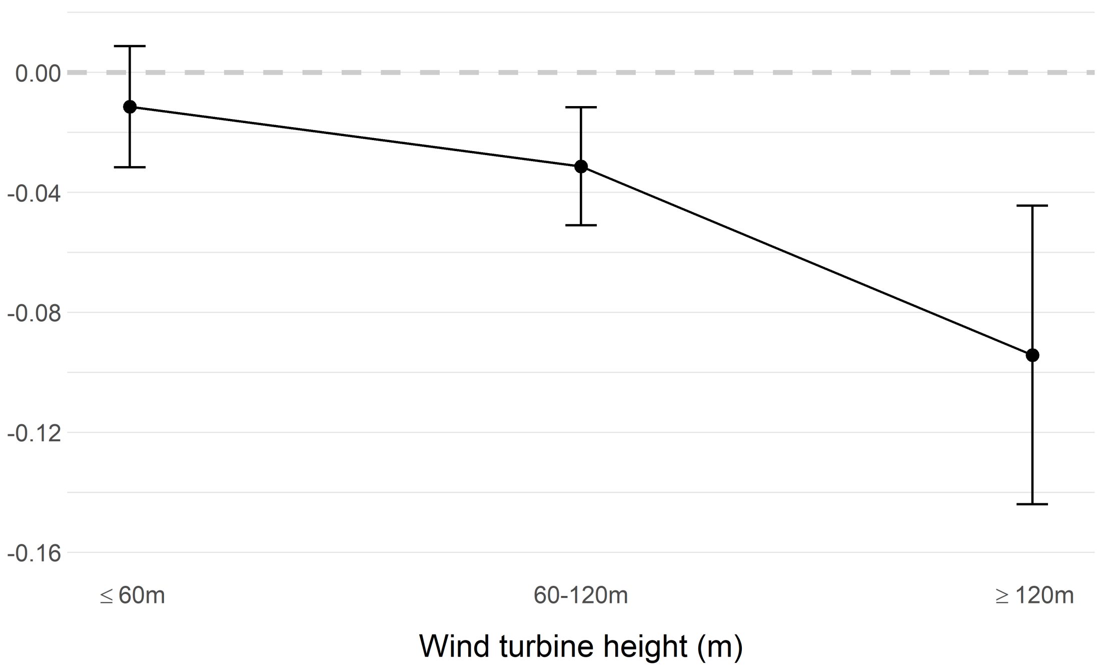

Figure 13: Effects on log(price) of nearest turbine - by height and distance

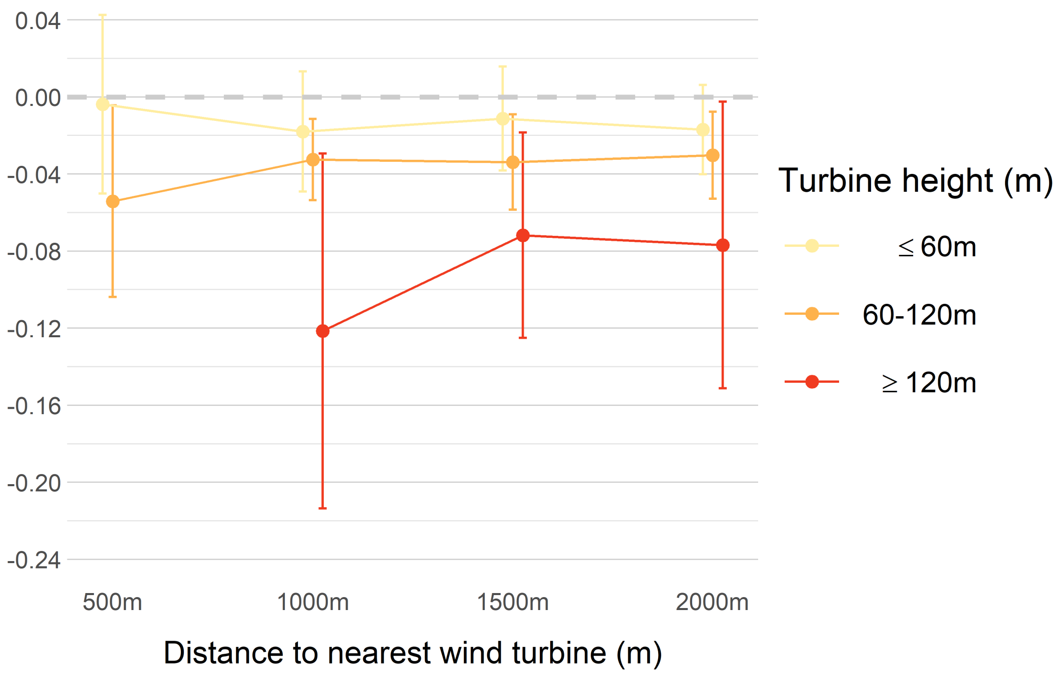

Table 5: Effect of shadow flicker on house prices for different specifications

Dependent Variable:	ln(Price)				
	(1)	(2)	(3)	(4)	(5)
Shadow Flicker	-0.078***	-0.069**	-0.052*	-0.051	-0.044
	(0.028)	(0.029)	(0.030)	(0.034)	(0.033)
Controls					
Turbine ≤ 2 km		Yes			
Turbine ≤ 2 km interacted with height			Yes		
Turbine distance and height in 500m bins				Yes	
Turbine distance and height in 250m bins					Yes
Municipality × Year	Yes	Yes	Yes	Yes	Yes
Address	Yes	Yes	Yes	Yes	Yes
Observations	541,184	541,184	541,184	541,184	541,184
R^2	0.81516	0.81519	0.81522	0.81522	0.81525

Standard-errors clustered at postal code in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Figure 14: Event study

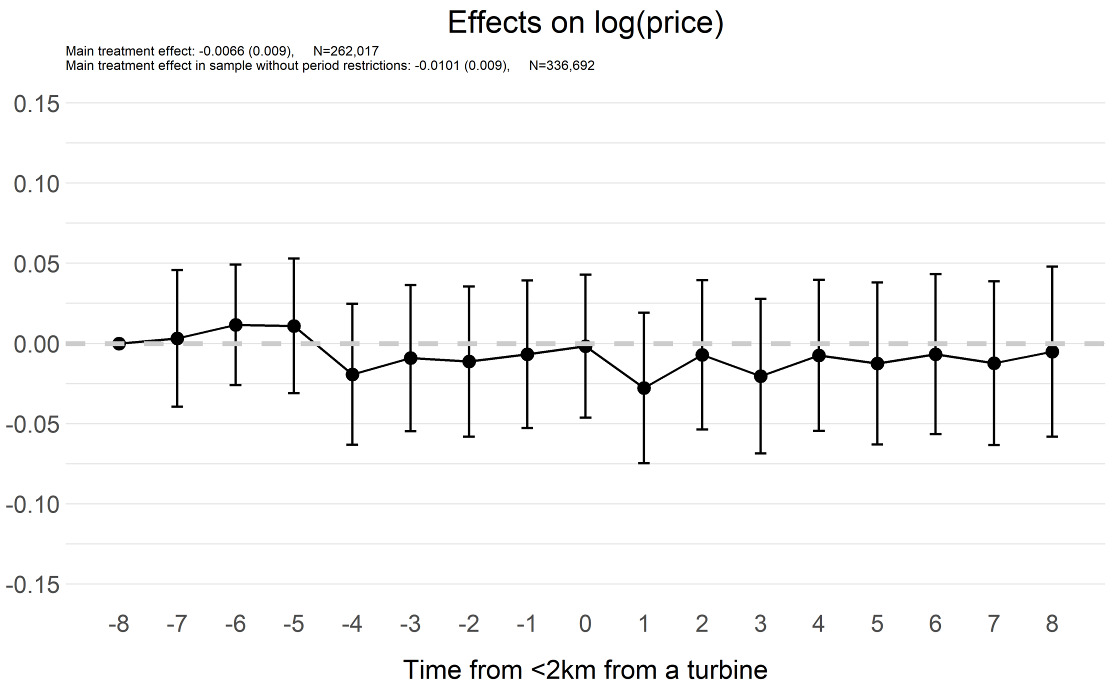

Table 6: Overall losses in house values and cost-benefit analysis (2020 euros)

	All turbines N=4,091	$\leq$ 60m N=884	60-120m N=2,561	$\geq$ 120m N=646
Panel A: Losses				
Total loss (millions)	1,657	0	1,344	313
Average loss per turbine	405,000	0	524,000	484,000
Median loss per turbine	46,000	0	112,000	89,000
Average loss per MWh	145	0	246	60
Median loss per MWh	16	0	54	11
Average loss per house	589			
Median loss per house	0			
Panel B: Cost-benefit analysis of producing one MWh with onshore turbines				
High SCC (200)	2615	2760	2514	2700
Medium SCC (125)	1580	1725	1479	1665
Low SCC (50)	545	690	444	630

Notes: Calculations are based on the specification in Figure 6. Insignificant variables from this regression are not included. There are 4,091 active onshore turbines in 2019. We assume a capacity factor of 30% so a wind turbine of 1 MWh delivers $365 \times 24 \times 0.30 = 2,628$ MWh per year.

Figure 15: Relative price effects by postal code in 2019

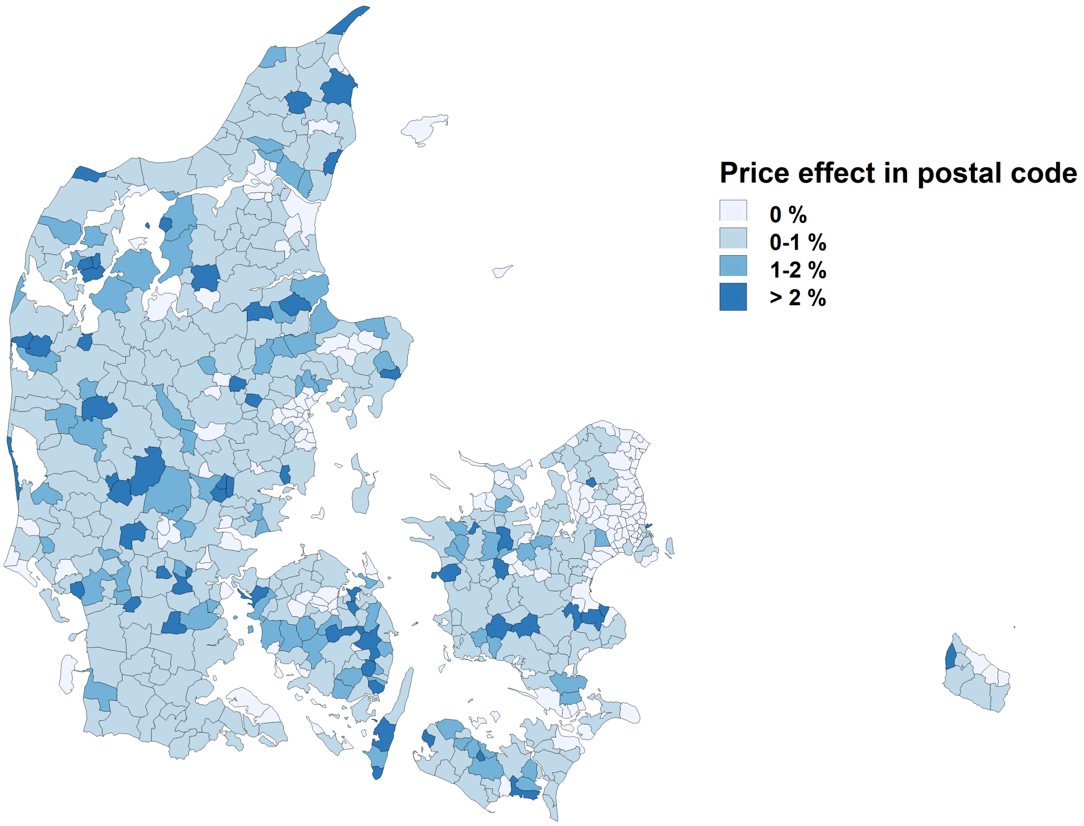

Notes: Figure 15 shows the 2019 average relative price effect by postal code based on the specification in Figure 6.

References

[1] Kim Abildgren. *Estimates of the national wealth of denmark 1845-2013*. Tech. rep. Danmarks Nationalbank Working Papers, 2015.

[2] Bent Jesper Christensen, Nabanita Datta Gupta, and Paolo Santucci de Magistris. "Measuring the impact of clean energy production on CO 2 abatement in Denmark: Upper bound estimation and forecasting". In: *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 184.1 (2021), pp. 118–149.

[3] Danish Energy Agency. *Overview of the Energy Sector*. https://ens.dk/service/statistik-data-noegletal-og-kort/data-oversigt-over-energisektoren. Last checked on May 26, 2021. May 2021. URL: https://ens.dk/service/statistik-data-noegletal-og-kort/data-oversigt-over-energisektoren.

[4] Clément De Chaisemartin and Xavier d'Haultfoeuille. "Two-way fixed effects estimators with heterogeneous treatment effects". In: *American Economic Review* 110.9 (2020), pp. 2964–96.

[5] Martijn I Dröes and Hans RA Koster. "Renewable energy and negative externalities: The effect of wind turbines on house prices". In: *Journal of Urban Economics* 96 (2016), pp. 121–141.

[6] Energinet. *BÅDE VIND OG SOL SLOG REKORD I 2017 – LEVEREDE 45,8 PROCENT AF STRØMMEN*. 2020. URL: https://energinet.dk/Om-nyheder/Nyheder/2018/01/09/Baade-vind-og-sol-slog-rekord-i-2017-leverede-459-procent-af-stroemmen. (accessed: 24.08.2020).

[7] Stephen Gibbons. "Gone with the wind: Valuing the visual impacts of wind turbines through house prices". In: *Journal of Environmental Economics and Management* 72 (2015), pp. 177–196.

[8] Andrew Goodman-Bacon. *Difference-in-differences with variation in treatment timing*. Tech. rep. National Bureau of Economic Research, 2018.

[9] Graham Harding, Pamela Harding, and Arnold Wilkins. "Wind turbines, flicker, and photosensitive epilepsy: Characterizing the flashing that may precipitate seizures and optimizing guidelines to prevent them". In: *Epilepsia* 49.6 (2008), pp. 1095–1098.

[10] Ben Hoen et al. "Wind energy facilities and residential properties: the effect of proximity and view on sales prices". In: *Journal of Real Estate Research* 33.3 (2011), pp. 279–316.

[11] IBRD. *Tracking SDG7–The Energy Progress Report*. Tech. rep. nternational Bank for Reconstruction and Development, 2020.

[12] IEA. *Renewables 2020 - Analysis and forecast to 2025*. Tech. rep. International Energy Agency, 2020.

[13] Cathrine Ulla Jensen, Toke Emil Panduro, and Thomas Hedemark Lundhede. "The vindication of Don Quixote: The impact of noise and visual pollution from wind turbines". In: *Land Economics* 90.4 (2014), pp. 668–682.

[14] Corey Lang, James J Opaluch, and George Sfinarolakis. "The windy city: Property value impacts of wind turbines in an urban setting". In: *Energy Economics* 44 (2014), pp. 413–421.

[15] Aslak Harbo Poulsen and Ole Raaschou-Nielsen. "Short-term nighttime wind turbine noise and cardiovascular events: a nationwide case-crossover study from Denmark". In: *Environment international* 114 (2018), pp. 160–166.

[16] Kurt Schmidheiny and Sebastian Siegloch. “On event studies and distributed-lags in two-way fixed effects models: Identification, equivalence, and generalization”. In: *ZEW-Centre for European Economic Research Discussion Paper* 20-017 (2020).

[17] Yasin Sunak and Reinhard Madlener. “The impact of wind farm visibility on property values: A spatial difference-in-differences analysis”. In: *Energy Economics* 55 (2016), pp. 79–91.

[18] Sonia A Voicescu et al. “Estimating annoyance to calculated wind turbine shadow flicker is improved when variables associated with wind turbine noise exposure are considered”. In: *The Journal of the Acoustical Society of America* 139.3 (2016), pp. 1480–1492.

[19] Richard J Vyn and Ryan M McCullough. “The effects of wind turbines on property values in Ontario: does public perception match empirical evidence?” In: *Canadian Journal of Agricultural Economics/Revue canadienne d’agroeconomie* 62.3 (2014), pp. 365–392.

[20] Alexander Zerrahn. “Wind power and externalities”. In: *Ecological Economics* 141 (2017), pp. 245–260.

[21] Eric Zou. *Wind Turbine Syndrome: The Impact of Wind Farms on Suicide*. Tech. rep. Working Paper, 2017. HPSA SATC Income Climate 36, 2017.