

2023

House of the rising sun: The effect of utility-scale solar arrays on housing prices

Vasundhara Gaur
University of Rhode Island

Corey Lang
University of Rhode Island, clang@uri.edu

Follow this and additional works at: https://digitalcommons.uri.edu/enre_facpubs

Citation/Publisher Attribution

Gaur, V., & Lang, C. (2023). House of the rising Sun: The effect of utility-scale solar arrays on housing prices. *Energy Economics*, 106699. <https://doi.org/10.1016/j.eneco.2023.106699>
Available at: <https://doi.org/10.1016/j.eneco.2023.106699>

This Article is brought to you by the University of Rhode Island. It has been accepted for inclusion in Environmental and Natural Resource Economics Faculty Publications by an authorized administrator of DigitalCommons@URI. For more information, please contact digitalcommons-group@uri.edu. For permission to reuse copyrighted content, contact the author directly.

House of the rising sun: The effect of utility-scale solar arrays on housing prices

The University of Rhode Island Faculty have made this article openly available.
Please let us know how Open Access to this research benefits you.

This is a pre-publication author manuscript of the final, published article.

Terms of Use

This article is made available under the terms and conditions applicable towards Open Access Policy Articles, as set forth in our [Terms of Use](#).

HOUSE OF THE RISING SUN:
THE EFFECT OF UTILITY-SCALE SOLAR ARRAYS ON HOUSING PRICES

Vasundhara Gaur and Corey Lang*

ABSTRACT

While utility-scale solar energy is important for reducing dependence on fossil fuels, solar arrays use significant amounts of land (about 5 acres per MW of capacity) and may create local land use disamenities. This paper seeks to quantify the externalities from nearby solar arrays using the hedonic method. We study the states of Massachusetts and Rhode Island, which have high population densities and ambitious renewable energy goals. Using difference-in-differences, repeat sales identification strategies, results suggest that houses within 0.6 miles depreciate 1.5-3.6% following construction of a solar array. However, additional analysis reveals that this average effect is primarily driven by solar developments on farm and forest lands and in rural areas, which is intuitive given the composite impact of solar, loss of open space, and loss of rural character. For these states, the local disamenities are the same order of magnitude as the global benefits of abated carbon emissions, which helps explain local opposition to siting.

Keywords: Solar energy; Utility-scale solar; Hedonic valuation; Difference-in-differences
JEL codes: Q24; Q42; Q51

* Contact author: 214 Coastal Institute, 1 Greenhouse Rd., University of Rhode Island, Kingston, RI, 02881. Email: clang@uri.edu.

We thank three anonymous referees, Ben Hoen, Salma Elmallah, and conference participants at AERE and NAREA for useful feedback. This work was supported by the USDA National Institute of Food and Agriculture, Agricultural and Food Research Initiative Competitive Program, Critical Agricultural Research and Extension, grant number 2019-68008-29826. The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

1 INTRODUCTION

Solar energy in the United States has grown at an average rate of 49% per year since 2009, making the US the second largest producer of solar energy in the world (EIA International Energy Outlook 2019). In 2019, solar energy accounted for 40% of all new capacity additions in the country, the largest ever in its history, and exceeding all other energy sources (Perea et al., 2020). By June 2020, the cumulative installed capacity of solar in the United States reached 81.4 gigawatts (GW), which is enough to power 15.7 million homes (Perea et al., 2020). Solar is predicted to overtake wind to become the largest source of renewable energy in the US by 2050, accounting for 46% of all energy produced from renewable sources (EIA Annual Energy Outlook 2018).

While there is a broad support for renewable energy in the United States (Bates and Firestone, 2015; Farhar, 1994; Firestone et al., 2018; Hoen et al., 2019; Krohn and Damborg, 1999), and for solar energy in particular (Carlisle et al., 2014, 2015; Farhar, 1994; Greenberg, 2009; Jacobe, 2013; Pew Research Center, 2019), the development of large-scale solar installations has not been obstacle free. One major hurdle to overcome before construction begins is the siting process. Solar installations require between two to ten times more land area than non-renewable sources to generate the same amount of energy, and the requirement of large tracts of land for their construction has become the largest cause of land use change in the United States (Trainor et al. 2016; Ong et al. 2013).¹ Recently, the siting of large solar projects has become contentious in some parts of the country due to concerns about visual disamenities, impacts on ecosystems, building new transmission lines, loss of a town's rural character, water pollution, fire risk, water use, and reduction in property values (Farhar et al., 2010; Gross, 2020; Lovich and Ennen, 2011). The debate is especially heated when solar development is proposed on existing farm and forest lands, which is common because these are the cheapest locations for development, but many consider antithetical to environmental objectives (Kuffner, 2018; Naylor, 2019).

The purpose of this paper is to quantify the externalities associated with proximity to utility-scale solar installations using hedonic valuation. Our objective is to provide policy

¹ Trainor et al. (2016) include the land area required for surface extraction of fossil fuels in their calculation of the impact of coal. However, they do not make the same consideration for solar panels, which also include surface-mined metals like silicon, aluminum, and silver.

relevant non-market cost estimates in order to help state and municipal policy makers implement policies and decisions that reflect public preferences.

We focus on the states of Massachusetts (MA) and Rhode Island (RI), which are ideal for two reasons. First, both states have recently experienced a sudden boom in the development of large-scale solar installations. This trend has been driven by the Renewable Portfolio Standards (RPS), regulations that require increased energy production from renewable energy sources, which have been adopted by both states. MA's RPS calls for 25% of electricity generated by renewable sources by 2030 and RI's RPS calls for 38.5% by 2035. Second, both states have high population density, ranked 2nd and 3rd among U.S. states (US Census Bureau, 2020). This level of development means that most solar sites are proximate to residential areas, which yields many observed transactions for precise estimates. Further, because so many households are impacted in these areas, our estimates are highly policy-relevant.

We analyze the impact of utility-scale solar installations sized 1 MW and above on nearby property prices in MA and RI.² We apply two empirical approaches. First, we use a traditional repeat sales, difference-in-differences (DID) identification strategy, which compares changes in housing prices after construction for nearby properties with those further away. We empirically estimate the spatial extent of treatment to be 0.6 miles from the solar installation and choose a cutoff for control properties of two miles. Our primary sample consists of 282 solar installations, 11,292 housing transactions occurring within 0.6 miles (treated group), and 95,999 transactions between 0.6 and two miles (never-treated control group). However, pre-treatment trends are not perfectly parallel which raise concerns about necessary assumptions holding. Given these concerns, we also estimate a DID model using only ever-treated properties, which relies entirely on temporal variation in construction dates. This method is preferred if there are endogeneity concerns about the siting of solar being correlated with trends in prices and not just levels. We present both models for all specifications and hedge about which is preferred.

Across a variety of specifications, our results suggest that solar installations negatively affect nearby property values. Results that average effects across all sites find negative impacts ranging from -1.5% to -3.6% (equal to reductions in housing values of between \$4,721 and \$11,330, prices adjusted to 2019 levels), with the models using only the ever-treated sample

² Following the U.S. Energy Information Administration (EIA), we define large-scale solar installations as those with an installed capacity of 1 MW or larger.

consistently indicating larger effects. However, we examine heterogeneity in treatment effects that lead to important insights. We posit that solar arrays on farm and forest lands (“greenfields”) cause greater externalities, given the combination of solar-specific disamenities and loss of open space amenities. Further, rural areas may be more impacted if industrial solar arrays are incongruent with highly valued rural character, but on the other hand space is scarcer in non-rural areas. We find that the average treatment effects are to a large part driven by greenfield arrays and arrays in rural areas. Specifically, greenfield developments lead to a decline in housing prices between 2.0% and 4.4%. For properties lying in the vicinity of solar installations in rural locations, the decrease in value is between 2.5% and 5.8% post solar installation construction. Coefficients on non-greenfield sites and non-rural sites are consistently negative, but never statistically significantly different than zero.

Our findings suggest that utility-scale solar arrays create local, negative externalities. This helps explain local concerns and opposition to new development and gives pause to current practices of not including proximate residents in siting decisions or compensating them after siting has occurred. While a full benefit-cost analysis is well beyond the scope of this paper, we can compare the local, negative externalities to the value of greenhouse gas reductions from the solar arrays, which is the major global benefit. Our back-of-the-envelope calculations imply a benefit-cost ratio ranging from 1.65 to 0.69 depending on the choice of model. While it is promising that the benefit-cost ratio can be greater than one, it is clear regardless that a substantial and uneven burden is imposed on local areas to achieve global benefits of a similar magnitude. However, benefit-cost ratios are likely to be more favorable in other states due to different sources of fossil fuels and sparser population.

The recent growth in utility-scale solar has been met with a wave of research focused on assessing externalities and siting preferences. Prior hedonic valuation research includes Abashidze (2019) who applies a DID methodology with treatment and control defined by proximity, similar to our first model. Using data from North Carolina, USA, she finds that property values decline 8.7% post-construction within 1 street-network mile of a solar array. She similarly tests for treatment effect heterogeneity by prior land use, but finds no statistical differences. Dröes and Koster (2021) apply both traditional DID and an alternative that relies on an ever-treated only sample, similar to our second model. Working in the context of the Netherlands, they find a 1.5-4.7% decline in value for properties within 1 km of a solar array.

Jarvis (2021) uses data from the United Kingdom and also applies a DID methodology, but uses properties near solar sites that were proposed but not built as the control group. He finds zero statistical impact on property values. Maddison et al. (2022) use a traditional DID methodology and find that property values in England and Wales decline by 5.6% within 750 meters of solar installations. However, this effect is driven exclusively by large (5 MW capacity or greater) installations lying south of proximate houses. They find no impacts for homes near smaller installations or for any other orientation.³ Abashidze and Taylor (2022) instead analyze the impact of nearby solar on agricultural land value. In general, they find no net impact, but farms close to electric transmission lines increase in value after a solar installation is constructed nearby, suggesting capitalization of an increased likelihood of lease payments.

In addition to the hedonic valuation studies, there are several stated preference studies that also examine externalities from utility-scale solar siting. Botelho et al. (2017) survey residents of Portugal using a contingent valuation approach and find that respondents in the proximity of large solar installations are willing to accept \$12.93 – \$56.64 per month on average as compensation, which is on par with our capitalization estimates converted to monthly payments. In addition, Botelho et al. conduct a nation-wide discrete choice experiment to delve into aspects of siting that drive the disamenity and estimate that respondents are willing to pay \$8.65, \$7.57, and \$5.15 per month to avoid negative impacts on flora and fauna, landscape, and glare effects, respectively. Kim et al. (2020) carry out a choice experiment in South Korea focused on land use and find large WTP (\$1,000-\$2,000 per household per month) for solar to be sited on rooftops and walls instead of farmlands, orchards, and mountainous areas. Gaur et al.

³ Our work is additionally closely related to the extensive hedonic applications assessing externalities of wind energy. Within the United States, studies that use data with large numbers of observations close to turbines find no significant impact on property prices, including Hoen and Atkinson-Palombo (2016) and Lang et al. (2014) using Massachusetts and Rhode Island data, respectively, and Hoen et al. (2015) examining wind farms across the country. In contrast, studies in European countries find that wind turbines have a significant negative impact on nearby properties, though the magnitude of the effect differs by region (Dröes and Koster, 2016, 2021; Gibbons, 2015; Jarvis, 2021; Sunak and Madlener, 2016). Vyn (2018) finds the Canadian experience to be heterogeneous and dependent on community acceptance. More recently, hedonic methods have focused on estimating externalities from offshore wind turbines. While this literature is still in its infancy, early studies indicate no negative impacts to property values or rental rates in the vicinity of offshore wind turbines (Jensen et al., 2018; Carr-Harris and Lang, 2019; Dong and Lang, 2022). Hedonic valuation has also been applied to residential rooftop solar. General consensus is that houses installed with rooftop photovoltaic panels sell for a premium, though there is regional variation in the size of the effect: 3.5% in California (Dastrup et al., 2012; Hoen et al., 2012), 5.4% in Hawaii (Wee, 2016), 17% in Arizona (Qiu et al. 2017), and 3.2% in Western Australia (Ma et al. 2016). However, this literature is only tangentially related as it is about quantifying internalities (e.g., valuation of personal financial benefits, warm glow), not externalities, and has nothing to do with land use.

(2022) develop a choice experiment focused also on land use, but also attributes of arrays such as visibility and property line setback. They survey residents of Rhode Island and find the largest determinant of approval is prior land use with positive WTP for arrays on non-greenfield sites (\$10 to \$21 per month per household) and negative WTP for arrays on greenfield sites (-\$13 to -\$49 per month per household). In addition, they find that households are willing to pay \$6-8 per month to avoid full visibility. We contribute to this emergent literature by offering another data point in the understanding of externalities and providing a test of convergent validity for stated preference work.

Our work also relates to sociology and psychology research on renewable energy acceptance. Patrick Devine-Wright, a leading scholar in this area, argues in several articles against a simplistic NIMBY explanation, instead “local opposition is conceived as a form of place-protective action, which arises when new developments disrupt pre-existing emotional attachments and threaten place-related identity processes” (Devine-Wright, 2009). Our examination of heterogeneous treatment effects by prior land use and population density are consistent with the ideas of Devine-Wright and others (e.g., Brittan, 2001; Firestone et al., 2018; Wester-Herber, 2004) in that we are finding larger negatives associated with solar arrays developed in areas where this technology contradicts ‘local character’ and substantially alters the ‘positive distinctiveness’ that people associate with such places.. We contribute to this literature by adding a revealed preference, market-based metric of place meaning. The combination of frameworks seems particularly important when thinking about optimal siting of new infrastructure or prioritizing over multiple objectives in the energy landscape.

2 DATA

To implement the hedonic analysis, we build a composite dataset that integrates: 1) the data on the location and attributes of all solar developments in MA and RI, and 2) the data on attributes and locations of residential properties in MA and RI.

2.1 Solar data

The dataset on solar installations is obtained from the Energy Information Administration’s (EIA’s) report EIA-860M, or the Monthly Update to the Annual Electric Generator Report. The EIA-860M contains data on the total capacity of electric generation

facilities in the United States that have a capacity of 1 MW and above, their point location (latitude and longitude), and the month and year that operation begins. Figure 1 represents a map of 284 solar installations constructed prior to August 2019, which is when we set the cutoff for being in our sample. The installations are well dispersed across all regions in both states, which increases confidence that estimates will not be affected by unobserved regional shocks. Figure 2 graphs new and cumulative solar capacity by year. The first installation began operation in June 2010. New capacity displays a continuous upward trend through 2014. There is a sharp fall in 2015, after which the trend rises again and peaks in 2017, before falling again in 2018. As of August 2019, the cumulative solar capacity of utility-scale arrays in MA and RI is 817 MW. Capacity factors for this region are about 16.5% (EIA 2019), which means these solar installations are collectively producing about 1180 GWh of electricity per year, which is enough to power 157,681 homes.

One limitation of our data is that we do not have shapefiles representing the exact footprint of the solar installations, thus we must approximate that using Geographic Information Systems (GIS) software. Solar installations require approximately 5 acres of land per MW of capacity (Denholm and Margolis, 2008; Ong et al., 2013). We assume that the point location is the centroid of the installation and then create a circle around it with an area equal to 5 times the capacity (in MW) of each array. We manually crosscheck the EIA data with Google Maps, and correct the latitude and longitude when they do not correspond to the centroid of the array. We recognize that this approximation of distance could lead some properties to be misclassified as treatment or control, inducing a small amount of classical measurement error in treatment status. As a result, our DID estimates may be slightly attenuated.

We hypothesize that prior land use may affect property value impacts. Specifically, houses in proximity to farms and forests that are developed into solar may depreciate more than houses in proximity to a brownfield or capped landfill that is developed into solar.⁴ Since farms, forests, and other open space are amenities and boost home values (Irwin, 2002; Lang, 2018), conversion of these types of lands may lead to larger price decreases because it is the combination of a loss of amenities and the gain of disamenities. To infer prior land use, we overlay the estimated circular footprints on 2005 land use data obtained from Massachusetts

⁴ Solar developers prefer farm and forest lands because they have substantially lower construction costs compared to alternative sites like brownfields, covered landfills, parking lot canopies, and industrial areas.

Bureau of Geographic Information and 2011 land use data obtained from Rhode Island Geographic Information System for the respective states. We then assign each installation a prior land use: ‘greenfield’ if it was formerly either a farm or forest land, and ‘non-greenfield’ if it was either a commercial site or a landfill.⁵ 63% of installations and 70% of capacity is classified as greenfield (see Figure A1 in the online appendix).

2.2 *Property data*

We use ZTRAX housing transaction data from Zillow (<http://www.zillow.com/data>), which include information on property location (latitude and longitude), sales price, date of transaction, and many property characteristics (lot size, square feet of living area, number of bedrooms, number of bathrooms, year built, number of fireplaces, central air-conditioning, and swimming pool). We begin with the universe of property transactions from January 2005 to June 2019 in the states of RI and MA, and then make several cuts, which we detail here. We drop observations with missing geo-coordinates (latitudes and longitudes), missing values for sales price, and those with prices of \$100 or less, which are clearly not arms-length transactions. Our sample at this point consists of 1,416,335 unique transactions. Sales prices are adjusted to 2019 levels using the Northeast regional housing Consumer Price Index from Bureau of Labor Statistics. We drop transactions in the top and bottom 5% of the sales price distribution to get rid of outliers, which reduces our sample size by 10%. Condominiums, which account for 26.70% of our sample size (after the previous cut), are dropped. We exclude houses with missing observations for bedrooms, full bathrooms, and half bathrooms. We also drop groups of properties with the same latitudes and longitudes, but different addresses because this indicates incorrect geocoordinates (Nolte et al., 2021). Further, we drop observations that have more than four stories, six bedrooms, five full bathrooms, or three half bathrooms. Houses that underwent major reconstruction are dropped since they may have different attributes in previous transactions. We exclude homes that sell before they were built, as there is evidence these are lot sales without improved property. Properties that transact more than once on the same date are likely to be subdivisions and are therefore excluded. We drop 11 observations that lie within the calculated solar array buffer. We also drop single-family residential properties with lot sizes

⁵ Details about the classification procedure are provided in the online appendix.

larger than two acres, since large plots could be potential sites for solar development and price impacts of nearby solar could be completely different. These steps together reduce our sample size by 35.70%, bringing the number of remaining observations down to 600,763. We also exclude all properties that transact only once in the chosen time frame because we focus on repeat sales analyses, which is considered the gold standard of all hedonic valuation (Banzhaf, 2021; Bishop et al., 2019; Nolte et al., 2021). We lose 50.32% of our observations after this cut, giving us a sample size of 298,453. We spatially merge the solar data with the property dataset by matching every property to the nearest eventual site of solar development to infer proximity. Finally, we exclude transactions occurring more than three miles of any eventual solar installation since properties further away are likely to differ significantly in several observable and unobservable characteristics. This removes 38.48% of observations, leaving us with a sample size of 183,566. Based on analysis presented in Section 3.1 in which we estimate the spatial extent of treatment, we further drop transactions greater than two miles away from an eventual solar site, and our final sample is 107,291 repeat-sales transactions representing 45,795 different properties.

Similar to prior land use, we hypothesize that existing development in areas surrounding solar arrays may impact property prices. Many rural areas pride themselves on their rural character and residents seek out that type of bucolic setting. Hence, construction of solar installations could be seen as an industrialization of the landscape and may cause larger negative impacts on property values. Whereas solar arrays in suburban and urban areas may be viewed as more congruent with existing surroundings. However, space is also more constrained in suburban and urban, which could lead to greater impacts there. We proxy for rural character with municipality-level population density, which comes from the 2010 Census. We define an indicator variable *Rural*, which equals one if the town has a population density of 850 people per square mile or fewer. We chose this cutoff because 850 is the average population density of MA, which forms the bulk of the observations in our dataset, and, at this cutoff, a little over a third of the properties and 66% of the solar installations are classified as rural, which we believe are reasonable proportions. It is important to note non-rural properties should not be thought of as urban, but more suburban. Very few utility-scale solar developments are built in urban areas as there is just not space.

3 METHODS

We use the difference-in-differences (DID) method in the hedonic framework to analyze the impact of solar installations on housing prices. We begin with a standard hedonic DID setup, in which we define treatment and control based on proximity. Properties located near large-scale solar installations are compared to similar properties that are further away from such installations, before and after construction.

We estimate repeat sales models that use within-property variation to identify the treatment effect by including property fixed effects:

$$P_{it} = \beta_1 Post_{it} + \beta_2 (Treated_i \times Post_{it}) + \mathbf{X}_{it}\boldsymbol{\gamma} + \alpha_i + \epsilon_{it} \quad (1)$$

Where P_{it} is the log sale price of house i at time t . $Post_{it}$ is an indicator for post-treatment, which equals 1 if a house sells after the treatment date, and $Treated_i$ is a dummy variable equal to 1 if a house is located near an eventual solar site and 0 otherwise. \mathbf{X}_{it} is a vector of temporal controls. α_i controls for time-invariant unobservables at the property level (e.g., school quality, proximity to other amenities and disamenities, traffic volume, walkability, property layout, curb appeal, etc.). Lastly, ϵ_{it} is the error term. In our basic specification, \mathbf{X}_{it} includes month-year fixed effects, which capture macroeconomic trends that affect the entire region that could be correlated with solar development trends. In addition to this specification, we estimate two more models. The first adds municipality-specific time trends to account for different housing price trends between municipalities. The second includes county-year fixed effects which allows for county-specific, nonparametric differences in housing market trends. In all models, we cluster standard errors at the census tract level to allow for correlated errors within a larger area. β_1 is the change in prices for control properties from before to after treatment. β_2 , the coefficient of interest, is the differential price change from before to after solar development for treated properties relative to control properties.

There are two aspects of this DID setup that are initially uncertain: the spatial extent of treatment, and the date on which treatment occurs. We define the treatment distance to be equal to 0.6 miles and provide evidence to support this choice in Section 3.1. Similarly, we specify the treatment date to be 6 months prior to when the solar array begins operating and provide supporting evidence in Section 3.2.

3.1 Spatial extent of treatment

Since the extent of treatment is unknown, we must identify d , the distance up to which the effects of constructing a solar installation persist, and this will define the boundary for our treatment group. Following similar strategies as Davis (2011), Muehlenbachs et al. (2015), and Boslett et al. (2019), we estimate a DID model similar to Equation (1), except with treatment defined in bins of successive tenth-mile increments and control always being 2-3 miles. Figure 3 plots the estimates for each tenth-mile distance bin ranging from zero to two miles. Results indicate large, negative impacts for houses within 0.1 mile, but with large standard errors. Point estimates are noisy, and some point estimates are close to zero. Bins 0.4-0.5 miles and 0.5-0.6 miles are negative and significant. This pattern is not consistent with expectations: we would expect to observe the effect gradually attenuating with distance.⁶ However, beyond 0.6 miles, all estimates are statistically insignificant. Other hedonic studies that examine the impact of solar installations on housing prices empirically estimate a treatment distance of 1 mile (Abashidze, 2019), 1 km (Dröes and Koster, 2021), and 750 m (Maddison et al., 2022). Our estimated treatment distance is almost exactly the same as the one estimated by Dröes and Koster (2021), and very close to the one estimated by Maddison et al. (2022), which gives some reassurance in our results despite the unintuitive distance decay figure. Given this evidence, we define the treatment group to be within 0.6 miles and the control group to be 0.6 – 2 miles.⁷

We only include transactions occurring within two miles of any eventual solar installation to increase similarities in observable and unobservable characteristics for sample properties. For properties lying within 0.6 miles of two installations, we omit those that transact before the closer of the two installations is built, but after the further one is built. This removes only 0.04%

⁶ To assess whether the significance of coefficients in the 0.4 – 0.5 and 0.5 – 0.6 mile bins could be a Type 1 error, we examine pre-treatment housing attributes in the online appendix. Table A1 presents pre-treatment housing attribute means. Table A2 tests for statistical differences in pre-treatment means of all housing attributes by tenth mile bins up to 1 mile, using properties between 1 – 2 miles as control. While some statistical differences exist, none point to anomalous observations in the 0.4 – 0.6 mile range. Table A3 presents the percentage and number of transactions in each distance bin by time period. There is no evidence that the volume of transactions is impacted by treatment in general nor for properties in the 0.4 – 0.6 mile range. Additionally, we conduct a test for the equality of all estimates in the first six distance bins and fail to reject the null, which provides evidence in support of the effect being constant in these bins.

⁷ Figure A2 in the online appendix plots the estimates from a similar regression, except with control defined as 1 – 2 miles. The results are qualitatively identical. Table A4 also examines robustness of results using different control groups based on different distances, and results are similar to the main findings.

of transactions and ensures a cleaner identification of the pre-treatment and post-treatment periods in our model.

3.2. Timing of treatment

The date on which treatment occurs in the minds of home buyers and sellers is ex ante unknown to us and is likely to pre-date the beginning of operation, which is the only milestone for which we have an exact date. To identify the treatment date, we conduct an event study that analyzes property price trends between the treated and control groups over time. Specifically, we define a time variable in terms of 6 month bins, starting from 6 years prior to operation date and up to 6 years post operation, and we choose 6 – 12 months prior as the reference category. We then estimate a DID model similar to Equation (1) in which we regress log sales prices on the treatment variable, the time bins, and their interaction, along with month-year and property fixed effects.

Figure 4 plots the coefficients and 95% confidence intervals of the event study model. There are two takeaways from these results. First, we find a large drop in prices in the 0-6 months prior bin relative to 6-18 months prior. The negative effect starting 6 months prior is sustained, though noisy, for the remainder of the post-operation period. As a result, we choose to define the treatment date as 6 months prior to operation date in all future specifications.⁸ This timing is in line with our expectations because it takes time for the array to be constructed, and thus disamenities will be apparent to potential buyers prior to operation. Second, while noisy, there is evidence of a pre-treatment downward trend in prices, suggesting properties near eventual solar sites may have been declining prior to construction. This trend is punctuated by the large negative difference found in the 18-24 months prior to operation time period, but is then reversed in the 6-18 months prior periods. One or both could be anomalous, but the graph raises concerns about the viability of the necessary parallel trends assumption.⁹ We discuss implications of this more in Sections 3.3-3.4.

⁸ While we have presented versions of Figures 3 and 4 with the eventual spatial extent and treatment and treatment date included, these findings are robust to different choices of one or the other.

⁹ Figure A3 in the online appendix presents a version of Figure 4 with time binned in increments of one year. There is less noise, but the qualitative findings that treatment begins six months prior to operation and relative prices are declining in treated areas pre-treatment hold.

3.3 Summary statistics and assumptions

Our final, composite dataset includes 107,291 repeat-sales transactions representing 45,795 unique properties around 282 solar installations.¹⁰ We observe 11,292 transactions within 0.6 miles, of which 34% are post-treatment.

The summary statistics for key variables are given in Table 1. The first column represents the mean values of our full sample. The mean sales price is \$314,710. The average property in our data has a lot size of 0.42 acres, has living area of just under 3000 square feet, approximately 3 bedrooms, and is about 58 years old. About 46% of the properties are matched to a greenfield development, and 35% are rural.

The critical assumption for the DID design to yield causal estimates is the parallel trends assumption, which requires that treatment and control properties would have the same trend in outcomes if treatment did not occur. We first assess the plausibility of this assumptions by comparing characteristics of treatment and control properties, with the logic that similar properties are likely to have similar price trajectories. The second and third columns in Table 1 compare pre-treatment housing attribute means between the 0 – 0.6 miles (treated) and 0.6 – 2 miles (control) observations. In the last column, we report the normalized differences in means, which is the difference in means between the treatment and control groups divided by the square root of the sum of their variances (Imbens and Wooldridge, 2009). Only one housing variable (Lot size) has a normalized difference exceeding 0.25, which is the limit beyond which the difference in means becomes substantial. However, this is not too concerning since the normalized difference just crosses the appropriate threshold. Additionally, our regression model uses property fixed effects, which effectively removes any concern about covariate overlap, except if price trends are correlated with housing characteristics.

Second, we examine pre-treatment trends in sales prices as seen in Figure 4. As discussed above, most coefficients hover near zero and are statistically insignificant in the pre-treatment period. However, the coefficient for 18-24 months prior to operation is negative, significant and large in magnitude. Additionally, there is some evidence of an overall negative trend in the pre-treatment coefficients, though the coefficients 6-12 and 12-18 months prior run counter to that

¹⁰ Our original dataset had 284 solar installations, but two are dropped because there are no repeat sales properties within 2 miles of them.

trend. Thus, the evidence is not convincing either in support or refutation of the parallel trends assumption.¹¹

Our identification strategy detailed in Equation (1) will mitigate bias from unobserved, time-invariant factors that are correlated with housing prices and solar siting. However, if the precise location of a solar array is endogenous and correlated not just with time-invariant unobserved attributes, but also correlated with price trends, then a comparison of treatment to control areas may be biased. Therefore, in the following section, we discuss an alternate DID estimator we employ that does not rely on a non-proximate control group and thus removes bias stemming from site selection being correlated with price trends.

3.4 *Alternative DID estimator*

We consider an alternative DID design that does not rely on a never-treated control group. We estimate a version of Equation (1) that includes only the 11,292 observations within 0.6 miles of an eventual solar installation. Identification in this model relies entirely on temporal variation in the construction of solar installations, instead of a combination of this variation and variation in trends between near and far houses. This is similar to the approach of Dröes and Koster (2021) who use this method in their hedonic study of solar arrays and wind turbines in the Netherlands out of concern for endogeneity of siting decisions.¹² This strategy is also recommended by Sun and Abraham (2021) for cases where never-treated units fail to satisfy the parallel trends assumption. The assumption we make when excluding never-treated observations is that locations with an eventual solar installation constructed nearby will have similar trends over time in the absence of a solar installation being constructed, and therefore not-yet-treated observations are a good counterfactual. One inconsequential change is that the variable $Post_{it}$ is

¹¹ Following Guignet and Nolte (2021), we estimate a series of regressions in Table A5 in the online appendix as an additional test for parallel trends. We include a linear trend variable (YearsPre) that measures the number of years between the treatment date and the property transaction date and interact it with Treated. We find that the coefficient on $Treated \times YearsPre$ is statistically insignificant and very small in magnitude across all regressions. This means that we cannot reject the hypothesis that the pre-treatment price trends are parallel, which lends credence to the DID assumptions.

¹² DID methods excluding the never-treated group have also been applied in other settings for alternative reasons. Beatty et al. (2021) only include treated gas stations in their preferred model of price impacts of hurricanes due to concerns about SUTVA violations. Lang and Cavanagh (2018) only include treated properties in their hedonic study of brownfield remediation because the density of brownfields made never-treated controls not proximate to treated observations and housing characteristics were dissimilar. See Callaway and Sant’Anna (2021), de Chaisemartin and D’Haultfœuille (2020), and Marcus and Sant’Anna (2020) for the theory behind using the ever-treated (or a ‘not-yet-treated’) sample as counterfactual.

collinear with $Treated_i \times Post_{it}$ and drops out from the model. We estimate the following model:

$$P_{it} = \beta_2(Treated_i \times Post_{it}) + \mathbf{X}_{it}\boldsymbol{\gamma} + \alpha_i + \epsilon_{it} \quad (2)$$

Our coefficient of interest is still β_2 and has the same interpretation. We only estimate this equation including either municipality-specific time trends or county-year fixed effects since there is no spatially proximate never-treated control group to capture local time trends.

Figure 5 presents a pre-treatment price trends analysis that tests the assumption we make for this alternative DID design. Since the never-treated observations are removed from the sample, the control group is formed by properties near eventual solar sites that have not yet been built. To create this figure, we compare the price trends by year for properties near solar arrays built early in our sample to properties near solar arrays built late in our sample. This is ad hoc, but an appropriate way to gauge pre-treatment price trend similarity. We use Figure 2 to divide the sample. First, we exclude properties near solar arrays built in 2010 and 2011 because this will allow more pre-treatment years to be examined and we lose relatively few observations by doing so. We define early treatment as an indicator variable equal to 1 if a house is proximate to a solar installation that will be built in years 2012 – 2016 and equal to zero if proximate to a solar installation that will be built in 2017 or later. We compare pre-treatment (2005-2012) price trends between properties that will have a solar installation built nearby in the years 2012-2016 (earlier-treated), to properties that will have a solar installation constructed in their vicinity in 2017 and beyond (later-treated). This approximates the identification strategy of the standard DID model with the ever-treated sample. Log sales prices are regressed on the interaction between early treatment and year dummy variables, along with month-year, property, and county-year fixed effects. Estimated coefficients giving differences between early treatment and early control properties over time and 95% confidence intervals are graphed. Similar to Figure 4, a graph showing no statistical differences between the earlier and later treated groups' pre-treatment trends supports the assumption that the not-yet-treated group is a strong counterfactual for the post-treatment group. We find that price trends are similar between the two groups, as evidenced by statistically insignificant coefficients (at the 5% level) in each year. The coefficient in 2008 is just barely insignificant (with a p-value of 0.051), but no other coefficient is remotely close to significant, suggesting that our assumption holds. There is also no downward trend that is present in Figure 4 (the standard DID model).

In Section 4, we present results using both the DID model that includes the never-treated properties and the DID model that excludes the never-treated properties. We remain equivocal about which is preferred and instead focus on the range of estimates.¹³

3.5 Heterogeneity in treatment effect

We extend the analysis to investigate heterogeneity in treatment effect in multiple ways. First, we investigate heterogeneity in treatment effect by two place-based characteristics: prior land use and rural character. This is done by interacting the treatment effect term in Equation (1) with variables for our characteristic of interest. The specifications are as follows:

$$P_{it} = \beta_1(Post_{it} \times Non - greenfield_i) + \beta_2(Treated_i \times Post_{it} \times Non - greenfield) + \beta_3(Post_{it} \times Greenfield_i) + \beta_4(Treated_i \times Post_{it} \times Greenfield_i) + \mathbf{X}_{it}\boldsymbol{\gamma} + \alpha_i + \epsilon_{it} \quad (3)$$

$$P_{it} = \beta_1(Post_{it} \times Non - rural_i) + \beta_2(Treated_i \times Post_{it} \times Non - rural) + \beta_3(Post_{it} \times Rural_i) + \beta_4(Treated_i \times Post_{it} \times Rural_i) + \mathbf{X}_{it}\boldsymbol{\gamma} + \alpha_i + \epsilon_{it} \quad (4)$$

where $Greenfield_i$ is an indicator variable equal to 1 if a property is located within the vicinity of a solar installation that was built on land that was formerly farmland or forested and $Non - greenfield_i = 1 - Greenfield_i$. $Rural_i$ is an indicator variable equal to 1 if property i lies in a town with a population density of 850 people per square mile or fewer, and $Non - rural_i = 1 - Rural_i$. The corresponding equations for the ever-treated sample are similar, except that the $Post_{it}$ variable drops out on account of collinearity.

Our coefficients of interest in Equations (3) and (4) are β_2 and β_4 . In Equation (3), we hypothesize that $\beta_4 < \beta_2 < 0$ because developments on farm and forest lands will lead to larger negative impacts on housing prices due to the more dramatic change in landscape compared to a

¹³ An additional concern with DID models with staggered treatment is that estimated coefficients can be biased if treatment effects are heterogeneous over time and some observations have negative weights (de Chaisemartin and D'Haultfoeuille, 2020). We analyze our data for the presence of negative weights and find relatively few. Applying the `twowayfweights` command in Stata (de Chaisemartin et al., 2019), 13.8% of treated observations have an associated negative weight and the sum of negative weights is -0.005. This compares favorably to the case study data used by de Chaisemartin and d'Haultfoeuille (2020) in which 40.1% of treated observations have an associated negative weight and the sum of negative weight is -0.533. Further, the time corrected wald estimator proposed by de Chaisemartin and D'Haultfoeuille (2020) produces estimates qualitatively identical to the standard DID. Thus, we are not concerned about this particular source of bias.

commercial site or landfill and the loss of open space amenities. In Equation (4), we again hypothesize that $\beta_4 < \beta_2 < 0$ because solar arrays are less congruent with rural settings and the contrast will lead to greater price declines, but there's more uncertainty here because of land scarcity in non-rural areas.

Second, we estimate a model that allows for heterogeneity in the impact based on distance. We identified treatment extending to 0.6 miles in Figure 3, but Figure 3 also suggests that treatment effects could be larger within 0.1 mile. To explore this possibility more formally, we develop a model that defines multiple distance bands. The first (outermost) band represents control properties located 1 – 2 miles away from the nearest solar installation. The second band is properties 0.6 – 1 mile away, which we differentiate from 1 – 2 miles to further test if the spatial extent of treatment does end at 0.6 miles. The third band includes treated properties located 0.1 – 0.6 miles from the nearest solar installation. Finally, the fourth (innermost) band consists of treated properties within a distance of 0.1 mile from the closest installation. Our specification is:

$$P_{it} = \beta_1 Post_{it} + \sum_{k=2}^4 \beta_2^k (dist_i^k \times Post_{it}) + \mathbf{X}_{it}\boldsymbol{\gamma} + \alpha_i + \epsilon_{it} \quad (5)$$

where $dist_i^k$ is a dummy variable equal to 1 if a property i lies within the k^{th} distance band. P_{it} , $Post_{it}$, X_{it} , and α_i are as defined in Equation (1). When estimating this model excluding never-treated properties, we only get estimates on the two inner rings, 0-0.1 miles and 0.1-0.6 miles.

4 RESULTS

4.1 Average treatment effects

We present our main results in Table 2. Columns 1 – 3 present the results obtained from estimating Equation (1) that includes the never-treated control group (distances of 0.6 – 2 miles). Columns 4 and 5 exclude the never-treated properties and present the results estimating Equation (2). All columns include month-year fixed effects and property fixed effects, Columns 2 and 4 additionally include municipality-year time trends, and Columns 3 and 5 replace those with county-year fixed effects. Including never-treated properties yields treatment effect coefficient estimates that range from -0.015 to -0.024. Excluding the never-treated properties yields coefficients that are about twice as large, ranging from -0.028 to -0.036. The smaller magnitudes observed in Columns 1 – 3 likely stem from the pre-treatment, downward trend in treated properties relative to never-treated properties seen in Figure 4. Overall, treatment effects are

negative and statistically significant across all models, confirming our hypothesis that nearby solar installations are, on average, a disamenity. Estimates suggest that houses lying within 0.6 miles of solar installations sell between 1.5% and 3.6% less post construction, all else equal.

We convert the percentage reduction to dollars by multiplying the coefficient and the average, pre-treatment property price for treated properties (\$314,710), which gives us a range of \$4,721 - \$11,330 (prices adjusted to 2019 levels).

4.2 *Heterogeneous treatment effects*

In Table 3, we examine the heterogeneity in treatment effect by three characteristics: prior land use, rural character of towns, and proximity to solar installations. Each panel presents two specifications, mirroring the sample and control variables in Columns 3 and 5 of Table 2.

In Panel A, we provide estimates from the model described by Equation (3) where we explore heterogeneity by prior land use. The results conform to our expectations; estimated treatment effects for greenfield and non-greenfield sites are both negative, but the treatment effects for greenfield sites is larger in magnitude. The coefficient on $Treated \times Post \times Greenfield$ ranges from -0.020 to -0.044 and is significant at the 5% level or higher in both specifications. In contrast, the coefficient on $Treated \times Post \times Non - greenfield$ ranges from -0.011 to -0.013 and is not statistically significant in either column.

There are two questions that arise from these results. First, do non-greenfield sites have zero externalities? Statistically, yes, we fail to reject a null hypothesis of no effect. However, the coefficients are consistently negative (also holds in Table 4 discussed below), so there may be some signal there, just not enough to overcome the noise. Additionally, across non-greenfield sites, there could be additional heterogeneity that we are unable to measure. For example, different arrays could have varying degrees of visibility. To improve our understanding, we can draw on Gaur et al. (2022), who recently conducted a choice experiment survey on preferences for solar siting attributes in Rhode Island. They estimate separate models for greenfield and non-greenfield solar sites and find that respondents have positive WTP to avoid fully visible arrays for both types (\$10.34/month for greenfields and \$4.42/month for non-greenfields).¹⁴ In addition,

¹⁴ This is not to be confused with Gaur et al. (2022)'s finding of a positive total WTP value for siting arrays on non-greenfield land types. This result is not related to site proximity and instead reflects overall preferences for siting solar arrays on non-greenfield land types.

respondents prefer further setback from property boundaries for non-greenfield sites, but are indifferent about setback on greenfield sites. Thus, these choice experiment results indicate that negative externalities can be present at non-greenfield sites. We argue that our hedonic estimates reflect those negative externalities, though we cannot be confident in the exact magnitude of those effects.

The second question raised by the results of Panel A Table 3 is whether the difference of the greenfield treatment effect relative to the non-greenfield treatment effect is driven entirely by loss of open space. This is a critical question because if the alternative to solar arrays is residential housing and that leads to the same disamenities, then there is no reason to be concerned about solar developments. In truth, we cannot definitively know, but we argue there are attributes of a solar array that lead to additional negative externalities beyond residential development of open space. Some portion of the wedge could be due to nearby residents feeling that solar arrays are incongruent with that type of landscape and it takes away from the aesthetic of that place in ways that common houses do not. Additionally, we can again point to the Gaur et al. (2022) results that show greater negative viewshed externalities on greenfield sites relative to non-greenfield sites, and these are on top of already substantial WTP to avoid development on greenfield sites to begin with. Relatedly, an unintended byproduct of this analysis is providing an upper bound on the value of privately held open space. Irwin (2002) and Geoghegan et al. (2006) both examine the property value impacts of developable open space (as well as permanently conserved open space), but they use a cross sectional approach, and their estimates vary substantially across models with some indicating developable open space is valued more than residential development and some the opposite. Our research offers better identification and bounds the impacts of loss of nearby developable open space as a small negative.

In Panel B, we examine heterogeneity by rural character of towns and report the coefficients from the specification defined in Equation (4). Similarly, these results conform to our expectations; estimated treatment effects for rural and non-rural sites are both negative and the treatment effect for rural sites is on average larger in magnitude. The coefficient on $Treated \times Post \times Rural$ ranges from -0.025 to -0.058 and is significant at the 5% level or higher in both columns. In contrast, the coefficient on $Treated \times Post \times Non - rural$ ranges

from -0.005 to -0.006 and is not statistically significant. The results suggest that nearby utility-scale solar causes housing prices to decline more in rural areas than suburban or urban areas.¹⁵

There is, as expected, a strong positive correlation (0.41) between greenfield and rural, which raises the question of whether large negative results observed in rural areas are just a function of the higher proportion of greenfields found there or vice versa. Table A7 in the online appendix estimates a quadruple interaction model to try to parse the effects of greenfields and arrays in rural areas. It is clear that the smallest impacts accrue to properties near non-rural, non-greenfield sites – in fact, we cannot reject no effect across all models. However, other orderings are inconsistent across columns, with each of three other categories yielding the largest negative impact in at least one specification. These results suggest we cannot attribute the results to greenfield sites or rural sites alone. Instead, there appears to be an additive effect. In sum, the results of Panels A and B indicate that valuation depends on context; surrounding land uses and place meaning contribute to the magnitude of price declines.

Lastly, in Panel C, we estimate the model described by Equation (5) that allows for heterogeneity in the impact on prices based on distance. The coefficient for the 0.6 – 1 mile band is statistically insignificant in Column 1, which is consistent with our assumption that treatment effects do not persist beyond 0.6 miles. The coefficients on the 0.1 – 0.6 mile band are significant and similar magnitude to the main results. The coefficients on the 0 – 0.1 mile band range between -0.038 to -0.042, which is between 1.5 to 2.4 times larger in magnitude than the 0.1 – 0.6 mile band, though insignificant. These results are suggestive that property values for homes lying within 0.1 mile from a solar installation may fall substantially, but our estimates are imprecise reflecting few observations within that distance band.¹⁶

In the online appendix, we also present results that test for heterogeneity by size of installation and time since construction (see Tables A8 and A9). We find that there are no statistically significant differences between categories, and results suggest that larger installations do not cause greater price declines and that treatment effects do not dissipate with time.

¹⁵ We examine different population density cutoffs for the definition of Rural in Table A6 in the online appendix. Results are consistent across different cutoffs.

¹⁶ There are 218 observations lying within 0.1 miles from a solar installation, of which 72 sell post construction.

4.3 Robustness checks

In Table 4 we present results from a series of robustness checks to ensure that our results are consistent to alternative data constructions or samples. We present results both for the average treatment effect models and models focused on greenfield heterogeneity. We include the latter because it is a critical piece of the story. Further, we present results for both the models that include the never-treated control properties (Panels A and B) and exclude the never-treated properties (Panels C and D) do the same, except using the ever-treated sample.

Columns 1 and 2 explore the assumption of the amount of land required per MW of installed capacity. Instead of 5 acres in our main models, Column 1 assumes 4 acres, and Column 2 assumes 6 acres. By contracting and expanding the assumed size of installations, the set of properties that are designated as treatment and control are altered. The estimated coefficients in these columns are qualitatively identical to the main and heterogeneity results, indicating that assumptions about the radius of arrays is not impacting results.

In Column 3, we control for the presence of wind turbines by including an indicator variable equal to 1 if a house lies within one mile of a built wind turbine. One may be concerned that solar and wind are co-located and disamenities from one may be captured in the estimated valuation of the other if not controlled for. The treatment coefficient is nearly identical to the main results. In MA and RI, there is little correlation in the siting of wind and solar energy, and solar is far more abundant (see Figure A4 in the online appendix).

Our main sample includes transactions in years 2005-2019. One may be concerned that this is too long of a time horizon and changes to the hedonic function can occur over that time. To address this concern, Column 4 only includes transactions occurring 2009-2019, and Column 5 only includes transactions that are within four years before or after the treatment date of the solar installation they are matched to. Both of these sample restrictions, particularly Column 5, greatly reduce our sample size in the repeat sales model because fewer properties transact multiple times in a short window. The average treatment effect estimates are larger, and the greenfield treatment effects are over twice as large.

In the online appendix, we check the robustness of our main results in four more ways. First, in Table A10 we test for anticipation effects two years prior to solar farm operation date and find no evidence of anticipation. Second, for the model that includes the never-treated properties, we vary the spatial extent of the control group (Table A4) and find that the treatment

effect is robust to different control group boundaries. Third, in Table A11 we examine whether regional price trends may be correlated with solar installation construction by including distance to city center by year trends in all our specifications. Our coefficients remain robust, suggesting that this is not a threat to identification. Finally, in Tables A12 – A14, we consider an alternative estimation approach for the sample that includes never-treated observations – the Coarsened Exact Matching (CEM) method (Blackwell et al., 2009; Iacus et al., 2012). This is a pre-regression matching approach that reweights our sample so that the treatment and control groups are balanced. We consider balancing by two spatial designations (county and solar site) and one temporal designation (year). Our findings suggest that the matching procedure improves multivariate and univariate balance by a very small magnitude, leading to results that are very robust to our main findings.

5 CONCLUSION

This paper estimates the valuation of externalities associated with nearby utility-scale solar installations using revealed preferences from the property market. Using the DID empirical technique, we define treatment by distance to the nearest solar installation, and compare treated properties to those lying between 0.6 and 2 miles from the installation (never-treated group), or to properties that receive a solar installation in their vicinity in the future. We observe 11,292 housing transactions occurring within 0.6 miles (treated group), and 95,999 transactions between 0.6 and two miles (never-treated control group) of 282 solar installations in MA and RI. One caveat of our analysis is that we find mixed evidence regarding the pre-treatment trends being parallel. Since we are unable to either confirm or deny whether the parallel trends assumption holds, we advise caution in interpreting our results as causal.

Our findings can be summarized as follows: there is a consistent negative average effect of proximity to utility-scale solar array, the estimates derived using the ever-treated sample are consistently larger than the ones that use the never-treated control group, and arrays sited on greenfields in rural areas cause larger negative impacts and drive the overall negative and significant average effects. Average treatment effects suggests that property values decline between 1.5% (\$4,721) and 3.6% (\$11,330) after the construction of a nearby solar installation, all else equal.

While a full benefit-cost analysis (BCA) of utility-scale solar arrays is beyond the scope of this paper, because we do not know anything about consumer and producer surplus¹⁷, we can at least benchmark the negative, local externalities against the global benefits of greenhouse gas (GHG) mitigation. We therefore conduct the following back-of-the-envelope calculations. While solar arrays typically have a lifetime of 25-30 years, there is uncertainty about what would happen after that time. Thus, we ignore those dynamic issues and only calculate costs and benefits for a single year. On the cost side, we first consider the point estimate from our preferred specifications, which translates to a loss between \$4,721 and \$11,330 per household (prices adjusted to 2019 levels). Assuming a 5% interest rate and an infinite time horizon, we get an annualized dollar value between \$236 and \$567 per household for treated homes close to solar installations.¹⁸ Our complete dataset (prior to any sample cuts) consists of 72,538 unique properties located within 0.6 miles of all solar installations in the dataset. Put together, we estimate an annual loss between \$17.12 and 41.13 million due to proximate solar installations in MA and RI.

To quantify the GHG benefits from solar installations, we first calculate net generation from solar installations. Assuming a capacity factor of 16.5%, the 817 MW of installed solar capacity in MA and RI generates is 1,180,892 MWh (megawatt hours) of electricity per year.¹⁹ Current non-renewable generation in MA and RI comes almost entirely from natural gas. According to the EIA, 0.42 mt (metric tons) of CO₂ are emitted from each MWh of electricity that is generated from natural gas, implying that a total of 495,975 mt of CO₂ are abated annually from solar energy generation. In addition, natural gas can leak in the distribution system, which releases methane, a much more potent greenhouse gas. Based on Hausman and Muehlenbachs (2018) and EIA, each MWh generated from natural gas is associated with 104.72 cf (cubic feet) of methane leaked. Thus, one year of solar generation mitigates an estimated 123,663 mcf of leaked methane, which has equivalent warming potential of 48,538 mt of CO₂. The EPA

¹⁷ To be sure, significant amounts of money are part of the market transactions. A developer quoted us that they offer landowners \$15-20,000 per MW per year of installed capacity. It is unknown how much is profit and whether some portion of that could be used to compensate proximate households.

¹⁸ We annualize the total property value impact by using the present value formula of a perpetuity: PV (present value) = Annual Cash Flow/Interest rate. Assuming an interest rate of 5%, we calculate $0.05 * (\$4,721)$ and $(0.05 * \$11,330)$ and get a property equity loss between \$236 and \$567 per household.

¹⁹ *Net generation (MWh) = % Capacityfactor × 365 days × 24 hours × Installed capacity (MW)*

estimates the current social cost of CO₂ is \$51.80 per metric ton, which places the value of annual greenhouse gas mitigation to be \$28.21 million (US EPA).

Combining the estimates, the benefit-cost ratio is between 1.65 and 0.69. In one scenario (using the DID estimates including the never-treated properties), the global benefits outweigh the local costs. However, using the DID estimates that exclude the never-treated properties, we come to the opposite conclusion. Regardless, in both cases it is clear that the local costs are substantial, bolstering local concerns about solar siting and clarifying the magnitude of costs borne by neighboring property owners. However, the benefit-cost ratio may be substantially better in other states that are less densely populated (fewer impacted houses) or more reliant on coal (greater carbon emissions displacement).

This research offers policy relevant findings. Communities in southern New England and elsewhere in the United States are currently grappling with contentious solar siting issues and will be for some time. These results quantify some of the opposition to certain siting decisions and allow those voices to enter into a state or local BCA. Further, our results suggest ways to reduce negative externalities that could be activated by state and local governments. In the case of siting on brownfields and covered landfills, developers may require additional subsidies to target those areas. Though non-financial costs, such as faster zoning approval may compensate them as well.

There are several directions of important future research. Similar hedonic studies should be completed elsewhere in the United States to assess similarity of valuation estimates and test our assertion that benefit-cost ratios will be more favorable elsewhere. Though, as discussed above, Abashidze (2019) finds even larger impacts than ours in North Carolina. In addition, examining valuation of smaller solar arrays (100 kW – 500 kW) could yield new insights. In southern New England, farms can install arrays of this size on marginal land and generate income that can help sustain the farm in the face of rising land costs (EcoRI, 2020). Our results may suggest that moving solar development away from greenfields and away from rural areas is a possible means to achieve an objective to minimize the average household external cost. However, implementing such a strategy could actually increase total external costs depending on the density of housing near alternative sites. Future research could also investigate community solar, which is a popular idea that is understudied in the context of siting. For example, a

contingent valuation survey could assess willingness to accept proximate community solar if it was structured such that nearby residents benefited financially.

REFERENCES

- Abashidze, N. (2019). *Essays on Economic and Health Effects of Land Use Externalities*. [North Carolina State University]. <https://repository.lib.ncsu.edu/handle/1840.20/38420>
- Abashidze, N., and Taylor, L. O. (2022). Utility-Scale Solar Farms and Agricultural Land Values. *Land Economics*. <https://doi.org/10.3368/le.99.3.102920-0165R>
- Banzhaf, H. S. (2021). Difference-in-Differences Hedonics. *Journal of Political Economy*, 129(8), 2385–2414. <https://doi.org/10.1086/714442>
- Bates, A., and Firestone, J. (2015). A comparative assessment of proposed offshore wind power demonstration projects in the United States. *Energy Research and Social Science*, 10, 192–205. <https://doi.org/10.1016/j.erss.2015.07.007>
- Beatty, T. K. M., Lade, G. E., and Shimshack, J. (2021). Hurricanes and Gasoline Price Gouging. *Journal of the Association of Environmental and Resource Economists*, 8(2), 347–374. <https://doi.org/10.1086/712419>
- Bishop, K. C., Kuminoff, N. V., Banzhaf, H. S., and Boyle, K. J. (2019). Best Practices in Using Hedonic Property Value Models for Welfare Measurement. *Review of Environmental Economics and Policy*, 43.
- Blackwell, M., Iacus, S., King, G., and Porro, G. (2009). Cem: Coarsened Exact Matching in Stata. *The Stata Journal*, 9(4), 524–546. <https://doi.org/10.1177/1536867X0900900402>
- Boslett, A., Guilfoos, T., and Lang, C. (2019). Valuation of the External Costs of Unconventional Oil and Gas Development: The Critical Importance of Mineral Rights Ownership. *Journal of the Association of Environmental and Resource Economists*, 6(3), 531–561. <https://doi.org/10.1086/702540>
- Botelho, A., Lourenço-Gomes, L., Pinto, L., Sousa, S., and Valente, M. (2017). Accounting for local impacts of photovoltaic farms: The application of two stated preferences approaches to a case-study in Portugal. *Energy Policy*, 109, 191–198. <https://doi.org/10.1016/j.enpol.2017.06.065>
- Brittan, G. G. (2001). Wind, energy, landscape: Reconciling nature and technology. *Philosophy and Geography*, 4(2), 169–184. <https://doi.org/10.1080/10903770124626>
- Callaway, B., and Sant’Anna, P. H. C. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230. <https://doi.org/10.1016/j.jeconom.2020.12.001>
- Carlisle, J. E., Kane, S. L., Solan, D., Bowman, M., and Joe, J. C. (2015). Public attitudes regarding large-scale solar energy development in the U.S. *Renewable and Sustainable Energy Reviews*, 48, 835–847. <https://doi.org/10.1016/j.rser.2015.04.047>
- Carlisle, J. E., Kane, S. L., Solan, D., and Joe, J. C. (2014). Support for solar energy: Examining sense of place and utility-scale development in California. *Energy Research and Social Science*, 3, 124–130. <https://doi.org/10.1016/j.erss.2014.07.006>

- Carr-Harris, A., and Lang, C. (2019). Sustainability and tourism: The effect of the United States' first offshore wind farm on the vacation rental market. *Resource and Energy Economics*, 57, 51–67. <https://doi.org/10.1016/j.reseneeco.2019.04.003>
- Dastrup, S. R., Graff Zivin, J., Costa, D. L., and Kahn, M. E. (2012). Understanding the Solar Home price premium: Electricity generation and “Green” social status. *European Economic Review*, 56(5), 961–973. <https://doi.org/10.1016/j.euroecorev.2012.02.006>
- Davis, L. W. (2011). THE EFFECT OF POWER PLANTS ON LOCAL HOUSING VALUES AND RENTS. *The Review of Economics and Statistics*, 93(4), 1391–1402.
- de Chaisemartin, C., and D’Haultfoeuille, X. (2020). Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, 110(9), 2964–2996. <https://doi.org/10.1257/aer.20181169>
- de Chaisemartin, C., D’Haultfoeuille, X., and Guyonvarch, Y. (2019). Fuzzy differences-in-differences with Stata. *The Stata Journal*, 19(2), 435–458. <https://doi.org/10.1177/1536867X19854019>
- Denholm, P., and Margolis, R. M. (2008). Land-use requirements and the per-capita solar footprint for photovoltaic generation in the United States. *Energy Policy*, 36(9), 3531–3543. <https://doi.org/10.1016/j.enpol.2008.05.035>
- Devine-Wright, P. (2009). Rethinking NIMBYism: The role of place attachment and place identity in explaining place-protective action. *Journal of Community and Applied Social Psychology*, 19(6), 426–441. <https://doi.org/10.1002/casp.1004>
- Dong, L., & Lang, C. (2022). Do views of offshore wind energy detract? A hedonic price analysis of the Block Island wind farm in Rhode Island. *Energy Policy*, 167, 113060.
- Dröes, M. I., and Koster, H. R. A. (2016). Renewable energy and negative externalities: The effect of wind turbines on house prices. *Journal of Urban Economics*, 96, 121–141. <https://doi.org/10.1016/j.jue.2016.09.001>
- Dröes, M. I., and Koster, H. R. A. (2021). Wind turbines, solar farms, and house prices. *Energy Policy*, 155, 112327. <https://doi.org/10.1016/j.enpol.2021.112327>
- EcoRI. (2020). 6 Acres of Solar Energy to Help Gloucester Farm Pay Bills. *EcoRI News*. <https://www.ecori.org/renewable-energy/2020/8/7/solar-energy-helps-farm-endure>
- EIA - Annual Energy Outlook 2018. (2018, July 17). <https://www.eia.gov/outlooks/aeo/>
- EIA International Energy Outlook. (2019). U.S. Energy Information Administration. <https://www.eia.gov/outlooks/ieo/>
- Farhar, B. C. (1994). Trends in US Public Perceptions and Preferences on Energy and Environmental Policy. *Annual Review of Energy and the Environment*, 19(1), 211–239. <https://doi.org/10.1146/annurev.eg.19.110194.001235>
- Farhar, B. C., Hunter, L. M., Kirkland, T. M., and Tierney, K. J. (2010). *Community Response to Concentrating Solar Power in the San Luis Valley: October 9, 2008 - March 31, 2010* (NREL/SR-550-48041). National Renewable Energy Lab. (NREL), Golden, CO (United States). <https://doi.org/10.2172/983406>
- Firestone, J., Bidwell, D., Gardner, M., and Knapp, L. (2018). Wind in the sails or choppy seas?: People-place relations, aesthetics and public support for the United States' first offshore wind project. *Energy Research and Social Science*, 40, 232–243. <https://doi.org/10.1016/j.erss.2018.02.017>
- Gaur, V., Lang, C., Howard, G., and Quainoo, R. (2022). When Energy Issues are Land Use Issues: Estimating Preferences for Utility-Scale Solar Energy Siting. *Land Economics*. <https://doi.org/10.3368/le.99.3.11221-0130R1>

- Geoghegan, J., Lynch, L., and Bucholtz, S. J. (2006). Are Agricultural Land Preservation Programs Self-financing? In *Economics and Contemporary Land Use Policy* (pp. 175–195). Routledge. <https://doi.org/10.4324/9781936331659-17>
- Gibbons, S. (2015). Gone with the wind: Valuing the visual impacts of wind turbines through house prices. *Journal of Environmental Economics and Management*, 72, 177–196. <https://doi.org/10.1016/j.jeem.2015.04.006>
- Greenberg, M. (2009). Energy sources, public policy, and public preferences: Analysis of US national and site-specific data. *Energy Policy*, 37(8), 3242–3249. <https://doi.org/10.1016/j.enpol.2009.04.020>
- Gross, S. (2020). *Renewables, land use, and local opposition in the United States* (p. 24). The Brookings Institution. https://docs.wind-watch.org/FP_20200113_renewables_land_use_local_opposition_gross.pdf
- Guignet, D., and Nolte, C. (2021). *Hazardous Waste and Home Values: An Analysis of Treatment and Disposal Sites in the U.S.* 56.
- Hausman, C., and Muehlenbachs, L. (2018). Price Regulation and Environmental Externalities: Evidence from Methane Leaks. *Journal of the Association of Environmental and Resource Economists*, 6(1), 73–109. <https://doi.org/10.1086/700301>
- Hoen, B., and Atkinson-Palombo, C. (2016). Wind Turbines, Amenities and Disamenities: A Study of Home Value Impacts in Densely Populated Massachusetts. *Journal of Real Estate Research*, 38(4), 473–504. <https://doi.org/10.5555/0896-5803-38.4.473>
- Hoen, B., Brown, J. P., Jackson, T., Thayer, M. A., Wisner, R., and Cappers, P. (2015). Spatial Hedonic Analysis of the Effects of US Wind Energy Facilities on Surrounding Property Values. *The Journal of Real Estate Finance and Economics*, 51(1), 22–51. <https://doi.org/10.1007/s11146-014-9477-9>
- Hoen, B., Firestone, J., Rand, J., Elliott, D., Hübner, G., Pohl, J., Wisner, R., Lantz, E., Haac, R., and Kaliski, K. (2019). Attitudes of U.S. Wind Turbine Neighbors: Analysis of a Nationwide Survey. *Energy Policy*, 134. <https://doi.org/10.1016/j.enpol.2019.110981>
- Hoen, B., Wisner, R., Thayer, M., and Cappers, P. (2012). Residential Photovoltaic Energy Systems in California: The Effect on Home Sales Prices. *Contemporary Economic Policy*, 31(4), 708–718. <https://doi.org/10.1111/j.1465-7287.2012.00340.x>
- Iacus, S. M., King, G., and Porro, G. (2012). Causal Inference without Balance Checking: Coarsened Exact Matching. *Political Analysis*, 20(1), 1–24. <https://doi.org/10.1093/pan/mpr013>
- Imbens, G. W., and Wooldridge, J. M. (2009). Recent Developments in the Econometrics of Program Evaluation. *Journal of Economic Literature*, 47(1), 5–86. <https://doi.org/10.1257/jel.47.1.5>
- Irwin, E. G. (2002). The Effects of Open Space on Residential Property Values. *Land Economics*, 78(4), 465–480. <https://doi.org/10.2307/3146847>
- Jacobe, D. (2013, March 27). *Americans Want More Emphasis on Solar, Wind, Natural Gas*. Gallup.Com. <https://news.gallup.com/poll/161519/americans-emphasis-solar-wind-natural-gas.aspx>
- Jarvis, S. (2021). *The economic costs of NIMBYism: Evidence from renewable energy projects*. 2.
- Jensen, C. U., Panduro, T. E., Lundhede, T. H., Nielsen, A. S. E., Dalsgaard, M., and Thorsen, B. J. (2018). The impact of on-shore and off-shore wind turbine farms on property prices. *Energy Policy*, 116, 50–59. <https://doi.org/10.1016/j.enpol.2018.01.046>

- Kim, K. J., Lee, H., and Koo, Y. (2020). Research on local acceptance cost of renewable energy in South Korea: A case study of photovoltaic and wind power projects. *Energy Policy*, 144, 111684. <https://doi.org/10.1016/j.enpol.2020.111684>
- Krohn, S., and Damborg, S. (1999). On public attitudes towards wind power. *Renewable Energy*, 16(1), 954–960. [https://doi.org/10.1016/S0960-1481\(98\)00339-5](https://doi.org/10.1016/S0960-1481(98)00339-5)
- Kuffner, A. (2018, March 16). *Worry over solar sprawl spreads across Rhode Island*. Providencejournal.Com. <https://www.providencejournal.com/news/20180316/worry-over-solar-sprawl-spreads-across-rhode-island>
- Lang, C. (2018). Assessing the efficiency of local open space provision. *Journal of Public Economics*, 158, 12–24. <https://doi.org/10.1016/j.jpubeco.2017.12.007>
- Lang, C., and Cavanagh, P. (2018). Incomplete Information and Adverse Impacts of Environmental Cleanup. *Land Economics*, 94(3), 386–404. <https://doi.org/10.3368/le.94.3.386>
- Lang, C., Opaluch, J. J., and Sfinarolakis, G. (2014). The windy city: Property value impacts of wind turbines in an urban setting. *Energy Economics*, 44, 413–421. <https://doi.org/10.1016/j.eneco.2014.05.010>
- Lovich, J. E., and Ennen, J. R. (2011). Wildlife Conservation and Solar Energy Development in the Desert Southwest, United States. *BioScience*, 61(12), 982–992. <https://doi.org/10.1525/bio.2011.61.12.8>
- Ma, C., Polyakov, M., and Pandit, R. (2016). Capitalisation of residential solar photovoltaic systems in Western Australia. *Australian Journal of Agricultural and Resource Economics*, 60(3), 366–385. <https://doi.org/10.1111/1467-8489.12126>
- Maddison, D., Ogier, R., and Beltrán, A. (2022). The Disamenity Impact of Solar Farms: A Hedonic Analysis. *Land Economics*, 071220. <https://doi.org/10.3368/le.071220-0105R>
- Marcus, M., and Sant’Anna, P. H. C. (2020). The Role of Parallel Trends in Event Study Settings: An Application to Environmental Economics. *Journal of the Association of Environmental and Resource Economists*, 8(2), 235–275. <https://doi.org/10.1086/711509>
- Muehlenbachs, L., Spiller, E., and Timmins, C. (2015). The Housing Market Impacts of Shale Gas Development. *American Economic Review*, 105(12), 3633–3659. <https://doi.org/10.1257/aer.20140079>
- Naylor, D. (2019, October 16). *West Greenwich residents air concerns over proposed solar project*. Providencejournal.Com. <https://www.providencejournal.com/news/20191016/west-greenwich-residents-air-concerns-over-proposed-solar-project>
- Nolte, C., Boyle, K. J., Chaudhry, A. M., Clapp, C. M., Guignet, D., Hennighausen, H., Kushner, I., Liao, Y., Mamun, S., Pollack, A., Richardson, J., Sundquist, S., Swedberg, K., and Uhl, J. H. (2021). *Studying the Impacts of Environmental Amenities and Hazards with Nationwide Property Data: Best Data Practices for Interpretable and Reproducible Analyses* (SSRN Scholarly Paper No. 3900806). <https://doi.org/10.2139/ssrn.3900806>
- Ong, S., Campbell, C., Denholm, P., Margolis, R., and Heath, G. (2013). *Land-Use Requirements for Solar Power Plants in the United States* (NREL/TP-6A20-56290, 1086349). <https://doi.org/10.2172/1086349>
- Perea, A., Smith, C., Davis, M., Sun, X., White, B., Cox, M., Curtin, G., Rumery, S., Holm, A., Goldstein, R., and Baca, J. (2020). *U.S. Solar Market Insight Executive summary*. Wood Mackenzie and Solar Energy Industries Association.

- Pew Research Center. (2019, November 25). *U.S. Public Views on Climate and Energy*.
<https://www.pewresearch.org/science/2019/11/25/u-s-public-views-on-climate-and-energy/>
- Qiu, Y., Wang, Y. D., and Wang, J. (2017). Soak up the sun: Impact of solar energy systems on residential home values in Arizona. *Energy Economics*, 66, 328–336.
<https://doi.org/10.1016/j.eneco.2017.07.001>
- Russell, A., Firestone, J., Bidwell, D., and Gardner, M. (2020). Place meaning and consistency with offshore wind: An island and coastal tale. *Renewable and Sustainable Energy Reviews*, 132, 110044. <https://doi.org/10.1016/j.rser.2020.110044>
- Sun, L., and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199.
<https://doi.org/10.1016/j.jeconom.2020.09.006>
- Sunak, Y., and Madlener, R. (2016). The impact of wind farm visibility on property values: A spatial difference-in-differences analysis. *Energy Economics*, 55, 79–91.
<https://doi.org/10.1016/j.eneco.2015.12.025>
- Trainor, A. M., McDonald, R. I., and Fargione, J. (2016). Energy Sprawl Is the Largest Driver of Land Use Change in United States. *PLOS ONE*, 11(9), e0162269.
<https://doi.org/10.1371/journal.pone.0162269>
- US Census Bureau. (2020). *Historical Population Density Data (1910-2020)*. Census.Gov.
<https://www.census.gov/data/tables/time-series/dec/density-data-text.html>
- U.S. Energy Information Administration (EIA). (n.d.). Retrieved June 24, 2020, from
<https://www.eia.gov/energyexplained/electricity/electricity-in-the-us-generation-capacity-and-sales.php>
- US EPA. (n.d.). *The Social Cost of Carbon* [Reports and Assessments]. Retrieved July 9, 2020, from /climatechange/social-cost-carbon
- Vyn, R. J. (2018). Property Value Impacts of Wind Turbines and the Influence of Attitudes toward Wind Energy. *Land Economics*, 94(4), 496–516.
<https://doi.org/10.3368/le.94.4.496>
- Wee, S. (2016). The effect of residential solar photovoltaic systems on home value: A case study of Hawai‘i. *Renewable Energy*, 91, 282–292.
<https://doi.org/10.1016/j.renene.2016.01.059>
- Wester-Herber, M. (2004). Underlying concerns in land-use conflicts—The role of place-identity in risk perception. *Environmental Science and Policy*, 7(2), 109–116.
<https://doi.org/10.1016/j.envsci.2003.12.001>

Figures and Tables

Figure 1: Map of utility-scale solar installations across Massachusetts and Rhode Island

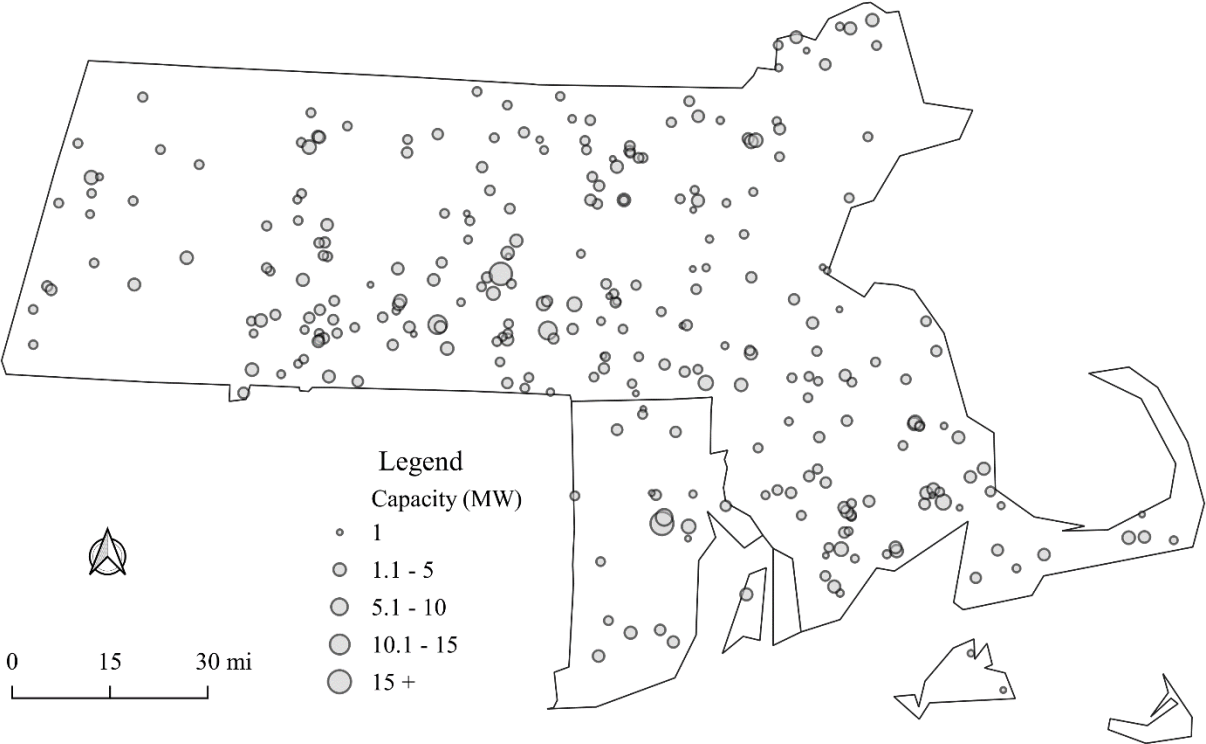


Figure 2: New and cumulative utility-scale solar capacity by year






Figure 3: Spatial extent of treatment




Notes: The treatment variable is defined as a bin variable, with treated properties lying within 1/10 mile distance bands up to 2 miles. Control properties are those lying 2 – 3 miles away from the nearest solar installation. Post = 1 if a house sells after the treatment date. The treatment date is defined a 6 months prior to solar installation operation date. The coefficients are obtained by estimating a DID model similar to Equation 1 that regresses log sales price on 1/10 mile distance bands up to 2 miles, along with month-year and property fixed effects. Resulting coefficients and 95% confidence intervals are graphed.

Figure 4: Event study of prices before and after solar installation operation date



Notes: The treatment variable is defined as a dummy variable equal to 1 if a house is within 0.6 miles of an eventual solar installation site. The time period variable is defined as a bin variable, starting from 6 years prior to solar installation operation date and up to 6 years post operation. Properties are sorted into the respective 6 month bin in which they transact, and the reference time period is 0.5 to 1 year prior to operation date. The coefficients are obtained by estimating a DID model similar to Equation 1 that regresses log sales price on the interaction between the treatment and the time period variables, along with month-year and property fixed effects. Resulting coefficients and 95% confidence intervals are graphed.

Figure 5: Pre-treatment price trends for DID model excluding never-treated properties



Notes: Sample size is 10,452 and includes properties within 0.6 miles of an eventual solar site built 2012 or later. Treatment is defined as an indicator variable equal to 1 if a house is proximate to a solar installation that will be built between 2012 and 2016 and equal to 0 if proximate to a solar installation built in 2017 or later. Log sales prices are regressed on the interaction between treatment and year dummy variables, along with month-year, property, and county-year fixed effects. Estimated coefficients giving differences between treatment and control properties over time and 95% confidence intervals are graphed.

Table 1: Housing attribute means by treatment status

Variable	Full sample	Pre-treatment means		Normalized difference in means
		0 - 0.6 miles	0.6 - 2 miles	
Price (000's)	314.71	305.36	316.37	-0.049
Lot size (acres)	0.42	0.47	0.41	0.103
House area (000's sq. feet)	2.92	2.86	2.92	-0.037
Bedrooms	3.09	3.04	3.09	-0.045
Full bathrooms	1.53	1.52	1.53	-0.012
Half bathrooms	0.48	0.45	0.48	-0.045
Age of home (years)	58.06	50.19	56.08	-0.121
Pool (1 = yes)	0.05	0.05	0.05	-0.005
Air conditioning (1 = yes)	0.35	0.38	0.34	0.057
Fireplace number	0.39	0.35	0.40	-0.058
Greenfield (1 = yes)	0.46	0.49	0.47	0.032
Rural (1 = yes)	0.35	0.39	0.35	0.146
Observations	107,291	7,448	64,322	

Notes: Sales prices are adjusted to 2019 levels using the CPI. Normalized difference in means calculated according to Imbens and Wooldridge (2009). Normalized differences exceeding 0.25 in absolute value are considered statistically different.

Table 2: Estimates of the impact of solar installations on property prices

Independent variables	Dependent variable: Sale price (ln)				
	Include never-treated			Exclude never-treated	
	(1)	(2)	(3)	(4)	(5)
Treated × Post	-0.024*** (0.008)	-0.018*** (0.007)	-0.015** (0.007)	-0.036*** (0.013)	-0.028** (0.013)
Controls					
Month-year fixed effects	Y	Y	Y	Y	Y
Property fixed effects	Y	Y	Y	Y	Y
Municipality time trends		Y		Y	
County-year fixed effects			Y		Y
Observations	107,291	107,291	107,291	11,292	11,292
R ²	0.871	0.876	0.878	0.889	0.891

Notes: Treated = 1 if a house is within 0.6 miles of an eventual solar installation site and Post = 1 if a house sells after the treatment date. The treatment date is defined as 6 months prior to solar installation operation date. Standard errors are clustered at the tract level and shown in parentheses. *, **, and *** indicate significance at 10%, 5%, and 1%, respectively.

Table 3: Heterogeneity of treatment effects

Independent variables	Dependent variable: Sale price (ln)	
	(1)	(2)
<i>Panel A: Heterogeneity by prior land use</i>		
Treated \times Post \times Non-greenfield	-0.011 (0.010)	-0.013 (0.015)
Treated \times Post \times Greenfield	-0.020** (0.009)	-0.044*** (0.014)
<i>Panel B: Heterogeneity by population density</i>		
Treated \times Post \times Non-rural	-0.006 (0.009)	-0.005 (0.015)
Treated \times Post \times Rural	-0.025** (0.010)	-0.058*** (0.014)
<i>Panel C: Heterogeneity by proximity</i>		
(0.6 – 1 mile) \times Post	-0.005 (0.006)	
(0.1 – 0.6 miles) \times Post	-0.016** (0.007)	-0.028** (0.013)
(0 – 0.1 miles) \times Post	-0.038 (0.052)	-0.042 (0.053)
Observations	107,291	11,292

Notes: All specifications include property, month-year, and county-year fixed effects. Treated = 1 if a house is within 0.6 miles of a solar construction and Post = 1 if a house sells after the treatment date. The treatment date is defined as 6 months prior to solar installation operation date. Greenfield = 1 if the prior land use is farm or forest land and Non-greenfield = (1-Greenfield). Rural = 1 if the population density per square mile is ≤ 850 and Non-rural = (1-Rural). In Column 1, the Panel A model also includes Post \times Greenfield and Post \times Non-greenfield, and the Panel B model includes Post \times Rural and Post \times Non-rural. In Panel C, (0.6 – 1 mile), (0.1 – 0.6 miles), and (0 – 0.1 mile) are dummy variables = 1 if properties lie within the respective distances from the nearest eventual solar installation, and distance bin for 1 – 2 miles is omitted in Column 1. Standard errors are clustered at the tract level and shown in parentheses. *, **, and *** indicate significance at 10%, 5%, and 1%, respectively.

Table 4: Robustness checks

Independent variables	Dependent variable: Sale price (ln)				
	1 MW = 4 acres	1 MW = 6 acres	Control for wind turbines	Drop pre-2009 transactions	Keep properties transacting +/- 4 years from treatment date
	(1)	(2)	(3)	(4)	(5)
Include never-treated					
<i>Panel A: Standard model</i>					
Treated × Post	-0.014* (0.007)	-0.014** (0.007)	-0.015** (0.007)	-0.025*** (0.009)	-0.021* (0.012)
<i>Panel B: Heterogeneity model</i>					
Treated × Post × Non-greenfield	-0.011 (0.010)	-0.009 (0.010)	-0.011 (0.010)	-0.014 (0.013)	-0.003 (0.016)
Treated × Post × Greenfield	-0.017* (0.009)	-0.019** (0.009)	-0.020** (0.009)	-0.037*** (0.012)	-0.043*** (0.016)
Exclude never-treated					
<i>Panel C: Standard model</i>					
Treated × Post	-0.024* (0.013)	-0.030** (0.013)	-0.027** (0.013)	-0.038** (0.016)	-0.066** (0.028)
<i>Panel D: Heterogeneity model</i>					
Treated × Post × Non-greenfield	-0.009 (0.016)	-0.015 (0.015)	-0.011 (0.015)	-0.010 (0.019)	-0.027 (0.031)
Treated × Post × Greenfield	-0.038*** (0.014)	-0.045*** (0.014)	-0.044*** (0.014)	-0.067*** (0.017)	-0.111*** (0.032)

Notes: Treated = 1 if a house is within 0.6 miles of a solar construction, and Post = 1 if a house sells after the treatment date. The treatment date is defined as 6 months prior to solar installation operation date. In Columns 1 and 2 we assume that the land area required for 1 MW of solar capacity is 4 acres and 6 acres, respectively. We control for the presence of wind turbines in Column 3 by including a dummy variable = 1 if a property lies within 1 mile of a post-construction wind turbine. Column 4 drops all transactions occurring before 2009, and Column 5 excludes all properties that transact more than 4 years before or after the treatment date of the nearest solar installation. All specifications include property, month-year, and county-year fixed effects. The number of observations for Panels A and B are: 106,552 in Column 1, 107,924 in Column 2, 107,291 in Column 3, 61,121 in Column 4, and 33,026 in Column 5. The number of observations for Panels C and D are: 10,965 in Column 1, 11,526 in Column 2, 11,292 in Column 3, 6,427 in Column 4, and 3,298 in Column 5. Standard errors are clustered at the tract level and shown in parentheses. *, **, and *** indicate significance at 10%, 5%, and 1%, respectively.