2020-04-28

Wind Turbines and Housing Prices: Valuing the Impact of Wind Farms on Transactions

Nathan Guzman

Honors Thesis

WIND TURBINES AND HOUSING PRICES:

VALUING THE IMPACT OF WIND FARMS ON TRANSACTIONS

by
Nathan Guzman

Submitted to Brigham Young University in partial fulfillment
of graduation requirements for University Honors

Economics Department
Brigham Young University
April 2020

Advisor: Jaren Pope

Honors Coordinator: John Stovall

ABSTRACT

WIND TURBINES AND HOUSING PRICES:
VALUING THE IMPACT OF WIND FARMS ON TRANSACTIONS

Nathan Guzman

Department of Economics

Bachelor of Science

As the number of wind farms worldwide increases, more and more homeowners will live near them. Some homeowners have found these wind turbines generate noise, obstruct views, and cast shadows on their property. Using data from 10 years of transactions in McLean County, Illinois, I investigate if a wind farm near a neighborhood impacts housing prices. After estimating a simple difference-in-differences model, I find that wind farms may have had a negative effect on home prices. I then discuss how this paper fits in with existing literature and provide suggestions about how further research could proceed.

TABLE OF CONTENTS

LIST OF TABLES AND FIGURES

I. Introduction

The share of renewable energy generated by wind farms worldwide is expected to increase from 8% in 2010 to 24% in 2035 (IEA 2012).[1] As they currently stand, wind turbines occupy large amounts of land, create noise, and cast shadows on the nearby area. As a result, many new farms have not only produced clean, dependable energy, but also generated complaints (Hardy and Eller 2017). While current scholarship focuses on the efficiency, implementation, and long-term viability of wind turbines as alternative energy sources, only a small contingent of scholars have examined how these turbines affect the people who live near them.

Economic theory suggests that the noise, view obstruction, and shadows caused by turbines could create negative externalities for homeowners. Consumers would likely object to these externalities, causing the home's transaction price to drop. This price decrease could originate with either the buyer or the seller. Buyers could decide to purchase fewer homes near wind farms, which would decrease demand and lower housing prices. Homeowners could also respond to the construction of wind turbines by putting their houses on the market, increasing supply and also lowering the market price.

In an alternative case, some buyers or sellers might encounter positive externalities when dealing with turbines. Some wind farms are built on private land, which means the company managing farm leases the land from its owners. Depending on the value of the lease, owning one could boost demand and potentially increase a home's

[1] While these projections do an excellent job of showing how much we will depend on wind energy, they do little to explain what these future farms will look like. Technological advancements will likely reduce a wind turbine's footprint, which could potentially speed adoption and mitigate any negative externalities. That being said, this paper approaches the relationship between wind farms and home prices as they currently stand.

price. The lease in this situation would not be an externality but contribute directly to the home's value. In another scenario, buyers and sellers could find the windmills attractive or associate them with the production of green energy, which could also increase prices. In the case where there is no price change, then wind farms are either not a major issue for buyers or sellers, or the positives and negatives cancel each other out.

Regardless of whether buyers or sellers drive the process, externalities originating from wind farms may manifest themselves in transactions. This paper quantifies these externalities through a quasi-experimental hedonic analysis using historical data from McLean County, Illinois. By way of this analysis, I find some evidence that wind farms negatively impact home prices. In the case of Twin Groves, the farm I focus on, its announcement may have affected prices more than the farm itself, which suggests that residents might have adjusted to the farm during its construction.

In order to come to this conclusion, I estimated a simple difference-in-differences model focusing on Twin Groves, a 240 turbine wind farm in McLean County. When built, Twin Groves was the largest wind farm east of the Mississippi and provided a significant shock to a county that had never invested in wind energy. I used a difference-in-differences model for my analysis because I could define treatment and control groups based on distance from a turbine. I had locational data on all the turbines in the farm and panel data for home purchases in the county from 2005 to 2018. Assuming the farm's construction was an exogenous shock (there was some uncertainty about when and if the turbines would be built), I used that data to compare the farm's effect on nearby homes to trends among more distant properties.

The regressions suggest there may be a negative relationship between the wind farm's construction and prices, although this relationship is not statistically significant. The relationship is strongest when focusing on the farm's announcement and persists even when varying the timeframe used in the analysis. A plausible explanation for this result is that residents objected to the project when it was introduced, but as more details about the farm surfaced—such as payments from EDP Renewables to property owners to lease land for turbines—those objections diminished. That being said, most of my coefficients are not statistically significant, meaning I cannot have statistical confidence that the true relationship is different than zero.

II. Literature Review

To date, a small number of scholars have explored this topic. Researchers have pursued this question using two primary methodologies: hedonic pricing models and choice experiments. Hedonic models have produced varying results. Dröes and Koster (2016) identified a 1.4% price decrease for houses within 2km of wind turbines in the Netherlands. Using a dataset of Danish households, Jensen, Panduro, and Lundhede (2014) found between a 7.3% and 14% depression in property price within a distance of 0.5 miles to the nearest turbine. While European studies tend to associate negative externalities with wind farms, others—specifically studies in the U.S.—have found no statistically significant change in housing prices after constructing a wind farm (Lang, Opaluch, and Sfinarokalis 2014). Castleberry and Greene (2018), for instance, find only isolated incidences of lower housing values.

Each of the above studies uses a difference-in-differences model to estimate the average treatment effect after the first wind turbine is constructed near a home. Because their methodologies are consistent, the different results suggest the populations they studied could have different preferences. In other words, European homeowners may see farms through different eyes than Americans.

Choice experiments also suggest heterogenous preferences among populations. European economists Meyerhoff, Ohl, and Hartje (2010) surveyed German respondents who lived near wind farms to try and determine their willingness to pay for wind power, concluding that “preferences vary significantly among respondents resulting in three segments with different preference structures.” Each of these preference structures depended on distance, with advocates receiving small benefits from moving turbines farther away from their homes, opponents experiencing great gains in utility from more distant turbines, and a middle group receiving moderate benefits from increased distance between themselves and wind farms. Brennan and Van Rensburg (2016) performed a similar study in Ireland and found that, while “the majority of respondents are willing to make (monetary) tradeoffs to allow for wind power initiatives,” the amount of compensation respondents required could be manipulated by changing the way the wind farm was implemented.

When taken together, the literature surrounding wind farms suggests they can, depending on inhabitants’ preferences, either create negative externalities or leave inhabitants indifferent. While this may be the case, the regional difference could come from how the project is implemented. If homeowners who live close to wind farms receive tax benefits, for example, their prices might not be affected by the turbines. For the purposes of my study, I focused on McLean County, Illinois.

III. Data Description

This paper uses housing data from McLean County, Illinois to perform its analysis. I chose McLean County because of the Twin Groves wind farm. The Twin Groves Project contains 218 turbines and finished on February 2nd, 2008. The farm was built to provide power to the Chicago metropolitan area. At the time of its construction, it was the largest wind farm east of the Mississippi. The farm was built over 19 months, a relatively short amount of time considering its size. The farm was also built during discrete timeframes, which allowed me to establish clear treatment areas and test the hypothesis of a negative impact on housing prices using precise pre- and post- periods. I put together a rough timeline of the Twin Groves farm's construction by searching articles from the Pantagraph, McLean county's local news outlet.[2] First, the county authorized the project on September 20th, 2005. Many residents opposed the farm, so they sued EDP Renewables. The ensuing legal battle pushed the construction start date to June 29th, 2006. EDP built the first tower on September 28th, 2006 and finished the first wind turbine on October 19th, 2006. The farm became operational on March 27th, 2007, and construction finished on February 2nd, 2008. Figure 1 gives a visual representation of the project's timeline. While considering all the dates, I chose to focus on October 19th, 2006—the day the first turbine was finished—as the cutoff for my pre- and post- periods. I chose October to accommodate my dataset's timeframe. I wanted 2 years of data before and after whatever date I chose to provide a symmetric window for my analysis, but my

[2] The Pantagraph has been McLean County's local news source since 1837. Its website contains all articles published since 2005. To find the dates I used in my regression, I searched for articles covering the Twin Groves wind farm and checked for milestones, including when construction started, when the first turbine was built, and when it started producing electricity. The Pantagraph can be accessed here: www.pantagraph.com

data begins in January of 2005. October 19th, 2006 almost creates a 2 year symmetric window. Additionally, the first turbine sent a strong message to consumers. It confirmed that the project would finish, and it created the first potential visual externality. As a result, I thought the October date would be the best for creating pre- and post- groups.

There were no wind turbines in the county at the time of the farm's construction, so the farm created a natural shock in McLean County's housing market. In order to measure this shock, I combined housing data with census tract data from the Census Bureau. I then added a layer of wind farm data, which came from the United States Wind Turbine Database: a collection of locations, construction dates, and other turbine information gathered by the U.S. Geological Survey. Map 1, which marks its turbines in purple, shows that although the project is far away from Bloomington—the county seat and large cluster of homes on the western side of the county—there are still a significant number of homes near its turbines. In fact, entire townships—Arrowsmith and Saybrook specifically—fit inside.

Another wind farm—White Oak—was built in McLean County around 2011. I could have used that wind farm for my study, but I had limited information on its construction timeline. Both projects had significant numbers of houses nearby. From 2005 to 2018, 70 transactions occurred within 1 mile of the White Oak project and 183 happened the same distance from the Twin Groves project. Although White Oak was closer to Bloomington than Twin Groves, fewer transactions occurred within 1 mile of White Oak than Twin Groves. White Oak was also the second to be built in the county, which means its introduction was less of a shock than Twin Groves. For these reasons I decided to focus on Twin Groves for my analysis.

The McLean dataset includes longitude and latitude numbers, transaction information, and general metrics for each property sold between 2005 and 2019. Before running my analysis, I looked for outliers, focusing on home price. In order to keep outliers from skewing my regressions, I dropped any observations that sold for less than \$25,000 or more than \$500,000.[3] Table 1 shows that the average lot without these outliers was 1.85 acres and had an 1,800.27 square foot building. Homes in McLean County sold for \$157,216.20 and were 9.74 miles from a wind turbine on average. Homes 0-1 mile from a turbine sold for \$132,894.90 on average. The properties were generally large, averaging 20.74 acres, and old, about 83.5 years. In contrast, lots 1-3 miles from a turbine were 12.13 acres on average and 74.64 years old. Despite those differences, prices were fairly similar; the average home sold for \$132,376.30.

A simple comparison of means suggests that Twin Groves had a negative effect on housing prices. Table 2 shows that the average price for homes 0-1 miles from a turbine decreased \$30,355 after October 19th, 2006, while the average price for homes 1-3 miles from a turbine increased \$32,981 in the post- period. The difference in these differences is -\$63,336.20, suggesting the wind turbines had a large negative effect on home prices. While this simple comparison of average housing prices in the treatment and control groups before and after the first turbine was built is suggestive, it does not formally control for differences in housing characteristics, nor does it provide us with statistical significance of the difference. For this, we proceed to a regression analysis.

[3] I chose those values based on what other researchers had done. Most of the papers I have cited restricted their data to minimize outliers. Homes that sold for less than \$25,000 were considered to not be in livable condition, while homes that sold for more than \$500,000 were considered “influential” in the sense that they would skew OLS estimates.

IV. Identification Strategy

In order to analyze the data, I constructed a difference-in-differences model that focuses on estimating the average treatment effect after the first turbine was built within *d* miles of a home. I defined the price of the property *i* in year *t* as p_{it}. I let f_i define the treatment group and x_t determine whether or not the transaction occurred after the first turbine was built. I also defined $f_i * x_t$ as the interaction variable between f_i and x_t. I let h_{it} be a set of housing characteristics including the lot size, building size, year built, and number of bathrooms. I implemented the model in the following format:

$$\ln(p_{it}) = \beta + \alpha(f_i * x_t) + \gamma f_i + \delta x_t + \mu h_{it} + \epsilon_{it} \tag{1}$$

where α captures the average treatment effect, β is the intercept, and ϵ_{it} is the error term. For the purposes of my regressions, *d* = 1. I restricted my sample to homes that were within 3 miles of a turbine to create a representative control group. This meant that for all my regressions, my treatment group contained homes 0-1 mile from a turbine and my control group included homes 1-3 miles from a turbine. I chose these distances so I could create two similar, relevant samples. McLean County is a heterogenous housing market: Bloomington is comparatively urban, while the area around Twin Groves holds townships and large farms. If I had included all or most of the county in my sample, my control group would have been most similar to Bloomington, while my treatment group would have looked more like Arrowsmith. If I left a gap between my treatment and control groups, I would have likely gotten a piece of Bloomington, which would have affected my control group. While other distances could have included more observations in my regressions, in the interest of accuracy and academic honesty I decided on the 0-1 mile and 1-3 mile bands.

The homes in the treatment group lie in the pink region in Map 1, while those in the control group sit in the lavender region. My first regression restricted observations to a 4-year window, 2 years before and after the first turbine was built, while the second used a 6-year window, 3 years before and after the first turbine was built. As explained before, I chose this first window because it was symmetric and centered on a date relevant to my analysis. I chose the second to add more observations into my regression. These specifications allowed me to estimate the effect of building a turbine for those houses most affected by the farm.

For this identification strategy to be valid, I assumed the treatment and control groups had the same trends in prices prior to the first turbine being built. This seemed reasonable given my treatment and control groups were very close together. In order to test this assumption, I created an event study with my data using the October date as my cutoff. Graph 1 shows the coefficients for the cohorts before and after the first turbine was built. Each cohort captured 6 months of data, with the 0 cohort capturing all transactions that occurred between 0 and 6 months of the turbine being built, the 1 cohort capturing all transactions that occurred between 6 and 12 months of the turbine being built, and so on. The graph supports the parallel trends assumption, showing that the differences in price between the treatment and control groups in the pre-period were fairly constant, while thy diverge in the post- period. Like my difference in means, my event study suggests the wind farm negatively impacted prices.

V. Results

Table 3 contains the coefficients from the first regression, which tests the treatment effect for homes within 1 mile of the Twin Groves farm using a 4-year window. According to the table, homes within 1 mile of Twin Groves sold for 19.8% less after the first turbine was constructed. While this term suggests a large decrease in home prices, it is not statistically significant. The 6-year window gave a more extreme estimate. Column 2 of Table 1 shows that homes within 1 mile of the wind farm sold for 32.1% less after the first turbine was built. Like in column 1, the First Turbine Built*One Mile coefficient is not statistically significant.

Large point estimates and a lack of statistical significance suggest that my identification strategy does not have adequate statistical power. Unfortunately, the first regression only has 53 useable observations, while the second only has 67 useable observations. The standard errors for the difference-in-differences estimate in both regressions are large, meaning I cannot statistically rule out the possibility of the wind farm having a zero, or even a positive effect on home prices.

Both coefficients are worth comparing. First, both are fairly large and negative. This shows that, while neither regression identifies a statistically significant relationship, the negative correlation between the farm's construction and home prices persists even when including homes sold more than two years after the farm finished construction. This supports the hypothesis that the farm did in fact impact home prices. Second, there is a significant gap between the coefficients. The first estimate is more conservative, while the second is 12% larger. This could be evidence that the farm continued to decrease prices as time went on. However, as previously mentioned my regressions are under-powered, which means a discussion about the differences in the coefficients is mostly speculative. It is possible that this 12% difference appears through chance variation in the data. In fact, further analysis suggests that Twin Groves may have had the greatest effect on home prices when it was announced.

To test this idea, I ran the same regression as before, but I used September 28th, 2005 as the cutoff for the post- period.[4] As shown in Column 3 of Table 3, The Project Approval*One Mile coefficient was -0.38 and was significant at the $\alpha = 0.1$ level. While this isn't close to conclusive evidence—the pre- period has very few observations and the post- period includes the effects of construction—this coefficient is both the largest of the three and the only to be statistically significant. This suggests that negative externalities could have started to appear when the farm was announced. Records from the Pantagraph also support this idea. An article written in May of 2006 reports that a Bloomington attorney filed a lawsuit against Horizon Energy, the builder, in an attempt to delay construction. Another article, written later in May, reports that construction was delayed until June 2006. The farm seems to have been a contentious issue with uncertainty about when and if it would be constructed, which suggests that home prices could have decreased the most after it was formally announced.

This change in prices could be related to the types of properties that sold during the observed period. McLean County is mostly farmland; the most populated part of the county is Bloomington, which sits several miles west of Twin Groves. Of all the covariates in each regression, only one—Lot Size—was statistically significant at the $\alpha = 0.01$ level. According to table 1, the average lot size for transactions within 1 mile of Twin Groves was 20.74 acres, considerably larger than most residential properties. These two facts suggest that buyers focused more on the land they were buying than the house

[4] Note that this is the date the county approved the farm and the Pantagraph announced its construction.

on it. Because the average lot size was so large, it's likely that several properties contained more than one wind turbine. Farmers probably realized this as soon as Twin Groves was announced, which likely complicated purchasing decisions. These farmers, most of which were unfamiliar with wind turbines, probably considered placing wind turbines on their land with distaste.

In order to placate residents, EDP Renewables signed leasing agreements with farmers. These agreements provided financial incentives for farmers to allow EDP to build on their property; as of now, EDP Renewables has paid more than $27.7 million to landowners.[5] These payments could have diminished the initial effect of announcing Twin Groves. An analysis that includes this leasing information could test the effect of different contract sizes on housing transactions and could find if the contracts reduced the impact Twin Groves had on home prices.

All that being said, the standard deviations for all three coefficients are large. As a result, I cannot be statistically confident of the validity of my negative estimates. Taken together, the three negative coefficients provide some evidence that Twin Groves may have had a negative effect on home prices. In order for me to find the effect size, I would need more observations that are both geographically close to the wind farm and temporally proximate to the October construction date.

VI. Conclusion

This paper examines the effect of the Twin Groves wind farm on home prices. Using housing data from McLean County, I estimate a difference-in-differences model

---

[5] See https/twingroveswindfarm.com for more information.

that focuses on houses within 1 mile of a turbine compared to houses a little farther away, before and after the first wind turbine was built. The resulting regressions show that Twin Groves may have had a negative effect on home prices; however, they do not give me the statistical confidence to say that the effects are different from zero. While there is some evidence of the effect's direction, the small sample size precludes me from having statistical confidence in the true effect size. Further research on Twin Groves should try and include data from more transactions, especially in the pre- period. Adding leasing information to the analysis could also allow researchers to examine the effect of financial incentives on externalities.

While writing this paper I learned about the importance of precision in research. My limited sample taught me that while I could influence the numbers, the credibility of my paper depended on my academic honesty. As a result, I worked to structure my methodology in a way that would best identify the relationship between wind turbines and housing prices, regardless of whether or not it was statistically significant. Through this process, I learned to carefully consider all the details that went into my identification strategy and later my paper. While there were frustrations along the way, I grew to appreciate the rigor inherent to academia, and I'm grateful for the opportunity to have written an Honors Thesis.

Reference List

Brennan, N., & Van Rensburg, T. M. (2016). Wind farm externalities and public preferences for community consultation in Ireland: A discrete choices approach. *Energy Policy, 94,* 355-365.

Castleberry, B. & Greene, J.S. (2018). Wind power and real estate prices in Oklahoma. *International Journal of Housing Markets and Analysis, 11, 5,* 808-827.

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.

Hardy, K. & Eller, D. (2017). 6 common complaints against Iowa wind turbines. Retrieved from https://www.desmoinesregister.com/story/money/business/2017/04/21/6-common-complaints-against-iowa-wind-turbines/100706178/

IEA (2012). World Energy Outlook 2012, Chapter 7, Renewable Energy Outlook. OECD/IEA, Paris, France.

Jensen, C., Panduro, T., & Lundhede, T. (2014). The Vindication of Don Quixote: The Impact of Noise and Visual Pollution from Wind Turbines. *Land Economics, 90, 4,* 668-682.

Lang, C., Opaluch, J., & Sfinarokalis, G. (2014). The windy city: Property value impacts of wind turbines in an urban setting. *Energy Economics, 44,* 413-421.

Meyerhoff, J., Ohl, C., & Hartje V. (2010). Landscape externalities from onshore wind power. *Energy Policy, 38, 1,* 82-92.

Appendix

Figure 1

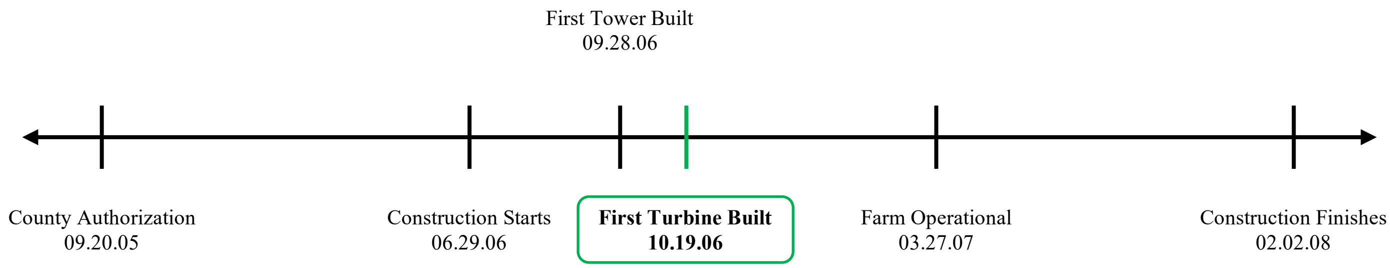

Graph 1

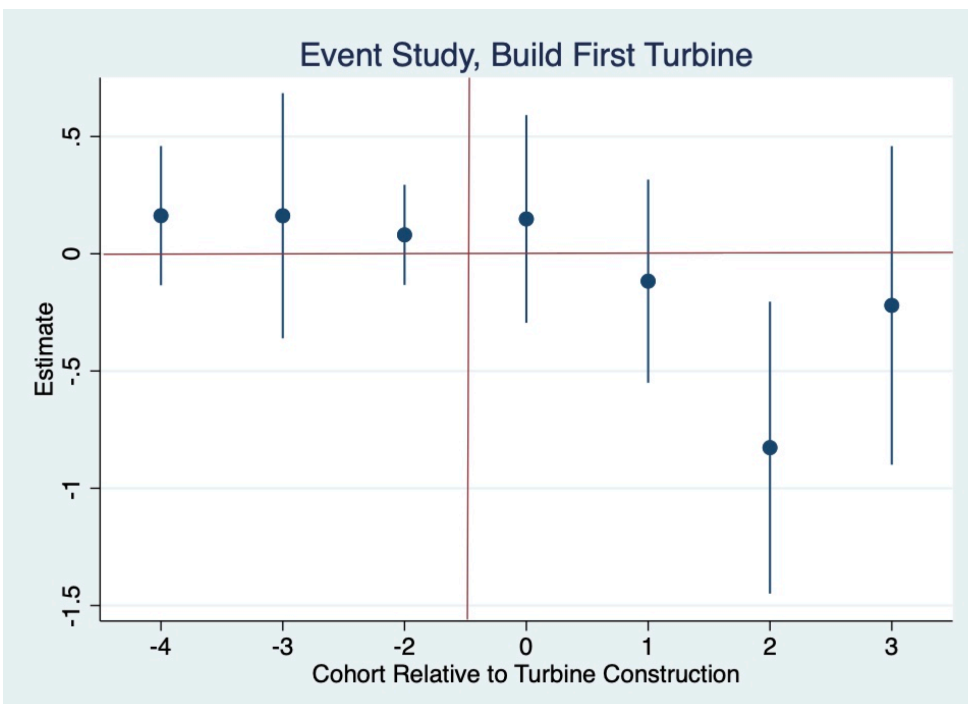

Notes: This graph plots the coefficients of a regression that compares differences in prices between homes 0 to 1 mile from a turbine and homes 1 one to 3 miles from a turbine. The results are split by cohort relative to the turbine being built. For instance, 0 represents the cohort that first experienced being close to a turbine. -4 includes all cohorts 18 to 24 months before the turbine was built; and 3 includes all cohorts 18 to 24 months after the turbine was built.

Map 1

McLean County

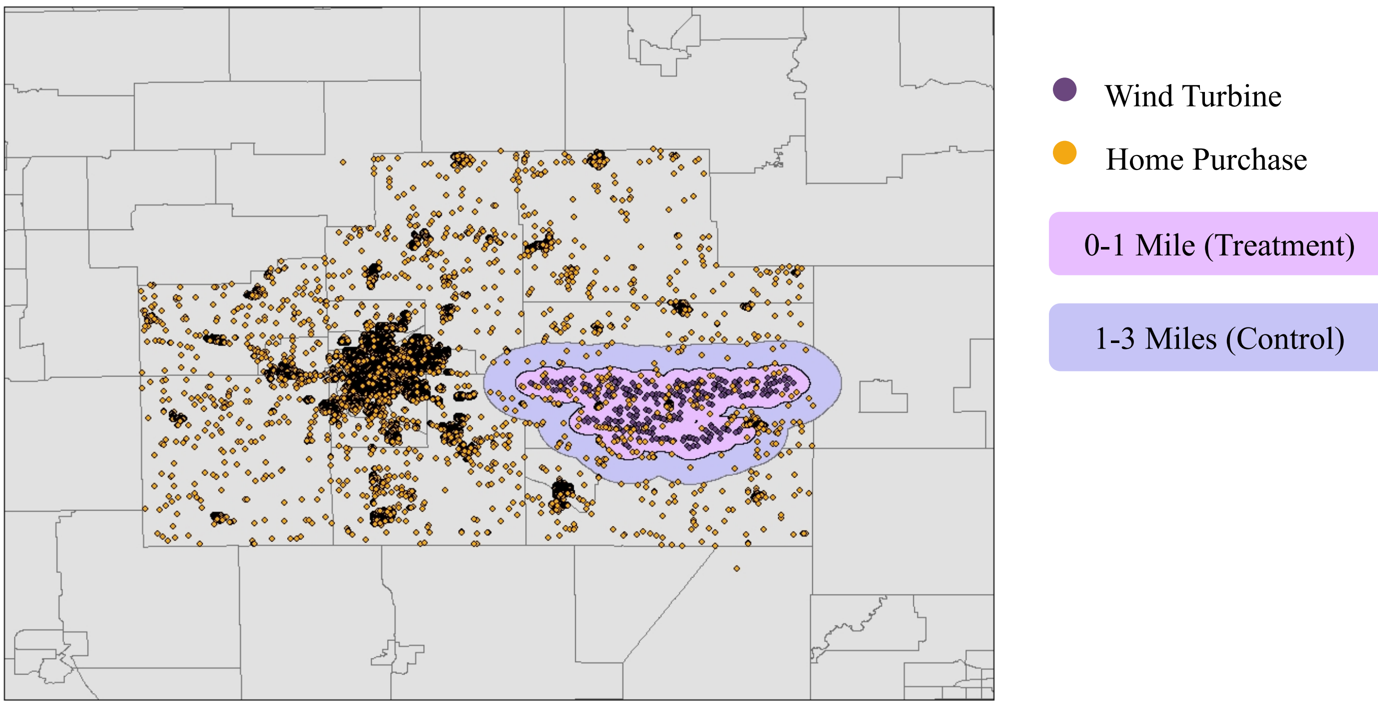

Notes: This map shows all the homes and turbines in my dataset. The transactions are colored in yellow, while the turbines are colored in purple. The lined area the data sit in is McLean County split into its census tracts. My treatment group, which contains only homes that are between 0 and 1 mile from a turbine, is highlighted in pink. My control group, which contains homes that are between 1 and 3 miles from a turbine, is highlighted in lavender.

Table 1

Summary Statistics

	Mean			Standard Deviation		
	McLean	0-1 Miles	1-3 Miles	McLean	0-1 Miles	1-3 Miles
Sales Price (Dollars)	157,216.2	132,894.9	132,376.3	87,767.3	88,886.3	93,568.2
Dist. to Turbine (Miles)	9.74	0.52	1.67	3.53	0.22	0.61
Lot Size (Acres)	1.85	20.74	12.13	10.42	20.74	26.91
Age (Years)	44.68	83.5	74.64	40.52	39.92	37.35
Building Area (Sq. Ft.)	1,800.27	1,187.87	1,234.47	3,660.90	359.37	491.04
Bathrooms	0.88	1.56	1.33	0.91	0.7	0.59
Num. of Observations	18,634	117	138			

Notes: This table summarizes the data used for this study and compares it to the county as a whole. The McLean columns gives values for all homes that sold in McLean County from 2005 to 2018. The 0-1 Miles columns give the same information, but only for homes in the treatment group. The 1-3 Miles columns give the corresponding values for the control group.

Table 2

Difference in Means

	Pre	Post	Difference
0-1 Miles	157,282.60	126,927.60	-30,355.00
1-3 Miles	104,175	137,156.20	32,981.2
Difference	53,107.60	-10,228.60	**-63,336.20**

Notes: This table gives the difference in means for the treatment and control groups in the pre- and post- periods. The table then takes the difference in those differences to give an estimate of the effect of building Twin Groves on home prices in the treatment group.

Table 3 - Homes Within 1 Mile of Twin Groves
Compared to Homes 2-3 Miles Out

VARIABLES	(1) 2 years before, after (10/19/06 cutoff)	(2) 3 years before, after (10/19/06 cutoff)	(3) 2 years before, after (9/20/05 cutoff)
First Turbine Built*One Mile	-0.198	-0.321	
	(0.236)	(0.215)	
First Turbine Built	-0.0441	-0.0239	
	(0.163)	(0.144)	
Project Approval*One Mile			-0.379*
			(0.222)
Project Approval			-0.0258
			(0.173)
One Mile from Twin Groves	0.289**	0.252	0.587***
	(0.140)	(0.151)	(0.203)
Lot Size (Acres)	0.111***	0.124***	0.135***
	(0.0244)	(0.0212)	(0.0246)
Building Area (Sq. Ft.)	0.000349	0.000375**	0.000649***
	(0.000226)	(0.000187)	(0.000226)
Total Calculated Bath Count	-0.102	-0.234	0.0110
	(0.186)	(0.176)	(0.183)
1-25 Years Old	0.441	0.458	-0.183
	(0.315)	(0.303)	(0.188)
50-100 Years Old	-0.0667	0.0324	-0.0318
	(0.141)	(0.123)	(0.143)
100+ Years Old	-0.0577	-0.0590	0.0338
	(0.142)	(0.139)	(0.149)
Constant	11.02***	11.17***	10.40***
	(0.402)	(0.393)	(0.441)
Observations	53	67	38
R-squared	0.443	0.471	0.623

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: This table considers the effect of being within 1 mile of a wind turbine on the natural log of the transaction price. First Turbine Built*One Mile is defined as a home that was within one mile of a turbine after October 19th, 2006. Project Approval*One Mile is defined as a home that was within one mile of a turbine after Sep. 20th, 2005. Column 1 considers transactions that occurred 2 years before or after the Oct. cutoff. Column 2 considers transactions that occurred 3 years before or after the Oct. cutoff. Column 3 considers transactions that occurred 2 years before or after the Sep. cutoff. In addition to these regressions, I ran one that considered a 2-year window around the October cutoff. That regression unfortunately had only 25 observations, so I dropped it.