## Online Appendix: Summaries of the 18 Core Studies

One of the earliest studies looking for a wind turbine effect on property values was done in the USA by Ben Hoen and a series of co-authors (Hoen et al. (2011)). The study uses 7459 sales transactions taken from within 10 miles of 24 facilities in nine states.[1] The study period is 1996 to 2007 and the coverage is broad, ranging from states in the North Pacific, Upper Midwest, Texas and Oklahoma, and the Northeast. These study areas were chosen purposefully to be both diverse and representative of the different kinds of wind facilities existing in the United States, and the study includes transactions up to several years before and after announcement and construction of the turbines. This timing dimension is a focal point of the study, and the authors have a good distribution of transactions across the different time periods. The study uses field observation for each transaction to judge the view of turbines for all properties in the sample, in addition to continuous distance and distance band measures. As is common in this earlier literature, the authors admit to a paucity of sales data close to the wind farms – 128 transactions within one mile and post-construction (PC) and fewer within one mile and in the post-announcement/pre-construction (PAPC) period. Keeping in mind that these data are spread over 24 different settings, in any individual community there are just a few transactions. A weakness of this analysis is that the authors pool the facilities and associated transactions into one dataset. The pooling helps improve the distribution of sales transactions across time periods and proximity to turbines, but might be masking heterogeneous impacts across facilities. Also, they do not employ a DiD analysis. They find no impacts from view and limited impacts from their distance measures in the PAPC period.

Hoen et al. (2015) follows a strategy similar to the 2011 study. They use a larger sample, a more rigorous methodological approach, but now only consider distance effects (measured in bands). The sample covers 51,276 transactions within 10 miles of wind farms in nine states (27 counties). They only include counties with at least 250 transactions within 10 miles of a wind turbine and the coverage of transactions near wind farms is better – 1198 sales within one mile with 376 of these post-construction and 331 within one-half mile with 104 post-construction. Again, these are spread over 27 different counties across the US, so at any single site the numbers are much smaller. They use traditional hedonic specifications as well as spatial econometric specifications in a DiD context using census tract fixed effects while also allowing county specific coefficients on structural variables but still assuming that wind turbine effects are homogeneous across the study areas. In their model controlling for spatial dependence, they find statistically insignificant reductions in housing prices within ½ mile of wind farms PAPC and PC. In their traditional hedonic specification and in all their robustness checks, they find no statistically significant negative impacts. Again, given limited data at any given site they cannot rule out such impacts at any particular facility, but they do have strong evidence that the "average" or typical finding in this setting is limited.

Heintzelman & Tuttle (2012) consists of three different, but related, case studies in rural Northern New York. Their transactions span 2000 to 2009 and include 6142, 3251, and 1938 transactions respectively. The data are configured such that a home in any of the three

[1] This represented 13% of the US wind power capacity in 2006.

communities (counties) uses the distance to the closest wind turbine as an argument in the hedonic price function. Since the analysis uses panel data and wind farms are constructed at different times, the closest wind farm for any observation varies over time. The most distant houses are 53 miles from the wind farms. The authors only consider distance in their analysis – a continuous measure (1/distance) and distance bands (half mile increments out to three miles) and the number of properties within three miles of a wind farm over the three cases is n=92, 118, and 251 – small relative to the full sample sizes. In a repeat sales analysis, sample sizes are scaled down by approximately 1/3. Since compensation was given to many of the households and/or communities in this study area, Heintzelman and Tuttle interpret the results in terms of whether compensation was "complete" or not. So, a negative impact from the hedonic model indicates under-compensation and a positive indicates over compensation.[2] They find significant negative impacts in two of the three study areas with both the continuous measure of distance and within some distance bands but, the bands do not give a strict monotonic decline in value with proximity to a wind farm which is of concern. The repeat sales analysis has a similar but somewhat weaker finding statistically. Again, like many earlier studies, the number of observations close to the wind turbines is quite small, which means that the accuracy of projecting impacts based on the distance variables is questionable. Important relative to the pooled studies in this review, they test a model where they pool the three, objectively quite similar, study areas into a single specification and reject the null hypothesis that the coefficients are identical across studies. This provides evidence that pooling transactions associated with multiple wind facilities in different areas may be problematic. In this particular case, the turbines in one of the study areas, where no property value impacts were identified, were set relatively far away from home clusters atop a plateau and, anecdotally, included a generous compensation package for the community, while in the other areas the turbines were more intermixed with the communities.

Vyn & Mccullough (2014) perform a single market difference-in-differences analysis similar to Heintzelman and Tuttle (2012) in a rural residential/agricultural landscape in Ontario, Canada. Their analysis considers a 133-turbine wind farm built in 2008 and the impacts on residential (n=5414) and agricultural (n=1590) properties using sales data from 2000-2010. They are the first study to include agricultural land in the analysis. Like others they express a concern for a limited number of transactions in close proximity to the turbines in the post-turbine periods. Their event dates are PA and PC. Unlike most, they consider two measures of PC (beginning of construction and completion of construction). With beginning of construction as the event date, they have 103 residential transactions within 5 km and 23 within 1 km. They use distance and view to measure impacts and consider an interactive model with distance and view. This allows the measure of view to have more intensity in closer proximity. Their view variable is a rating (following Hoen (2009)) that ranges from 1 to 3 and increases as the share of the turbine(s) that are visible rises, and is similarly based on field observation. A strength of this paper is the wide variety of specifications, including continuous distance measures, distance bands in kilometers (< 1, 1-3, and 3-5), measures of turbine density, a spatial autoregressive model, repeat sales, all in a difference-in-differences context. In none of their specifications do they find significant impacts of the wind turbines, residential or agricultural. They call attention to low population

[2] There was no formula or data for compensation by household, so it could not be explicitly incorporated into the model.

density in the area of their application and note that this a common thread with Heintzelman and Tuttle (2012), who find no impact in one of their three cases and it happens to be the case where population density is lower.

Heintzelman, Vyn, and Guth (2017) is a bit of a hybrid between the Vyn and McCullough (2014) paper and that of Heintzelman and Tuttle (2012) and focuses on a single industrial wind farm with 86 turbines that affects two distinct markets separated by the U.S./Canada border. The Wolfe Island wind facility is on a Canadian island in the St. Lawrence River that is visible from many properties on the American side of the river. Most of the study area is rural, with a significant population of seasonal/vacation home properties. The study includes 6,017 properties in the American market, and 2,262 in the Canadian market. Of these, less than 3% of the parcels in Canada, and less than 1% of those in the USA had views of the turbines, and there are only 47 parcel transactions within 5 miles PC in Canada and 58 in New York, with 39 and 15, respectively, with views PC. The small numbers problem is acute in this study. View is determined by field inspection of transactions within 5 miles of the turbines. The authors use a standard difference-in-differences approach, but have to deal with a significant confounding factor on the US side which is that those homes with the “best” views of the turbines are likely to be waterfront or very close to the water. They address this using a waterfront variable included in the model. The authors find significant PC effects for those homes with a view of the turbines on the US side in the full sample, restricting to homes within 20 miles, and again within 10 miles of the turbines. They find some evidence of negative effects from a continuous proximity (inverse distance) variable on the US side as well. Interestingly, they find no significant effects on the Canadian market. They speculate that this heterogeneity could be caused by a number of factors: compensation to landowners in Canada, inclusion in the siting process, a larger portion of vacation homes in the US market, and the fact that waterfront vacation homes on the Canadian side (largely on the island itself) may be more likely to face away from the turbines even if they do have a view of the turbines.

Skenteris et al. (2019) look at the impacts of wind turbines on two different Greek islands, both of which are popular tourist destinations. The two islands have 14 and 4 wind farms respectively, and with 400 and 1416 transactions over the period 2006 to 2016. All of their transactions are treated as if they are PC, so there is no causal identification (turbines are reported to have been built between 2001 and 2008). They use both continuous distance and distance bands to measure correlations, and find a significant negative effect of proximity to wind turbines on one of the two islands. The distance bands suggest that this effect is limited to be within 2km of the turbines. There is no measure of visibility. On South Evia, the island where they find impacts within 2km of approximately 14%, the authors suggest that turbines are, in general, in closer proximity to homes and less isolated from population centers, which likely explains the significant impact.

Sunak and Madlener (2016) focuses on four small to medium sized wind facilities (5 to 26 turbines each) in a small region of Germany. They study PC impacts, with 905 of 2141 total transactions occurring after construction, and consider view, proximity, and density of turbines for each property. These dimensions of impact are combined into an index variable (VIL) ranging from 1-6 where 1 indicates no view and 6 indicates an extreme view of, on average, 10 turbines at close distance. Of the 905 post-turbine transactions, about 42% have at least a medium view of turbines (VIL>3), while 32.8% have no view. This is a good amount of variation, and the inclusion of both visibility and distance in a single measure helps to overcome the lack of variation seen in so many studies that look at only one of these dimensions. Visibility is determined through the use of a high resolution Digital Surface Model (DSM). Parcels are all between 700m and 6km from the turbines. The prices used are land prices, separate from the home values in the transactions, although the mechanism for how these land values are calculated is unclear and there is some concern that they are not market-based. The empirical approach in Sunak and Madlener (2016) is somewhat different from most other studies in employing spatial econometric approaches including spatial lag and error models and a spatial Durbin model. These approaches are also compared to the standard spatial fixed effects models employed in most of the rest of the reviewed studies. They find generally negative and significant impacts of having at least a medium view of turbines PC. The spatial econometric approaches reveal some additional significant estimates relative to the non-spatial model at VIL=4 and VIL=6, although the qualitative results are quite similar. Unfortunately, they do not report results of decomposed proximity or view impacts for the purposes of comparison to other studies.

Sunak and Madlener (2017) use a subset of the data used in the 2016 paper, and apply a locally-weighted regression (LWR) technique which allows for coefficient estimates to vary over space, using a nearest-neighbor weights matrix in the estimation. They compare this to a standard fixed effects model. This study focuses on one farm with 9 turbines, 1405 property sales each between 982m and 5.2km from the turbines, all of which happen PC, so this is not a DiD approach, limiting causal identification. They do not use the VIL index here, but instead continuous distance and proximity dummy variables. They again use a DSM model to estimate visibility and include a visibility dummy variable and a count variable of visible turbines. They find significant negative impacts for properties less than 1 km away (visible and not) and a significant negative impact using the continuous distance variable in the standard models. They do, however, suffer the usual small numbers problems with fewer than 1% of their observations within 1km of a turbine. The LWR results, using continuous distance and number of visible turbines variables, are hard to summarize (results are only displayed on maps), but similarly suggest negative and significant impacts across the space and, unsurprisingly, these appear to be stronger in close proximity and when more turbines are visible.

Jensen, Panduro, & Lundhede (2014) represents a larger study, more in Hoen et al. (2011, 2015) style, looking at 24 study areas spread across Denmark, each containing property transactions within 2.5 km of turbines, some of which occur before turbine construction (it is not clear how many, and they do not attempt a DiD approach). In total, they include 12,640 transactions over 12 years from 2000-2011. They also employ spatial econometric approaches in the form of spatial error and spatial autoregressive models. Like other studies, they test proximity and visibility (using a DSM approach, as in Sunak and Madlener (2016, 2017)), but they also explicitly investigate noise effects using an equation linking turbine characteristics and distance to estimate noise effects under "ideal" conditions for noise transmission – the worst case noise scenario. They do not report numbers of observations at various distances from turbines, but do report that 33% of their observations have a view of the turbines. In addition, 68% of their observations are modeled to experience turbine noise exceeding 20dB. This is impressive coverage of affected parcels. They find that, simultaneously, both noise and visibility have significant negative impacts on property values and that view effects are magnified by closer proximity. Their qualitative results are unchanged by the spatial approaches as compared to simple OLS, although the magnitudes of the visibility effects are substantially lower in the spatial models. It is worth noting, however, that they do not report non-spatial fixed effects models. The main contribution of this study is the explicit measurement of noise impacts which most other studies simply imply through proximity measures.

Jensen et al. (2018) expands on the 2014 study with a larger sample of transactions, and the inclusion of both on- and off-shore wind turbines with a large range of characteristics. They are one of only three papers we are aware of that study offshore turbines. Although they only partially report on their turbine data, they include turbines ranging in height from 22-140 meters and in capacity from 11-3000kW. They conduct separate analyses for the on- and off-shore turbines, and separate analyses for primary and secondary (vacation) homes in each of five separate markets. For the onshore turbines, they only have PC data, so they cannot make any causal inference about the impact on property values. The econometric approach is a spatial semi-parametric Generalized Additive Model (GAM) to help control for spatial dependence and the usual endogeneity and omitted variables concerns. Impact of the onshore turbines is measured in the number of turbines within 3km and a measure of turbine density that increases with proximity in that same range, forgoing the more detailed measures of impact used in the 2014 study. Including both impact variables in their models (which seems to present a significant unnecessary collinearity risk) they find negative correlation with property values for both number and density of turbines in all five regions, and almost all of these coefficients are statistically significant for primary homes. For secondary homes they only include the density variable and find significant negative correlations in three of five regions. Overall, they have more than 85,000 observations for the onshore analysis, but these are split into 10 separate analyses with numbers ranging from 408 to 25,301 observations per analysis. The analysis of offshore turbines relies on a much smaller sample, still split between primary and secondary homes (and run separately for the two farms), but has the advantage of including both pre- and post-construction transactions, allowing for a DiD analysis. These analyses focus on view impacts PC, and include 275/91/43/9 homes with turbine views out of 1611/703/2712/1316 total transactions respectively. Of these, still smaller numbers of homes with views were traded after construction, so there is a small numbers problem in this study. It isn't clear how view was determined. The turbines studied in this analysis are between 3.5 and 9.5 km from shore and are comprised of a total of 162 turbines with hub heights of 80m. They find no significant PC property value impacts from the offshore turbines.

Lang, Opaluch, & Sfinarolakis (2014) are the first to study impacts in a more densely populated area, considering 10 separate wind power facilities in Rhode Island. Another difference between their study and those that came before it is that, given the more urban/suburban setting, nine of the facilities they study are single turbines and one is a cluster of three, so not traditional industrial scale wind facilities.[3] The facilities include turbines on school grounds, residential

[3] A related study using the same dataset by Gorelick (2014) generates very similar results.

communities, and mixed residential/commercial communities. The authors note that some are "…coupled with an existing disamenity such as proximity to a highway or water treatment plant." They consider 48,545 transactions within 5 miles of a turbine and they have 584 transactions within ½ mile and 3254 within 1 mile. The PAPC and PC counts in the ½ mile range are 75 and 74, which are low relative to the full sample but comparable to, or an improvement on some other studies. Impact is measured using distance bands and view in separate hedonic models. They field-verify viewsheds for transactions within two miles of the turbines and use five categories to define view. They consider announcement and construction as events and use a standard DiD approach, as well as a repeat-sales analysis. They find negative impacts in the distance bands within three miles but all are statistically insignificant, small, and nonmonotonic (failed to increase in impact moving closer to a turbine). This provides at best weak evidence of impact. In the repeat-sales analysis they find a significant PAPC impact period in a ½ - 1 mile distance band (5.9% price impact) but not in the other bands including the closer 0 – ½ distance band. The authors obtain similarly insignificant results in considering the view of the turbines. This analysis, along with Hoen & Atkinson-Palombo (2016) which we discuss next, provides evidence that facilities with small numbers of turbines are less likely to have adverse property value impacts.

Hoen & Atkinson-Palombo (2016) similarly study the impacts of small-scale wind farms (mostly 3-turbine farms or smaller) on residential properties in densely populated areas spread across Massachusetts from 1998 to 2012. The analysis of more than 122,000 transactions is pooled over 21 markets and includes properties in coastal, mountainous, and highly developed areas in the state. There are a total of 41 industrial turbines included in the study. Compared to many of the previous studies, they have a relatively large number of affected homes, with 1107 transactions within ½ mile of a turbine. Of these 230 are PC and 224 are PAPC. They consider only proximity impacts, focused on the effects of being within ½ mile of a turbine, PAPC and PC, and their base control group extends to five miles from any turbines, including robustness checks on both extents (within ¼ mile and extending the control out to 10 miles). Controlling for housing attributes and many locational amenities and disamenities, they first find strong evidence that, as one might expect, turbines are more likely to be sited in areas with lower *ex ante* property values. This endogeneity problem is a warning to the literature at large that careful consideration must be taken to control for these effects. Using a DiD analysis to control for this endogeneity, then, they find some limited evidence of impact in the PAPC period, but no evidence of impacts in the PC period. Interestingly, they do find evidence of impacts from other spatial variables including proximity to highways, beaches, power lines, and landfills, which suggests that the method is valid and strengthens their case that the turbines in their study area are not having an impact. It is important to note, again, however, that these turbines are not in utility scale farms, but in much smaller clusters.

Vyn (2018) is in the same spirit of the larger studies covering a large land area in a single study, in his case the Canadian province of Ontario. He uses 22,159 observations within 20km of a wind turbine as his sample, covering 37 wind farms ranging from 3 to 110 turbines in rural areas. By comparison to most other studies, Vyn (2018) has a good distribution of transactions across PA/PAPC/PC periods and about 20% of the transactions are within 5km of the turbines. This distribution allows for a robust estimate of property value impacts using a standard DiD design. Impact is measured using distance bands between the transaction and the nearest turbine, as well as a measure of turbine density. The novelty of this study is the use of an additional policy variable to disaggregate the sample – whether or not a community has labeled itself as an "unwilling host" community to wind turbines. This appears to be a bit of a movement across Ontario where some communities have chosen to signal their unhappiness with wind development through local non-binding resolutions. Importantly, the power to site wind facilities is held not by the local community, but by the province. Overall, Vyn (2018) finds negative and significant impacts from proximity across the province in both PAPC and PC eras, but when these impacts are disaggregated into unwilling host communities and others, the negative impacts are broader and more significant in unwilling host communities. Measured impacts of turbine density are mostly only negative and significant in unwilling communities. There is no information about the timing of the unwilling host designation relative to announcement and construction of the turbines, so it is hard to know the direction of the correlation between negative impacts and self-designation as an unwilling host.

The last four studies we consider in this review have in common their very large sample sizes and geographic scales. Gibbons (2015) uses a large sample (1.7M transactions) from 2000-2011 for England and Wales to examine the impacts of nearly all wind developments in the study region – some 148 wind farms ranging in size from 1 to 106 turbines. Included transactions are those within 14km of wind facilities. The transactions are averaged at the post code level to generate quarterly post code average prices over the sample period. He then uses a digital elevation model to approximate visibility using post code and wind farm centroids. Visibility is then interacted with a PC dummy variable and a series of proximity dummy variables to indicate the effect of visibility on proximity at various distances. Two control groups are created to compare to treated homes in the DiD analyses employing post code and temporal fixed effects. In the first analysis, treated observations are compared to observations in post codes that will be treated by the end of the sample period. This analysis essentially isolates the effect of construction on a set of observations that will eventually have visible wind turbines. The second analysis compares the same set of treated observations to observations that are in the same proximity bands to operational turbines but are estimated to not have them in their viewshed. This analysis isolates the effect of visibility (the spatial dimension) rather that construction (the temporal dimension). Because of the scale of their dataset, this study does not suffer the small numbers problem common in much of the literature. By using postcode averages as observations instead of individual transactions, and calculating visibility approximately at the same scale, there are many more observations in close proximity to turbines and within the viewsheds. As a result, they identify 1,142 postcodes within 1km of a wind facility and 1,125 of these are estimated to be able to view the turbines. It is important to recognize that the aggregation, even at as small a scale as post code, and the approximation of viewsheds both introduce error into their analysis although it isn't clear in which direction this might bias the estimates. They find consistent significantly negative impacts of visible and operational wind turbines on property values in both of their analyses, and these are robust to various specifications of control variables. They also include an analysis accounting for the size of the wind facilities and find that larger facilities have larger negative impacts.

Droes and Koster (2016) use a similarly large dataset of some 2.2M transactions in the Netherlands spanning 1985-2011 to study the impact of the 1,898 onshore turbines that have been constructed in that period. They use a standard DiD analysis to explore PC impacts for homes within 2km of turbines versus all other homes. They do not attempt any analysis of visibility, nor do they attempt to find impacts beyond the 2km boundary. In addition to the large dataset, what sets this study apart is the large number of robustness checks varying the fixed effects approaches, the extent of the sample, and they use interaction terms to allow the effect to vary over some different temporal subsamples, by home setting, and by the number of turbines within the 2km treatment zone. Overall, this is perhaps the most comprehensive study of wind turbine impacts on property values, and they find very consistent negative and significant impacts across nearly all specifications of the analysis.

Eichholtz, Kok, Langen & van Vulpen (2018), use virtually the same dataset as Droes and Koster (2016) to investigate the effects of all power plants, not just, but including, wind turbines. The strength of the paper is the ability to take account of plant openings and closings, and to allow for the direct comparison of impacts from wind facilities to other facilities including coal, gas, and biomass. They use a traditional DiD approach focused on proximity, but define their treatment and control groups carefully. In particular, they count all transactions within 2.5 km of a power plant treated, then allow for a 1.5 km buffer distance of excluded transactions before their control group starts at 4 km, and extends out to 20 km. This helps to prevent leakage of impacts across the 2.5 km boundary from biasing estimated impacts. There is no estimate of visibility or other impacts. In addition to the impacts of these various facilities, they also allow the effects to vary based on the size of the plant (# of turbines in the case of wind) and on whether or not the plant is in an urban setting. They use location fixed effects and standard controls of home characteristics. They have more than 300,000 observations in their treatment group, with more than 1.4 million in the full control group. In general they find that gas and wind turbine openings have a negative impact on property values in both urban and non-urban settings, while biomass plants have positive impacts in non-urban settings but negative impacts in urban settings. Wind turbine effects are generally about half or less the magnitude of gas plant effects. Closing effects are less significant (there are fewer closings than openings), in general, but basically have the opposite impacts of plant openings. A repeat sales analysis also provides similar results.

Similar to Eichholtz et al. (2018), Jarvis (2021) investigates not just wind turbines, but also solar installations. In addition to measuring the presence of these renewable energy facilities, they measure the effect of size. Their primary treatment variable is a difference-in-differences measure of size of nearby facilities where proximity is measured with distance bins and size in MW. They use a log-log specification so that they are measuring the change in property values for a 1% change in installed capacity within a specified distance.