NBER WORKING PAPER SERIES

DISTRIBUTIONAL EQUITY IN THE EMPLOYMENT AND WAGE IMPACTS OF ENERGY TRANSITIONS

Ben Gilbert
Hannah Gagarin
Ben Hoen

Working Paper 31608
http://www.nber.org/papers/w31608

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2023

Any views expressed are those of the authors and not those of the U.S. Census Bureau or the National Bureau of Economic Research. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 1947 (CBDRB-FY22-P1947-R9663). This analysis was funded by the Wind Energy Technologies Office of the U.S. Department of Energy (DOE) under Contract No. DE-AC02-05CH11231. This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation. We also thank the Rocky Mountain Research Data Center staff and in particular Philip Pendergast, Kas McLean, and Carlos Becerra for their support. We are grateful for detailed and helpful comments from Johannes Schmieder and all of the participants at the NBER Conference and Pre-Conference Workshop on Distributional Impacts of New Energy Policies, which benefited from funding from the Alfred P. Sloan Foundation. We could not have performed this work without the support of the DOE, specifically Patrick Gilman, Rin Ball, and Maggie Yancey.

Distributional Equity in the Employment and Wage Impacts of Energy Transitions
Ben Gilbert, Hannah Gagarin, and Ben Hoen
NBER Working Paper No. 31608
August 2023
JEL No. Q4,Q42,Q43,R11,R12

ABSTRACT

We use restricted-access, geocoded data on the near-universe of workers in 23 U.S. states in order to quantify the impact of wind energy development on local earnings and employment, by race, ethnicity, sex, and educational attainment. We find the largest relative impacts for workers without a high school education, or workers with a college education, in addition to other systematic differences across sub-populations. We compare these results to estimates using county aggregates of the worker-level data, such as can be obtained using publicly available data. We find that (a) county-level estimates are dramatically dampened relative to geocoded worker-level estimates, and (b) the degree of bias differs by sub-population such that qualitative comparisons of impacts are not consistent using restricted-access data versus county-level data for most sub-populations. We discuss implications for achieving equity goals within energy transition policies.

Ben Gilbert
Colorado School of Mines
1500 Illinois St
Golden, CO 80401
bgilbert@mines.edu

Hannah Gagarin
Sandia National Labs
hrgagar@sandia.gov

Ben Hoen
Electricity Markets and Policy Group
Lawrence Berkeley National Laboratory
bhoen@lbl.gov

1 Introduction

Considering the current Biden administration's climate goals, which include a net-zero carbon pollution-free power sector by 2030 and a net-zero carbon pollution-free economy by 2050, the US is poised to undergo an energy transition away from fossil fuels and towards renewable energy technologies. This has the potential to change the landscape of local labor markets as large-scale renewable energy technologies are deployed in communities across the US. Renewable energy development may demand different skills, and occur in different places, than legacy energy sources. Renewable energy can also change the composition of the local economy through rent and royalty payments to landowners, tax payments to local jurisdictions, or local demand from manufacturing facilities for renewable energy components (Kim, 2019). Some scholars and policy makers raise concerns, however, that these investments still occur within the same political and cultural contexts that historically create an uneven distribution of economic welfare, along racial, gender, educational, or other socioeconomic dimensions, which has the potential to continue marginalizing vulnerable or underrepresented populations (Mueller and Brooks, 2020; Mejía-Montero et al., 2021). The concept of incorporating a justice element into existing and future US energy policy is often referred to as a "just transition" and lies at the center of a branch of social science research called environmental and energy justice (EEJ). EEJ provides a conceptual, analytical, and decision-making framework for understanding how to pursue a sustainable energy system that redistributes welfare to avoid undue burdens on marginalized communities (McCauley et al., 2019; Carley and Konisky, 2020).

In this paper, we investigate how large are the local employment and earnings gains for workers from wind energy development, and to whom do these gains accrue in terms of different worker sub-populations by race, ethnicity, sex, and educational attainment. We estimate impacts separately for black, American Indian/Native Alaskan, and white workers; for Hispanic workers; for male and female workers; and for workers in four educational attainment brackets. As a secondary research question, we also ask: is their exclusivity in the research community regarding the ability to correctly answer these questions? Specifically, we quantify how well can we measure these impacts using county-level data, such as is publicly available from the American Community Survey, the Bureau of Labor Statistics, etc., as compared to estimates using individual, geo-coded, worker-level data such as is only available in restricted-access settings for researchers with the institutional capacity or connections required to obtain such data.

In order to answer these questions, we combine geocoded data on the near-universe of workers in 23 U.S. states from 2000 to 2020 with geocoded data on all U.S. wind projects. We estimate the causal effect of the arrival of utility-scale wind capacity within 20 miles of a worker's residence on that worker's earnings and employment status, controlling for wind capacity in place at greater distances. Specifically, we use restricted-access worker-level data provided by the US Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) dataset for 23 participating states, which contains worker residence coordinates, employment status, and earnings, as well as age, sex, race, ethnicity, and educational attainment for all workers who paid into, or received benefits from, their state's unemployment insurance program – more than 96 percent of workers. We combine this with the U.S. Wind Turbine Database, which contains the coordinates, arrival year, and capacity rating of wind plants.

In order to deal with the massive computational needs for a dataset of this size, we employ two separate empirical approaches. We first implement a unique shift-share IV design within a spatial lag model using a subset of the data. This approach estimates the effect of wind capacity in increasing 20-mile donut-distances from each residence on that worker's outcomes. We instrument for capacities in each donut using predicted capacities that come from a shift-share design. In this design, we predict wind capacity in localized hexagonal grid cells using local average wind speeds interacted with national aggregate wind capacity expansion trends and global prices of metals and energy commodities. We then aggregate these localized predictions to donut-distances around each worker's residence to construct the IVs for actual capacity in each donut. Results from this exercise indicate that the majority of the employment and earnings impacts are captured by wind development within the first 20-mile ring, with much smaller coefficients and larger standard errors at greater distances. However, this approach is very computationally intensive on a large geocoded dataset.

In order to reduce computation time and draw inference from the full dataset, we then use a brand new difference-in-differences estimator: Local Projections Difference-in-Differences (LPDID) (Dube et al., 2023). The LPDID estimator is a regression-based approach that avoids the biases associated with Two-Way Fixed Effects in the presence of staggered treatment and heterogeneous, dynamic treatment effects and provides computational advantages over recent alternative estimators (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Cengiz et al., 2019) while allowing for continuous treatment variables. Having established in the initial shift-share approach that most impacts occur within 20 miles, we use the 20-mile capacity as our "treatment" variable in the diff-in-diff approach, and control for capacity at greater distances as additional regressors. Despite the computational advantages of LPDID, estimation on the full dataset is still not feasible. Because we have the near-population of workers, however, we repeatedly randomly sample 1 million unique workers (with replacement) and re-estimate our models 100 times on each random sample in order to report the "true" variation in coefficient estimates rather than using analytical standard errors calculated from a single sample.

We estimate impacts for all workers, and then separately for black, American Indian/Native Alaskan, and white workers, Hispanic/Latino workers, male and female workers, and workers with either no high school completion, a high school diploma, some college or Associates degree, and a bachelor's degree or higher. With causal estimates of the effect of wind capacity at the geocoded worker level in hand, we then aggregate the worker-level data to county-level averages and re-estimate analogous models on the aggregate data. In future versions of this paper, we will also estimate the models on publicly available county-level American Community Survey data from the same states from 2010 to 2020 for comparison. [1]

We find that wind power development within 20 miles of a given worker causes a statistically and economically significant increase in earnings and employment. These impacts vary meaningfully across sub-populations. Black workers experience the largest proportional marginal impact despite there being few black workers within 20 miles of a wind project. Men enjoy larger gains than women. As another example, workers without high school completion have the largest proportional impacts, followed by workers with a college degree.

We further find that estimates using county-level aggregates are dramatically attenuated compared to the analogous estimates from worker-level data, especially for earnings. While earnings impacts are statistically and economically significant for all sub-populations in the restricted-access data, our estimates are effectively zero using the same data aggregated to the county level. The degree of attenuation is also not uniform across sub-populations, nor does it vary in predictable ways. For example, the employment impact for workers with some college coursework are 16 percent as large using county-level vs. worker-level data, whereas the impacts for workers with a bachelor's degree are 60 percent as large, and the impact estimates for Hispanic/Latino and American Indian/Native Alaskan workers are actually larger using county-level data.[2] This suggests that

[1]Aggregating to the county level, versus aggregating to the county plus any other county with a centroid within 20 miles produces very similar effect estimates, so we report county-level estimates here as a comparison to the worker-level plus 20 miles results.

[2]This is in line with recent work by Colmer, Voorheis, and Williams, 2021, who find that county-level estimates

researchers from lower-resourced institutions, without access to such data resources, face additional barriers to accurately measuring the employment impacts of new energy investments that may be important to the populations that their institutions serve.

Our findings contribute to a growing literature on the equity implications of large scale energy transitions. The impacts that we estimate are far larger than can be explained by maintenance workers at the wind plant alone. Our findings confirm other recent studies showing larger local impacts (Gilbert, Gagarin, and Ben Hoen, 2023; E. J. Brunner and Schwegman, 2022) than earlier studies using data from the 2000's (Brown et al., 2012a). However, while our employment estimates are in line with these recent studies, our earnings estimates are much larger and our evidence suggests that using county-level data understates earnings impacts by a far wider margin than employment impacts. The magnitudes of these effects suggest that the majority of local economic impacts occur through indirect channels. First, wind projects pay royalties to local landowners that can either be fixed annual payments or shares of the annual wind revenue. To the extent that local landowners spend money in the community, this generates a local economic impact multiplier. Second, wind installations contribute to the local tax base, raising revenue for school districts (E. Brunner, Ben Hoen, and Hyman, 2022), counties (Castleberry and Scott Greene, 2017) and other local jurisdictions. Expenditures by these jurisdictions on local services also generate economic impact multipliers (Shoeib, Infield, and Henry C. Renski, 2021). Third, the infusion of capital expenditures during the construction phase of a wind project may stimulate the opening of new establishments that persist when the construction phase ends, including the opening of manufacturing facilities for energy components (Kim, 2019). Although our coefficients are larger than previous wind energy literature, however, when we convert these coefficients to local economic multipliers they are still modest in size compared to central estimates in the literature on earnings and employment multipliers for fiscal stimulus or industry investment from sectors other than wind energy. Back-of-the-envelope calculations suggest an employment multiplier of 0.46 jobs per million dollars of wind capacity investment, whereas employment multipliers in most other sectors are more than an order of magnitude larger. Our calculated earnings multiplier is approximately 0.14 dollars of local worker earnings per dollar of wind capacity investment, which is about 1/5 or 1/6 as large as central estimates of fiscal stimulus multipliers. Because direct employment at operating wind projects is very small, these are likely to be almost entirely indirect multiplier effects.

We also contribute to the burgeoning field of EEJ. EEJ examines how systemic injustices embedded in societal and cultural norms can still be present in energy transitions and addresses how to approach solving this from a policy perspective, to ensure a more "just" energy transition. However, there is currently a knowledge gap in the EEJ literature in terms of being able to carefully measure impacts of energy development on disadvantaged sub-populations at a granular level. This paper attempts to alleviate this gap by comparing results from geocoded worker-level data on disadvantaged sub-populations to those from county-level aggregates in order to establish the quality with which we have previously measured these impacts.

The remainder of this paper is organized as follows. Section 2 reviews the literature on regional economic impacts from energy in relation to the EEJ literature. Section 4 discusses the empirical methods, and section 3 describes the data used in the analyses. Section 5 presents and discusses the results while section 6 concludes.

---

of air pollution impacts on health may obscure some important within-county heterogeneity in responses.

2 Literature Review

EEJ recognizes that energy systems are deeply entrenched in social, economic, and political life which can lead to large-scale impacts on communities (Miller, Richter, and O'Leary, 2015). Broadly, EEJ focuses on the ways in which underrepresented communities are harmed in the context of energy systems. It provides a framework for policymakers and stakeholders to ensure a "just transition" for all members of the community. There is a breadth of EEJ literature, spanning historical and potential energy developments, so that this literature review is by no means exhaustive; instead, the purpose is to outline literature relevant to this current paper, as well as highlight potential literature gaps.

2.1 EEJ In a US Context

Within the context of a US transition to a low-carbon energy sector, renewable energy can often be portrayed as a more socially-just form of energy due to decarbonization (Mejía-Montero et al., 2021); however, Mueller and Brooks, 2020 point out that there can still be an uneven distribution along social dimensions such as race/ethnicity and other marginalized populations. Mainly, while an energy system might transition, the cultural and societal practices and norms that existed in the "old" energy system (in this case, a fossil-based energy system) still remain ingrained, maintaining the continued marginalization and under-representation of communities in the "new" energy system. Some EEJ scholars and stakeholders argue that policymakers should approach the roll out of green technologies with the intent of ensuring that these new systems better recognize and affirm said communities, and better reflect their needs in energy policy, to produce a more just energy system. Understanding how renewable energy investments have altered local labor markets, and in particular, the position of disadvantaged sub-populations within local labor markets, is therefore important for evaluating the EEJ dimensions of the energy transition.

There are numerous studies that examine the impact of renewable energies, and their impacts on marginalized communities in the US. For example, Borenstein and Davis, 2016 show that the distribution and use of renewable energies tends to skew towards white, higher-income households in the US. And, higher-income households tend to receive the bulk of tax credits for home weatherization, solar panel installation, hybrid and EV cars and other "clean energy" investments, so that this policy mechanism fails to bring a clean energy technology to those that need it most. Sunter, Castellanos, and Kammen, 2019 examine the deployment of solar energy, and their results show that rooftop solar PV is less likely to be adopted by Black- and Hispanic-majority census tracts, likely due to racial and ethnic differences in household income and home ownership. This work is important in demonstrating the potential "climate gap" that can occur in the introduction of clean energy technologies, where only one community benefits.

2.2 Wind Development and EEJ

There is a wealth of literature regarding energy development and its impact on economic outcomes in local communities. Some focus on more historical energy sources, such as coal and the feature of boom and bust cycles (Black, McKinnish, and Sanders, 2005). More recent work has examined the shale oil and gas boom on the US and economic impacts (Tsvetkova and Partridge, 2016; Weinstein, Partridge, and Tsvetkova, 2018; Maniloff and Mastromonaco, 2017). In thinking about wind development impacts on economic outcomes alone, both Varela-Vázquez and Sánchez-Carreira, 2015 and Brown et al., 2012b provide estimates regarding wind projects and impacts to communities, and show that there are positive economic outcomes associated with a wind project development. However, these works examine energy development impacts without a justice element.

In terms of wind energy development and the EEJ literature, these studies tend to focus on what is referred to as a "social gap": where there is widespread public support, but localized opposition (Acheampong, Erdiaw-Kwasie, and Abunyewah, 2021; Fergen and Jacquet, 2016). Thus, the EEJ literature emphasizes that local acceptance of wind energy as an important factor to consider (Guo et al., 2015; Kontogianni et al., 2014; Jones and Eiser, 2010) in the overall success of developing a wind project. Indeed, public perception of wind projects tends to be a theme in EEJ-related wind development literature (Shoeib, Henry C Renski, and Hamin Infield, 2022; Barry, Ellis, and Robinson, 2008; Zárate-Toledo, Patiño, and Fraga, 2019; Mills, Bessette, and H. Smith, 2019). But there are some works that focus not on the overall acceptance and perception of a wind project, but on the outcomes of underrepresented communities located near wind projects. Acheampong, Dzator, and Shahbaz, 2021 posit that wind development is likely sited disproportionately in areas where residents have lower social and financial capital, arguing that this places the cost and burden of wind energy on those that are already historically marginalized. Indeed, half of wind projects are located in non-metro, predominantly rural counties Shoeib, Henry C Renski, and Hamin Infield, 2022, which is important since rural areas in the US tend to suffer from a higher rate of poverty compared to urban areas Farrigan, 2021. Similarly, Mueller and Brooks, 2020 argue that siting wind projects in locations with lower labor force participation is an example of energy injustice. To the extent that wind development generates local employment, income, tax revenue, and community services, however, the benefits may exceed the costs.

Shoeib, Henry C Renski, and Hamin Infield, 2022 argues that larger, more urban counties may be better able to internalize the potential economic benefits from an energy development, in that they are able to provide specialized services such as banking, legal, and secondary inputs such as skilled workers and raw inputs. Shoeib, Henry C Renski, and Infield, 2022 similarly find that wind development only increases farm income in rural areas, whereas earnings, employment, and poverty improve on average when studying all counties. Mauritzen, 2020 supports this idea of uneven impacts between rural and urban areas, showing that rural areas characterized by low-incomes are least likely to capitalize on positive economic benefits of wind power development, compared to more urban, metro areas. Indeed, Pedden, 2006 also demonstrates this, with their results showing that smaller communities see more leakage into nearby areas that are better able to provide more services. Mauritzen, 2020 also find, however, that wages rise in rural areas, arguing that the likely mechanism is from landowner lease payments and local tax revenue.

A puzzle for studies such as ours that find sizeable net employment and earnings impacts of wind, is by what mechanism does the wind installation generate these local benefits? Ben Hoen et al., 2015 finds that home prices near wind projects are not negatively impacted, while E. J. Brunner and Schwegman, 2022 finds that wind development increases nearby home values. E. J. Brunner and Schwegman, 2022 also estimate sizeable positive impacts from wind using county-level data, including increases in GDP per capita, income per capita, and median household income, with employment shifts into construction and manufacturing jobs and out of farm employment. Castleberry and Scott Greene, 2017 and Shoeib, Henry C Renski, and Hamin Infield, 2022 find that wind projects increase local tax revenue and availability of community services, which can generate and sustain employment. E. Brunner, Ben Hoen, and Hyman, 2022 find that wind development causes sustained increases in revenue and expenditures per student in local school districts.

2.3 Restricted-Access LEHD Data

Many of the local labor market and community well-being questions raised above can be better answered with more granular data on workers and communities. The U.S. Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) database infrastructure provides the inputs required to produce the quarterly workforce indicators (QWI) (Abowd et al., 2013). Researchers with either paid access or institutional access to a Federal Statistical Research Data Center can propose research questions using the underlying individual worker-level and establishment-level data. For example, McKinney and Abowd, 2020 use the LEHD to study changes in male earnings before and after the Great Recession, and Card, Rothstein, and Yi, 2021 use the LEHD to study the impact of location on worker earnings. The LEHD is released in multi-year "snapshots." This paper uses the 2014 snapshot, with data from 2000 through 2014, as well as the recently-released 2021 snapshot with data from 2000 through 2020.

3 Data sources and characteristics

We combined several geocoded datasets in order to estimate the models described in section 4. Wind capacity data came from the U.S. Wind Turbine Database, which contains latitude and longitude coordinates, year of operation, nameplate capacity, and other technical specifications for all wind turbines in the United States over the study period (BD Hoen et al., 2021; Rand et al., 2020). Project-specific ID numbers used to group turbines into plants were provided by Lawrence Berkeley National Lab.

We worked with the U.S. Census Bureau's Federal Statistical Research Data Center (FSRDC) program to access the LEHD infrastructure files from within a secure location. The LEHD is built on quarterly state-level unemployment insurance rolls, which cover more than 96 percent of workers who reside in the U.S. We aggregate to the annual frequency for analysis. Because the LEHD is constructed from state-level unemployment insurance programs, state-by-state approval to access the data is required. The 23 states shown in blue in Figure 1 agreed to participate in our approved project. As this figure indicates, although we do not have data for Texas or Minnesota, we do have access to states with a significant share of wind capacity. For this paper, we used the LEHD's Individual Characteristics File which indicates race, ethnicity, sex, and educational attainment of individual workers. The data do not indicate whether a worker identifies with multiple races or ethnicities, or is gender non-binary. This is a limitation of our study. We also used the Employment History Files, which includes a worker's quarterly earnings from all jobs held in that quarter. We aggregated earnings to the annual level. We calculated our employment variable as the number of quarters in each year in which the worker had positive earnings, divided by 4. We also used the ICF LEHD Residence Files, which contain latitude and longitude coordinates for the worker's residence in each year.[3]

For this paper, we report shift-share IV results using the 2014 LEHD Snapshot with data from 2000 to 2014, and diff-in-diff results using the recently released 2021 LEHD Snapshot, which includes employment data through 2021, but residence data through only 2020. Our preferred results are the diff-in-diff estimates from the 2021 Snapshot, but we use our shift-share IV results from the earlier snapshot in order to guide our estimation approach as we discuss in Section 4. Because of migration that occurred during the COVID-19 pandemic, we were not confident in imputing residence locations for 2021. We therefore limit our analysis to the years 2000 through 2020.

In each of these datasets we use georeferenced locations (i.e. latitude and longitude) of both worker residences or centroids of wind projects in order find the aggregate wind capacity at various distance bands from each worker residence. This produced an annual worker-level panel with

[3]To the extent that workers live in one place and periodically travel to distant work sites with energy installations, our approach would not pick these people up, although our sense is that this is more common in the oil and gas industry than in the wind industry. One could possibly look for workers in the LEHD whose establishments of employment are farther from their residence than a reasonable commuting distance, but this is not straightfoward and is beyond the scope of this study.

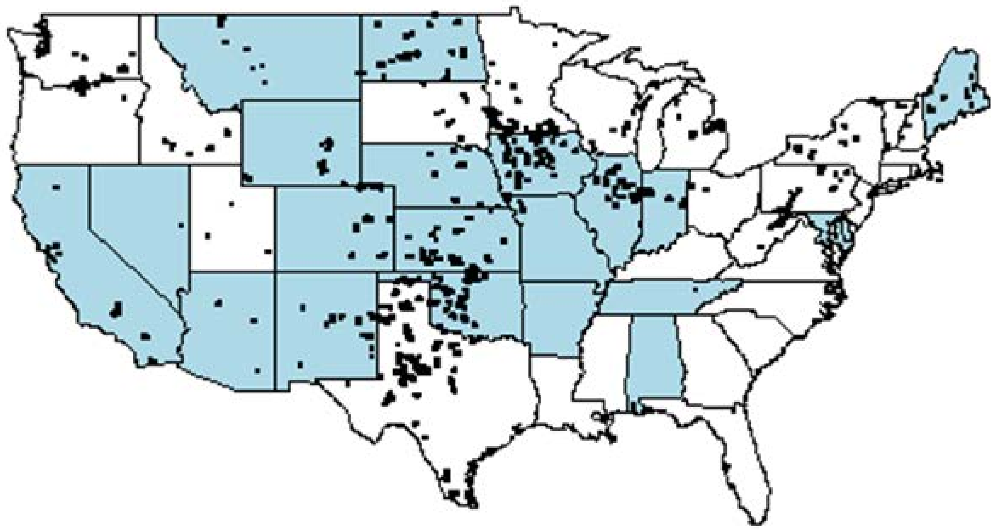

Figure 1: US States With Wind Projects

Notes: States highlighted in blue are the states examined in this paper; the black dots represent US wind project installations. Source: Author, BD Hoen et al., 2021

employment and earnings outcomes, demographic characteristics, and wind capacity exposure at various distances. Summary statistics for this combined panel are given in Table 1.

4 Methods

Worker-level residential location choice and labor supply decisions may be endogenous with respect to determinants of local wind energy development because of location preferences, local economic shocks, or state and local policies affecting local labor markets. We therefore need a careful identification strategy to recover the causal effect of wind development on local labor outcomes.

In addition to these potential endogeneity concerns, our empirical approach needs to overcome two more important challenges. First, is the computational limits on performing many geospatial calculations of individual worker distances to wind projects and estimating multiple regression models with over a billion observations. Second, it is not obvious who belongs in the "treatment" versus "control" groups given previously documented spatial spillovers of energy development on labor markets (Feyrer, Mansur, and Sacerdote, 2017; James and B. Smith, 2020).

It is not straightforward to classify worker residences as "wind" vs. "non-wind" households as other individual-level studies have done (Jacobsen, Parker, and Winikoff, 2023) because of this second concern. As Figure 1 indicates, wind projects are spatially clustered. Many workers may live within a short distance of more than one wind project, so assigning them as "treated" by their closest wind project will understate exposure. This is one reason we aggregate capacity in distances around worker residences, despite the computational intensity of doing so.[4] One option to determine

[4] Another benefit of aggregating capacity around residences, using the residence as the unit of analysis, rather than using the wind project as the unit of analysis and including all nearby households in the "treatment group" is that we can include capacity that may be just across state borders in states for which we do not have worker data, but which may still impact workers that we do observe. For example, plants in southern Minnesota may affect workers

Table 1: Summary Statistics

	Worker-Level				County-Level			
	Treated		Untreated		Treated		Untreated	
	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev
All Workers	$N \cdot T = 1430000000$				$N \cdot T = 29750$			
Earnings per Worker	31100	110600	32540	186400	33670	8545	33820	10290
Employment	0.68	0.43	0.66	0.43	0.68	0.06	0.64	0.10
Post Treatment GW	0.31	0.32			0.2	0.24		
N (Workers or Counties)	9399000		103000000		240		1180	
Black	$N \cdot T = 165300000$							
Earnings per Worker	22660	150400	22010	83090	24750	10190	25340	9825
Employment	0.62	0.44	0.63	0.44	0.53	0.15	0.55	0.16
Post Treatment GW	0.40	0.38						
N (Workers)	677000		12700000					
American Indian	$N \cdot T = 23100000$							
Earnings per Worker	18340	31460	18990	44110	25620	8630	26750	9736
Employment	0.58	0.44	0.59	0.45	0.59	0.12	0.56	0.14
Post Treatment GW	0.30	0.31						
N (Workers)	172000		1604000					
White	$N \cdot T = 1110000000$							
Earnings per Worker	31030	109100	33830	201700	34170	8812	35070	11100
Employment	0.69	0.43	0.67	0.43	0.68	0.06	0.65	0.10
Post Treatment GW	0.30	0.31						
N (Workers)	7659000		78400000					
Hispanic	$N \cdot T = 237200000$							
Earnings per Worker	24780	50000	23730	88440	26770	7085	26850	7957
Employment	0.65	0.43	0.66	0.43	0.62	0.09	0.59	0.12
Post Treatment GW	0.38	0.39						
N (Workers)	1306000		17900000					
Female	$N \cdot T = 706100000$							
Earnings per Worker	24140	42270	25610	92130	26430	6673	26990	7959
Employment	0.67	0.43	0.66	0.44	0.69	0.06	0.65	0.10
Post Treatment GW	0.31	0.32						
N (Workers)	4554000		50100000					

Notes: This table reports summary statistics from the LEHD worker data from 2000 to 2020 in the 23 states to which we have access. At the worker level, $N \cdot T$ denotes the total number of observations in each group in the near population, from which summary statistics are calculated. N indicates the number of workers (or counties) in each group. At the worker level, our Monte Carlo regression estimates randomly sample from N and take all years available for the selected workers. At the Worker level, "Treated" indicates workers who have at least 10 MW of wind capacity within 20 miles of their residence for at least one year during the study period. At the County level, "Treated" indicates counties that have at least 10 MW of wind capacity in the county for at least one year during the study period. Earnings per Worker denotes the average annual earnings for individuals who had non-zero earnings for at least two quarters of the year. Employment denotes the number of quarters per year in which a worker reported non-zero earnings, divided by 4. Post Treatment GW denotes the average exposure to wind capacity within 20 miles of an individual worker, after at least 10 MW has been installed. The County-Level statistics are calculated from the same underlying data, aggregated to the county level.

Table 1: Continued, Summary Statistics

	Worker-Level				County-Level			
	Treated		Untreated		Treated		Untreated	
	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev
Male	$N \cdot T = 721900000$							
Earnings per Worker	37800	148900	39340	245700	40690	10540	40440	12720
Employment	0.68	0.42	0.67	0.43	0.67	0.07	0.64	0.11
Post Treatment GW	0.31	0.32						
N (Workers)	4845000		53000000					
No High School	$N \cdot T = 218300000$							
Earnings per Worker	19000	33610	18460	84260	25350	7503	24690	7245
Employment	0.6	0.44	0.6	0.45	0.6	0.06	0.58	0.10
Post Treatment GW	0.33	0.35						
N (Workers)	1343000		16600000					
High School	$N \cdot T = 378300000$							
Earnings per Worker	24000	79300	23690	75920	30220	7590	29770	7686
Employment	0.67	0.43	0.65	0.44	0.67	0.06	0.64	0.11
Post Treatment GW	0.30	0.31						
N (Workers)	2605000		26900000					
Some College	$N \cdot T = 447100000$							
Earnings per Worker	29280	62590	29630	96840	33110	8054	32960	8687
Employment	0.69	0.42	0.67	0.43	0.69	0.06	0.66	0.11
Post Treatment GW	0.31	0.32						
N (Workers)	3057000		31500000					
College	$N \cdot T = 384300000$							
Earnings per Worker	48230	188900	52490	327200	43220	10990	44490	14560
Employment	0.72	0.41	0.7	0.42	0.71	0.06	0.67	0.11
Post Treatment GW	0.31	0.31						
N (Workers)	2393000		28100000					

Notes: This table reports summary statistics from the LEHD worker data from 2000 to 2020 in the 23 states to which we have access. At the worker level, $N \cdot T$ denotes the total number of observations in each group in the near population, from which summary statistics are calculated. N indicates the number of workers (or counties) in each group. At the worker level, our Monte Carlo regression estimates randomly sample from N and take all years available for the selected workers. At the Worker level, "Treated" indicates workers who have at least 10 MW of wind capacity within 20 miles of their residence for at least one year during the study period. At the County level, "Treated" indicates counties that have at least 10 MW of wind capacity in the county for at least one year during the study period. Earnings per Worker denotes the average annual earnings for individuals who had non-zero earnings for at least two quarters of the year. Employment denotes the number of quarters per year in which a worker reported non-zero earnings, divided by 4. Post Treatment GW denotes the average exposure to wind capacity within 20 miles of an individual worker, after at least 10 MW has been installed. The County-Level statistics are calculated from the same underlying data, aggregated to the county level.

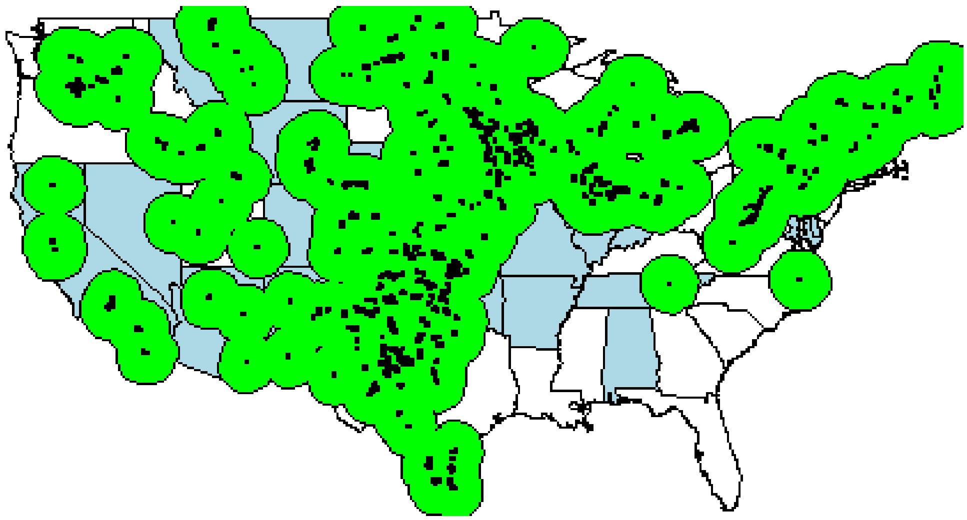

Figure 2: 100-mile buffers around each wind project

Notes: States highlighted in blue are the states examined in this paper; the black dots represent US wind project locations, and green circles represent 100 mile radii around each wind project. Source: Author, BD Hoen et al., 2021

treatment status might be to select a certain distance cutoff beyond which we assume exposure is zero. In the context of oil and gas, previous studies have found economically and statistically significant impacts as far as 60 to 100 miles away (Feyrer, Mansur, and Sacerdote, 2017; James and B. Smith, 2020). As Figure 2 illustrates, a 100-mile buffer around all wind projects leaves very few areas from which to select a control group. In order to deal with both the endogeneity issues and the spillover concerns with treatment assignment, we first estimate a spatial lag model a la James and B. Smith, 2020, but with a shift-share IV at each spatial lag that is akin to Feyrer, Mansur, and Sacerdote, 2017. This approach allows us to establish the extent of spillover as a guide to determining treatment status. However, the approach is incredibly computationally intensive and we are only able to implement it on a small random subset of the full dataset.

As we will show, this exercise suggests that most impacts are captured within the first 20 miles. Impacts at greater distances are much smaller, but not zero. We therefore proceed with a more computationally feasible estimation approach that uses capacity within 20 miles as the main treatment variable, while still controlling for capacity at greater distances, and that leverages the entire dataset.

The solution that we implement, and the source of our preferred results, is the Local Projections Difference-in-Differences (LPDID) framework recently proposed by Dube et al., 2023, described in detail below. In order to deal with remaining computational constraints and capture the sampling uncertainty in our estimates, we re-estimate our LPDID models 100 times using random draws of worker IDs from the near-population in our dataset. Finally, we aggregate the full dataset, as well as relevant sub-populations, to the county level and re-estimate the LPDID coefficients for comparison.

from Northern Iowa who are in our sample.

4.1 Shift-share IV

We estimate our spatial lag regressions with shift-share instrumental variables on a 0.1 percent random sample of the 2014 LEHD Snapshot dataset. The "share" component of the IV, or the exogenous cross-sectional variation, comes from average local wind speeds, whereas the "shift" component, or time series shock, comes from national aggregate trends in expansion of wind capacity, number of turbines, and number of wind plants, as well as commodity prices of crude oil, natural gas, aluminum, and a rare earth metals composite index. In a shift-share estimator, identification can come from either exogeneity of the "shares" or from many exogenous "shocks" or "shifts". In our case, we argue that identification comes from both sources. Exogenous spatial variation in wind resource has been argued in E. Brunner, Ben Hoen, and Hyman, 2022 and Brown et al., 2012a as well as for other resources like oil and gas deposits in Maniloff and Mastromonaco, 2017 and Feyrer, Mansur, and Sacerdote, 2017. We augment this with exogenous time series shocks in the form of national or global trends in commodity prices that might affect local wind development, and national wind development trends that capture the national policy and investment climate in a given year.

Specifically, we estimate the following spatial lag model:

$$Y_{ict} = \sum_{d \in \{20,40,60,80,100\}} \gamma_d D_{ict}^d + \alpha_{ic} + \mu_{st} + \epsilon_{ict} \tag{1}$$

where D_{ict}^d is the amount of wind capacity in donut distance d from person i's residence in year t, with d ranging from 0 to 20 miles, 20 to 40, etc. through 80 to 100 miles.

We instrument for each D_{ict}^d using predicted capacity in each donut/ring around each person's residence. We construct these predicted capacity instruments using the following approach. We segment the lower 48 U.S. states into approximately 216,000 evenly sized hexagons, and gather data on the average wind speed within each hexagon. These hexagons are about the size of the average Census block group. Using the geocoded U.S. Wind Turbine Database, we then calculate the total number and capacity of turbines in each hexagon in each year. The vast majority of hexagons have zero wind capacity, so we use a panel quasi-Poisson regression to predict total wind capacity and number of turbines in each hexagon in each year. The right hand side of the panel Poisson regression includes hexagon fixed effects, state-by-year fixed effects, and a cubic polynomial in average wind speed interacted with the total national U.S. wind capacity in each year, the total number of wind turbines in the U.S. in that year, the number of U.S. wind plants in that year, and the vector of commodity prices. We use a cubic function of wind speed in order to capture the fact that wind development is not feasible at either very low or very high average speeds. The variation in predicted hexagon-year wind development is then driven by unobserved hexagon-specific effects such as suitability for transmission access, secular state-by-year macroeconomic or policy trends such as state-level renewable energy incentives or mandates, and a nonlinear shift-share component driven by the interaction of wind speed with national aggregate investment behavior, essentially allocating nationwide investment in each year to the locations that are most suitable for development according to wind speed.

A problem with the panel Poisson approach with hexagon fixed effects is that the predicted value in any hexagon that never has wind capacity - the vast majority of them - is zero. As such, the predicted values are extremely highly correlated with the actual values. Aggregating these predicted values around worker residences in order to construct instruments for equation (1) thus produces instruments that essentially reproduce the OLS estimates.

In order to deal with this issue, we further aggregate both actual wind capacity from the U.S. Wind Turbine Database and predicted wind capacity from the Poisson regression on each hexagon, to the county-year level. We also calculate average wind speed in each county. We then predict capacity in each county-year using a linear OLS-FE regression with county fixed effects, state-by-year fixed effects, and interactions between the county-average wind speeds and the sum of hexagon-level predictions within each county. This procedure generates county-year variation in predicted wind capacity. Finally, we use county centroids to aggregate the county-year predicted capacity values in donut distances around each worker's residence. This produces a nonlinear transformation of unique fixed county and hexagon characteristics, and state-level macroeconomic trends, and shift-share variation from local wind speed and national aggregate wind development trends. These county-level predictions aggregated to 20-mile donut-distances from worker residences to surrounding county centroids become our IVs for equation (1) in a standard just-identified linear IV framework.[5]

We perform this shift-share estimation approach for equation (1) on both a 0.1 percent random sample and on county-aggregated data for comparison.

4.2 Shift-share IV Results

Here we report results from the estimation of equation 1 using a 0.1 percent sample of the worker data from 2000 through 2014 (i.e., using the 2014 snapshot of the LEHD). Standard errors are clustered at the worker level in worker-level regressions, and at the county level in county-level regressions.

In both Tables 2 and 3 we can see that at the worker level (Panels A) the magnitude of the point estimate for capacity declines steeply at distances greater than 20 miles. This provides confidence that results for the 20 mile ring reported are capturing most of the relevant impacts. By comparing the first column of Tables 2 and 3 we also see that within 20 miles, the difference between the county-level estimate and the worker-level estimate are much greater for log earnings than for employment. We will see a similar general pattern when we compare individual and county-level estimates using the LPDID estimator on the data through 2020.

Although these are not our preferred estimates, the magnitudes are also comparable to the LPDID results we will see in Section 5, despite the different sample and estimation method. The first column in Panel A of Table 2 indicates that an additional 100 MW of wind capacity within 20 miles causes employment of the average worker to increase by 0.26 percentage points. The comparable effect using LPDID on the 2021 LEHD Snapshot is 0.12 percentage points from a 100 MW increase, taken from the first column of Panel A in Table 6 in the Appendix. Similarly, the first column of Panel A of Table 3 shows that an additional 100 MW of nearby wind capacity increases earnings by 2 percent. The comparable estimate of the continuous treatment effect using LPDID on the more recent data is in the first column of Panel A of Table 7 in the Appendix, which corresponds to a 1.1 percent increase from 100 MW of wind.

4.3 Local Projections Diff-in-Diff

It is not computationally feasible to estimate the shift-share IV approach described in Section 4.1 on the full 2021 LEHD Snapshot, or even to downsample and perform the geospatial calculations repeatedly on large random subsamples. We therefore turn to the LPDID method for our preferred results.

Under LPDID, the estimation of a diff-in-diff event study is analogous to the estimation of an impulse response function by local projections in a time series analysis. One estimates a sequence

[5] It may be possible to estimate this in a nonlinear approach in one step, but computational limitations make the linear IV framework more feasible to implement.

Table 2: Spatial Lag Coefficients (GW w/in each 20 mile donut): Employment

	0 to 20 m.	20 to 40 m.	40 to 60 m.	60 to 80 m.	80 to 100 m.
			Panel A: Worker-Level		
Installed	2.6	-0.66	-0.33	-0.14	-0.05
Capacity (GW)	(2.1)	(0.37))	(0.15)	(0.086)	(0.052)
R^2	0.20				
K-P Wald F-stat	265				
$N \cdot T$	1438000				
			Panel B: County-Level		
Installed	3.8	1.8	2.1	1.4	1.0
Capacity (GW)	(3.3)	(1.7))	(1.0)	(0.71)	(0.53)
R^2	0.22				
K-P Wald F-stat	2110				
$N \cdot T$	29750				

Notes: This table reports the spatial lag coefficients from equation (1) using the shift-share IV procedure described in section 4.1. The dependent variable is the percentage of the year employed (quarters with nonzero earnings divided by 4). The treatment variables are GW of capacity within each 20-mile donut distance of a worker's residence. We used a 0.1 percent random sample of workers from the 2014 snapshot of the LEHD. Standard errors are clustered at the worker level in Panel A, and at the county level in Panel B.

Table 3: Spatial Lag Coefficients (GW w/in each 20 mile donut): Log Earnings

	0 to 20 m.	20 to 40 m.	40 to 60 m.	60 to 80 m.	80 to 100 m.
			Panel A: Worker-Level		
Installed	20	1.7	0.12	-0.27	-0.075
Capacity (GW)	(12)	(2.4))	(0.96)	(0.57)	(0.34)
R^2	0.18				
K-P Wald F-stat	265				
$N \cdot T$	1438000				
			Panel B: County-Level		
Installed	4.1	2.2	1.2	0.75	0.51
Capacity (GW)	(2.3)	(0.84))	(0.45)	(0.31)	(0.22)
R^2	0.90				
K-P Wald F-stat	2110				
$N \cdot T$	29750				

Notes: This table reports the spatial lag coefficients from equation (1) using the shift-share IV procedure described in section 4.1. The dependent variable is the log of earnings, with the sample limited to employed workers (those with at least two quarters of non-zero earnings in a year). The treatment variables are GW of capacity within each 20-mile donut distance of a worker's residence. We used a 0.1 percent random sample of workers from the 2014 snapshot of the LEHD. Standard errors are clustered at the worker level in Panel A, and at the county level in Panel B.

of long-difference regressions with increasingly long differences in the outcome from the treatment date. Treatment coefficients from each regression in the sequence make up each element of the event study relative to the treatment date. In each regression in the sequence, one can obtain "clean controls" (control units that have not previously been treated) by limiting the sample to those units whose treatment status changed in the current period, and those that have never been treated during the time horizon of the given long difference.

For this approach, we assume that workers living within 20 miles of utility-scale wind capacity are "treated", and all other workers are not. We include capacity exposure at greater distances as control variables in our regressions. It is possible that some workers with a utility scale wind plant 19 miles away may be in our treatment group while others with utility scale wind 21 miles away are not, despite very similar exposure levels. However, this is a limitation of any study that chooses a hard spatial cutoff for treatment exposure. We deal with this by using the regression to explicitly control for the exposure that any worker who is just outside the edge of our treatment group might experience.

We pull from the entire dataset outside 20 miles for our control group. While some authors argue that the control group should first be narrowed using propensity score matching or inverse propensity weighting, we are comfortable with our approach for several reasons. First, most workers have at least some level of exposure to wind even if they don't live within the 20-mile band, as we can see in Figure 2 and Table 3, even if the exposure is small. Second, wind projects have a mix of urban, suburban, and rural exposure as can be seen in Figure 1 with wind projects near Chicago, Los Angeles, and Denver as well as small-to-medium-sized cities like Des Moine and Omaha. Third, the observable characteristics between treated and control groups are not incredibly different even in the near-population (Table 1. Finally, although our randomly selected control workers might have systematically different characteristics in levels from the treated workers, the parallel trends assumption is strongly satisfied in all but a few subsamples, as we will see in the results section. However, later versions of this paper will report results with a narrower selection of control group workers using propensity score methods.[6]

Specifically, our model for outcomes is

$$Y_{ict} = \gamma D_{ict} + X'_{ict}\beta + \alpha_{ic} + \mu_{st} + \epsilon_{ict} \tag{2}$$

where Y_{ict} is either the fraction of the year that person i is employed in year t and county c, or the log of earnings. In logged earnings regressions we restrict the sample to workers having nonzero earnings in at least two quarters of the year, so coefficients are impacts on earnings conditional on being employed.[7] In this case, γ is a semi-elasticity approximating the percent change in earnings for each unit change in the treatment variable. D_{ict} is a treatment variable defined as either an indicator for the presence of at least 10 MW of wind capacity in year t within a 20 mile radius of person i's residence, or as a continuous variable measuring the total capacity in gigawatts within 20 miles that year. 10MW is roughly equivalent to the 5th percentile of project size during the study period, and can be considered a minimum size threshold for a commercial wind project (BD Hoen et al., 2021; E. Brunner, Ben Hoen, and Hyman, 2022). The average exposure to wind capacity for treated workers within our sample is approximately 300MW, so in regressions with continuous treatment we can interpret $\gamma \cdot 0.3$ as the semi-elasticity at the mean, from the addition of one

[6] We thank Johannes Schmieder for this suggestion.

[7] Future versions of this paper will explore alternative definitions of earnings, such as the nonzero earnings at a given employer when the worker is also employed at the same job in the preceding and following quarter, or the annual earnings relative to a baseline year. We will also scale the left and right hand side variables by the size of the local economy so that impact estimates are applicable to communities of different sizes. We thank Johannes Schmieder for these suggestions.

average-sized wind project. X_{ict} is a vector of control variables. In our case, if wind development is spatially autocorrelated and there are regional economic spillovers, we would need to control for wind capacity at greater distances than 20 miles. X_{ict} is the spatial lag of wind capacity in increasing 20-mile donut distance bins from person i's residence in year t, e.g., capacity 20 to 40 miles away, 40 to 60, etc., out to 100 miles. As it is challenging to causally identify multiple spatial lags within the same diff-in-diff framework, we do not report these coefficients or give them a causal interpretation but rather control for them in order to causally identify γ. α_{ic} is a worker-by-county fixed effect, capturing time-invariant characteristics of the worker-place combination, such as the worker's productivity within a particular set of local work opportunities. Finally, μ_{st} is a state-by-year fixed effect capturing regional macroeconomic trends at the state level that may be correlated with both worker outcomes and wind energy development. With both sets of fixed effects, variation in the treatment variable comes from changes in wind development near a person's residence, during spells between major cross-county or cross-state moves, that occur independently of state-level macroeconomic trends.

In order to estimate equation (2) by LPDID, we take successively long differences of (2) and estimate them one at a time. Specifically, using the full near-population dataset and each sub-population of interest, we estimate

$$Y_{ic,t+h} - Y_{ict} = \delta_h \Delta D_{ict} + \Delta X'_{ict} \beta^h + \Delta \mu^h_{st} + \Delta \epsilon^h_{ict} \tag{3}$$

sequentially for different values of h, limiting the sample each time to newly treated ($\Delta D_{ict} > 0$) and not-yet-treated ($D_{ic,t+h} = 0$). The treatment coefficient δ_h in each regression is the event study estimate for event-time h. Averaging δ_h over the post-treatment periods gives an estimate of the average treatment effect, γ, from equation (2). Examining δ_h for negative values of h can help evaluate the parallel trends assumption, such as by plotting the individual coefficients on an event study graph and/or testing the significance of the cumulative sum of δ_h over pre-treatment periods. We report both approaches. Including temporal lags of Y_{ict} on the right hand side of the regression can help control away pre-existing differences in trends in order to obtain conditional parallel trends, while also controlling for potentially endogenous selection into treatment based on pre-treatment outcomes. Parallel trends tests are satisfied in almost all of our worker-level specifications, so we do not include temporal lags of Y_{ict} in this version of the paper. Note that we have differenced out the worker-by-county fixed effect α_{ic} which improves computation time, but this also means that each long difference is within a spell in which a worker lives in a particular county. Using separate worker fixed effects and county fixed effects, and differencing within worker (but not necessarily within county if a worker migrates), produces very similar estimates which we omit for the sake of brevity.

This approach economizes on computational resources, making estimation on very large datasets feasible, for several reasons. First, a very large set of worker-by-county fixed effects are removed by differencing before estimation. Second, because each regression in the sequence limits the sample to clean controls, the "stacking" approach of Cengiz et al., 2019 is not required. Stacking involves appending panels of treated units with panels of units that are not treated within the same time window, such that each treated unit has a set of clean controls. This approach significantly magnifies the size of the analysis dataset, requiring additional memory and RAM. In our setting with 1.4 billion observations on the near-universe of workers in 23 states observed over 20 years, this approach is not feasible. Third, because this is a regression-based estimator, the computationally intensive task of averaging thousands of individual 2 × 2 pre/post treatment comparisons, as in Callaway and Sant'Anna, 2021, or estimating shares of cohort weights as in Sun and Abraham, 2021 is not necessary. In fact, Dube et al., 2023 demonstrate that computation time for LPDID is comparable to two-way fixed effects estimation, and much faster compared to Callaway and Sant'Anna, 2021; Sun and Abraham, 2021 for which computation time is more than two orders of magnitude greater. This is a crucial benefit in our case using a massive restricted access dataset. Further, LPDID achieves the same reduction in treatment effect bias as these alternative new estimators when treatment arrival is exogenous, while achieving even less biased results when selection into treatment depends on past outcomes.

We estimate the regressions in equation (3) using the entire dataset, then estimate them separately for black, American Indian/Native Alaskan, white, Hispanic, male, and female workers, as well as those without a high school education, with a high school education, some college coursework or an Associates degree, and those with a Bachelors degree or higher.

4.4 Monte Carlo Estimation

Even with the computational advantages of LPDID, estimating equation (3) on the full LEHD dataset is not computationally feasible. In order to solve this problem, we repeatedly randomly sampled (with replacement) 1 million unique worker IDs. For each random 1-million-worker sample, we estimated the sequence of LPDID regressions in (3) for all available time periods and retained the δ_h coefficients. We repeated this exercise 100 times, and calculated the mean and standard deviation of coefficient estimates across the 100 draws. These coefficient means and standard deviations across draws are what we report in the results section below, rather than using analytical standard errors from a particular draw.

In order to deal with the fact that sub-populations are unevenly geographically distributed across the United States, we stratified our random samples by county. Further, we over-sampled counties with fewer workers in particular sub-population categories. Specifically, when randomly sampling sub-populations, we specified a fixed minimum number of workers to be drawn from each county such that the total number of unique workers would sum to one million. If too few workers in a given sub-population resided in a given county, we took all workers from that sub-population in that county, and reallocated the remaining workers to other counties in proportion to their population shares in order to reach one million workers in each draw.

4.5 County-level aggregation

We also wish to illustrate how the use of data aggregated to arbitrary and differently-sized administrative boundaries such as counties can bias impact estimates in general, and investigate the extent of this bias for different sub-populations. In order to do so, we aggregate the individual-level outcomes within a county and re-estimate the sequence of regressions from equation (3). In order to define analogous outcome variables at the county level, we calculate the average fraction of the year employed for all workers in each county-year, and we calculate the log of average earnings per worker who had non-zero earnings in at least two quarters of the year. We perform this calculation for the entire dataset and for each sub-population of interest. We then modify equation (2) by using county fixed effects (rather than worker-by-county fixed effects). We define the treatment variable as either a binary indicator for whether or not a given county has at least 10 MW of wind capacity, or the continuous number of gigawatts of wind capacity in the county.[8] We then

[8]To be consistent with the 20 mile ring used in the individual-level regressions, we also estimated regressions in which we defined treatment by aggregating capacity in all counties with centroids within 20 miles of a given county's centroid. Results are not noticeably different - most counties do not have another county centroid within 20 miles of their own centroid.

estimate the regressions from equation (3), as well as the shift-share exercise for equation (1) on this county-level aggregate data.

5 Results

We first report results for the average treatment effects at the individual worker, and county aggregate levels. These are averages of the δ_h coefficients from equation (3) in the post-treatment period, which are estimates of the γ parameter from equation (2). We then illustrate the dynamics of these effects through event study graphs.

5.1 Average Effects

Tables 4 and 5 show the average treatment effects (average of event study coefficients in the post-treatment period) for binary treatment on employment and earnings, respectively. Tables 6 and 7 in the Appendix report an analogous set of results for continuous wind capacity treatment. Panel A in each table shows the effects from the same model estimated on 100 Monte Carlo random draws of the geocoded worker-level data, reporting means and standard deviations of parameter estimates across draws. Panel B in each table shows the effects using the full data aggregated to the county level, with analytical standard errors clustered at the county level. The first column in each table reports these results for all workers, whereas each subsequent column reports results for each of the sub-populations that we study.

The first row and column of Table 4, for example, shows that being exposed to utility scale wind installations within 20 miles increases employment by 0.38 percentage points for the average worker. Given that approximately 55,000 workers live within 20 miles of a utility scale wind project, this translates to an average effect of approximately 208 jobs per project. We use similar calculations to translate coefficients into jobs for each subpopulation in Table 4. By contrast, previous studies find local impacts of approximately 50 to 90 jobs per project (Brown et al., 2012b; Gilbert, Gagarin, and Ben Hoen, 2023). In terms of an employment multiplier, this translates to approximately 0.46 jobs per million dollars in wind project investment.[9] This is a fairly modest employment multiplier compared to those in other industries or from federal stimulus that tend to be at least an order of magnitude larger (Chodorow-Reich, 2019).

Using a continuous treatment variable in GW of capacity yields similar results. The first row and column of Table 6 in the Appendix indicates that each additional GW of capacity within 20 miles increases the average worker's employment by 1.2 percentage points. The average capacity within 20 miles of a treated worker is 0.3 GW, which yields a treatment effect at the mean that is very similar to the binary treatment coefficient.

Among race and ethnicity subgroups, this impact is proportionately largest for black workers at 0.61 percentage points as compared to white workers at 0.33 percentage points. However, because more white workers live near wind projects, the total jobs impact is still much larger for white workers (147 jobs) versus black workers (24 jobs). The impact is much larger for male than for female workers (0.42 versus 0.29 percentage points increase in employment, translating to 119 versus 78 jobs). The effect is also proportionately largest for workers with either very low skill (no high school education) or high skill (college education), while again the aggregate jobs impact is still greater for more-educated workers (roughly 40 jobs for people with high school education or below,

[9]We arrive at this multiplier by the following calculation: utility scale wind installation costs are approximately \$1.5 million per MW, and the average exposure of treated workers is 300 MW within 20 miles, or \$450 million of wind investment. Dividing 208 jobs by \$450 million gives the result.

versus 60 for people with at least some college). In the worker level data, the cumulative pretrends are not statistically significant for any sub-population except for American Indian/Native American workers. The other key feature to note from Table 4 is that the county-level impact estimates are severely attenuated as compared to the worker-level. They are not statistically significant for any sub-population or overall, and with the exception of Hispanic and American Indian workers, the point estimates are quite a bit smaller than in the worker-level regressions - coming in at about 1/3 of the worker-level estimates on average.

Table 4: Average Treatment Effects (Binary > 10 MW w/in 20 miles): Employment

	All	Black	Am. Ind	White	Hispanic	Female	Male	No High	High Sch.	Some Coll	College
					Panel A: Worker-Level						
ATE	0.38	0.61	0.36	0.33	0.40	0.29	0.42	0.50	0.28	0.33	0.43
	(0.13)	(0.23)	(0.18)	(0.13)	(0.16)	(0.11)	(0.12)	(0.12)	(0.12)	(0.11)	(0.10)
Jobs	208	24	3.6	146.7	30	77.6	119.4	39.6	42	59.7	60.4
cumulative	0.0037	-0.33	-1.2	-0.0005	0.32	0.12	-0.011	-0.091	0.16	0.041	0.019
pretrend	(0.43)	(0.63)	(0.59)	(0.35)	(0.47)	(0.39)	(0.33)	(0.39)	(0.32)	(0.35)	(0.31)
					Panel B: County-Level						
ATE	0.14	0.073	0.48	0.15	0.48	0.13	0.15	0.19	0.13	0.054	0.26
	(0.12)	(0.96)	(0.89)	(0.11)	(0.49)	(0.12)	(0.14)	(0.23)	(0.14)	(0.13)	(0.15)
cumulative	0.22	-0.62	-0.63	0.25	0.66	0.15	0.29	0.44	0.14	0.070	0.63
pretrend	(0.36)	(3.5)	(2.8)	(0.34)	(1.7)	(0.38)	(0.43)	(0.71)	(0.44)	(0.43)	(0.58)

Notes: This table reports the average of event study coefficients (δ_h) in the post-treatment period as "Average Treatment Effects", for both worker-level and county-level regressions. These are estimates of γ from equation (2). The dependent variable is the percentage of the year in which a worker had non-zero earnings. The treatment is a dummy variable equal to 1 if the worker had at least 10 MW of capacity within 20 miles of their residence, or if the county had at least 10 MW of capacity in county-level regressions. Aggregating to 20 miles around county centroids produced similar results which are omitted here. Worker level estimates are parameter averages and standard deviations across 100 model estimates from repeated random draws from the near-population. County-level estimates use the full dataset aggregated to the county level, with standard errors clustered at the county level. The cumulative pre-trends test reports the sum of event study coefficients over the pre-treatment period, and its standard deviation across draws (worker level) or analytical standard error (county level).

Similar patterns arise when examining Table 5 regarding impacts on log earnings (conditional on being employed). The first row and column of Table 5 shows that earnings increase by 3.6 percent following the arrival of a nearby wind project. This is also somewhat larger than previous literature which has fairly mixed results for earnings impacts ranging from no impact to about three percent (E. Brunner, Ben Hoen, and Hyman, 2022). Considering the average earnings of treated workers in our sample is $31,100, this translates to an impact of approximately $1,111 per employed worker per year. In terms of an earnings multiplier, this is equivalent to about 0.14 dollars in worker earnings per dollar invested in local wind capacity. This is also modest compared to similar infusions of spending such as government fiscal stimulus, which has earnings multipliers estimated to range from 0.3 to 2, but with most estimates falling between 0.6 and 1 (Ramey, 2019).

As with employment, using a continuous treatment variable in GW of capacity yields similar results for earnings as the binary treatment variable. The first row and column of Table 7 in the Appendix indicates that each additional GW of capacity within 20 miles increases the average employed worker's earnings by 11 percentage points. With an average treated capacity of 0.3 GW the treatment effect at the mean is again very similar to the binary treatment coefficient.

These impacts are proportionately largest for black workers and workers without a high school education, and are heavily imbalanced for men versus women. Unlike the employment impacts, these proportionate impacts are also reflected in the actual dollar magnitudes of earnings impacts. Black workers increase earnings by over $1,200 per year while workers without a high school education gain over $1,000. However, the increase in male earnings is almost 3 times that of female earnings, and the largest earnings increase is enjoyed by college graduates at $1,718 per year. These impacts are statistically significant for all sub-populations except for Hispanic and American Indian workers. Cumulative pre-trends are again only significant for American Indian workers. When comparing earnings impacts at the worker level in Panel A to the county level in Panel B, however, we see even more dramatic attenuation at the county level for earnings than for employment. Almost all county-level estimates are close to zero and many are the opposite sign. None are statistically significant. They are particularly biased downward for black and Hispanic workers, and workers either with a college degree or without a high school diploma.

Table 5: Average Treatment Effects (Binary > 10 MW w/in 20 miles): Earnings

	All	Black	Am. Ind	White	Hispanic	Female	Male	No High	High Sch.	Some Coll	College
					Panel A: Worker-Level						
ATE	3.6	5.4	3.7	3.1	3.0	2.5	4.5	5.4	2.6	3.5	3.6
	(1.5)	(2.4)	(2.1)	(1.4)	(1.6)	(1.2)	(1.4)	(1.3)	(1.2)	(1.2)	(1.2)
Earnings ($)	1,111	1,218	682	965	743	599	1,682	1,019	629	1,036	1,718
cumulative	-2.7	-0.56	-13	-2.0	3.4	-1.7	-1.6	-2.5	0.74	-1.9	-4.4
pretrend	(4.4)	(6.9)	(6.5)	(3.8)	(4.7)	(4.1)	(3.9)	(4.1)	(3.0)	(3.6)	(3.2)
					Panel B: County-Level						
ATE	-0.13	-1.1	1.4	-0.16	-0.52	-0.077	-0.13	-0.69	-0.071	0.0039	-0.72
	(0.31)	(3.5)	(2.2)	(0.31)	(1.2)	(0.26)	(0.39)	(0.63)	(0.39)	(0.35)	(0.42)
cumulative	0.44	4.1	6.6	0.45	-3.3	0.25	0.63	-1.2	0.72	-0.21	1.2
pretrend	(1.0)	(12)	(7.2)	(0.98)	(4.3)	(0.82)	(1.3)	(1.8)	(1.4)	(1.1)	(1.3)

Notes: This table reports the average of event study coefficients (δ_h) in the post-treatment period as "Average Treatment Effects", for both worker-level and county-level regressions. These are estimates of γ from equation (2). The dependent variable is the log of earnings, with the sample limited to employed workers (those with at least two quarters of non-zero earnings in a year). The treatment is a dummy variable equal to 1 if the worker had at least 10 MW of capacity within 20 miles of their residence, or if the county had at least 10 MW of capacity in county-level regressions. Aggregating to 20 miles around county centroids produced similar results which are omitted here. Worker level estimates are parameter averages and standard deviations across 100 model estimates from repeated random draws from the near-population. County-level estimates use the full dataset aggregated to the county level, with standard errors clustered at the county level. The cumulative pre-trends test reports the sum of event study coefficients over the pre-treatment period, and its standard deviation across draws (worker level) or analytical standard error (county level).

5.2 Event Study Graphs

Figures 3 and 4 plot the LPDID event study coefficients for employment and log of earnings, for worker-level and county-aggregate data. We normalize to two-years before any wind project is operational in order to capture potential construction effects in the year preceding the first year of operation.

We can see several important findings in both figures. First, we can see from both figures that impacts are not concentrated around the construction phase of event-years -1 and 0. Rather, impacts grow over time and persist as many as six years after the project becomes operational. This is consistent with the hypothesis that there are many indirect channels through which renewable energy can improve the local economy, not just related to direct employment at the renewable energy installation. For example, because of local landowners spending royalty payments, additional community services stimulated by additional local tax payments, or other indirect channels.

Second, we can see from both figures that error bars on the worker-level estimates are relatively large compared to county-level error bars, despite the fact that worker-level sample sizes are in the millions while county-level sample sizes are in the thousands. This suggests that there is significant heterogeneity in impacts that is picked up through our repeated random sampling procedure.

Third, the figures help visualize the degree of attenuation or downard bias from using county-level aggregates. Figure 3 shows that county-level estimates are smaller, and still quite noisy, yet point estimates are on average positive and less than half of the worker-level estimates. By contrast, Figure 4 shows that earnings estimates at using county-level data are, essentially, precisely estimated zeroes. This suggests that earnings impacts as reported in the literature are likely understating true earnings impacts by a wider margin than employment estimates.

Figure 3: Impact of Wind Capacity on Employment: Worker-Level vs. County-Level

Notes: These are event study estimates from Equation (3) using worker-level data versus county-aggregate data. Treatment is a binary variable for whether at least 10 MW of wind has arrived within 20 miles of the worker's residence, or arrived within the county. Event year 0 equates to the year wind project operations began. Wind project construction may have occurred in the preceding year, so we normalize to event-year -2.

Figure 4: Impact of Wind Capacity on Log Earnings: Worker-Level vs. County-Level

Notes: These are event study estimates from Equation (3) using worker-level data versus county-aggregate data. Treatment is a binary variable for whether at least 10 MW of wind has arrived within 20 miles of the worker's residence, or arrived within the county. Event year 0 equates to the year wind project operations began. Wind project construction may have occurred in the preceding year, so we normalize to event-year -2. The dependent variable is logged earnings, so impact estimates are semi-elasticities, or approximately percentage changes in annual earnings.

5.2.1 Event Studies: Race and Ethnicity

Figures 5 and 6 also confirm the average treatment effects in the previous tables. The effect on black workers for both employment and earnings is both larger and more persistent than for other races and ethnicities; as we saw in Table 1, despite there being relatively few black workers with wind projects within 20 miles, the average "treated" black worker is also exposed to significantly more capacity. Yet Tables 6 and 7 in the Appendix show that the marginal impact of an additional gigawatt of wind capacity is also larger for black workers.

5.2.2 Event Studies: Sex

Figures 7 and 8 further show that impacts on employment and earnings are consistently larger for male workers than female workers. However, there is much overlap in the confidence intervals at each time step. Qualitatively, the impacts on men start sooner, are larger, and persist longer than for women.

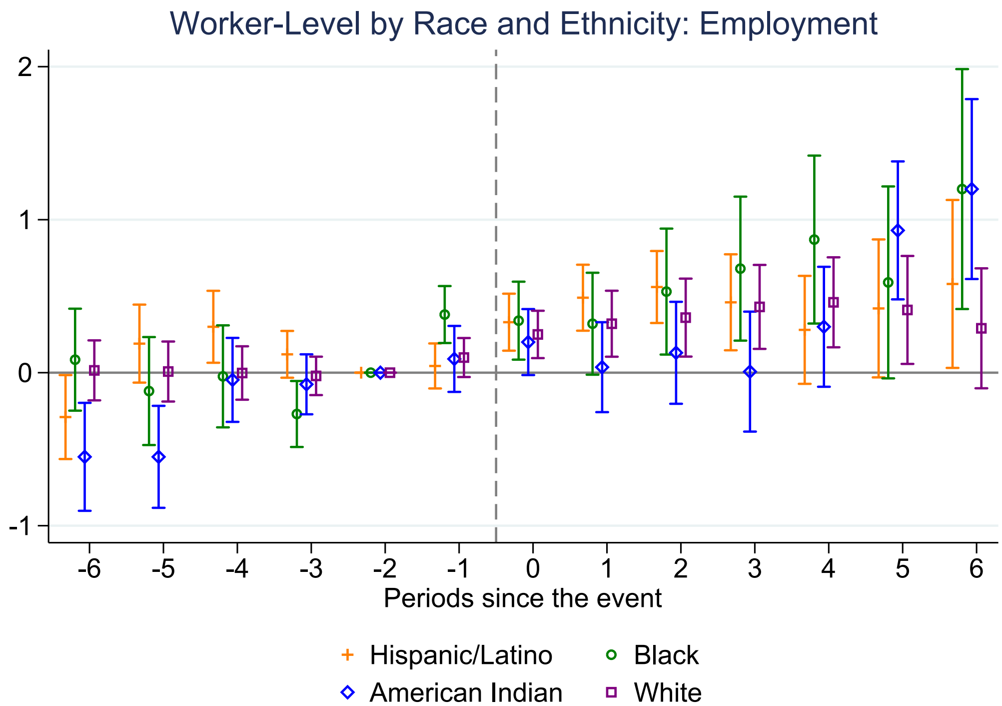

Figure 5: Impact of Wind Capacity on Employment: Worker-Level vs. County-Level

Notes: These are event study estimates from Equation (3) using worker-level data versus county-aggregate data. Treatment is a binary variable for whether at least 10 MW of wind has arrived within 20 miles of the worker's residence, or arrived within the county. Event year 0 equates to the year wind project operations began. Wind project construction may have occurred in the preceding year, so we normalize to event-year -2.

5.2.3 Event Studies: Educational Attainment

Finally, figures 9 and 10 again show persistent impacts that are largest for workers without a high school education, or workers with a college degree. However, impacts for all educational categories are statistically significant and persistent for many years following the arrival of a wind project.

6 Conclusion

As the United States continues to make unprecedented investments in renewable energy in order to meet carbon emissions goals, this will shift the demand for skilled and unskilled labor and generate new sources of income, tax revenues, and expenditures, possibly in places that have not previously been major energy producing communities. These developments could either continue to allocate benefits to privileged groups while continuing to restrict access to disadvantaged groups, or they could increase access to economic opportunity among vulnerable populations.

In this paper we use restricted-access geocoded data on the near-universe of workers in 23 U.S. states in order to estimate the local earnings and employment impacts of wind energy development. We estimate these effects for all workers, and separately for black, American Indian/Native Alaskan, white, and Hispanic workers, male versus female workers, and those without a high school diploma, with a high school diploma, some college coursework, and with a college degree. We then aggregate this data to the county level in order to compare our estimates with those we would have obtained

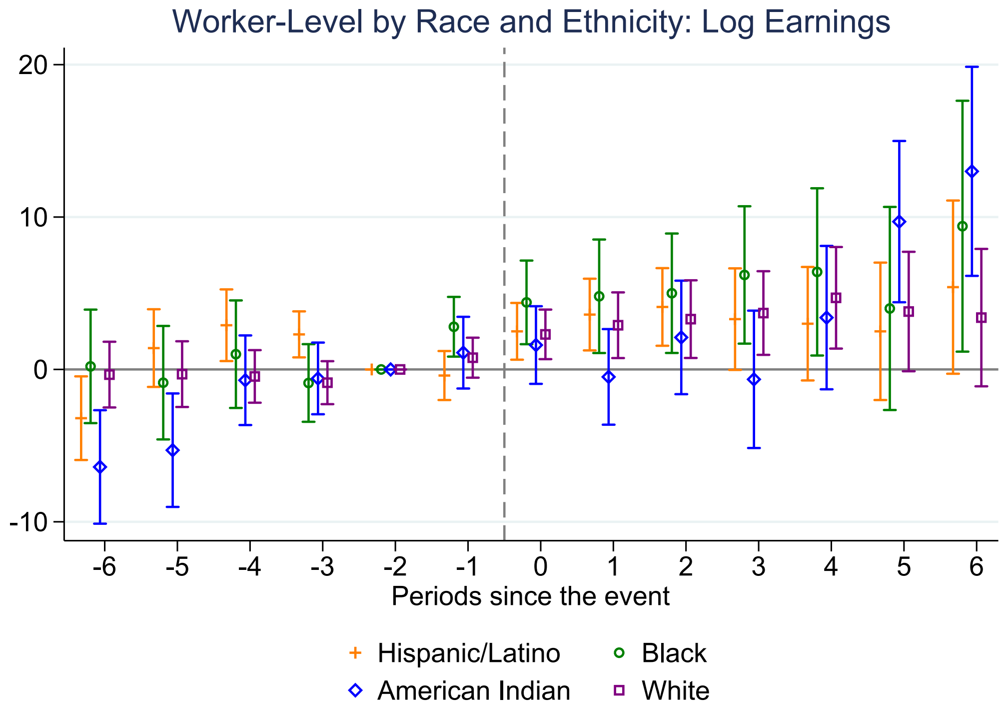

Figure 6: Impact of Wind Capacity on Log Earnings: Worker-Level vs. County-Level

Notes: These are event study estimates from Equation (3) using worker-level data versus county-aggregate data. Treatment is a binary variable for whether at least 10 MW of wind has arrived within 20 miles of the worker's residence, or arrived within the county. Event year 0 equates to the year wind project operations began. Wind project construction may have occurred in the preceding year, so we normalize to event-year -2. The dependent variable is logged earnings, so impact estimates are semi-elasticities, or approximately percentage changes in annual earnings.

with county-level aggregates using data such as is available in the public domain.

We find economically and statistically significant employment and earnings gains from wind development within 20 miles of a worker's residence. We also find that these impacts are relatively more pronounced for black workers, men, and very low skilled or high skilled workers. These impacts persist for years after the construction phase ends, suggesting that there may be multiple indirect channels through which wind capacity in place in a community can generate benefits.

We also find that impacts are drastically understated when using county-level aggregate data, and that the degree of downward bias varies across sub-populations in ways that are not obviously predictable. This finding suggests that there is also in inequity within the research community in terms of which researchers have access to better data in order to generate accurate estimates of impacts on the communities who may be served by their institutions.

Our study is limited in several respects. We were not able to access data on two major wind energy states such as Texas and Minnesota. We were also not able to explore impacts of other energy sources or compare a variety of identification strategies because of limitations on computation time using such a large dataset. We were also not able to look at more intersectional outcomes, starting with more race and ethnicity categories but also including workers with more than one race or ethnicity, non-gender-binary workers, or more granularity in worker skill level. We further did not consider impacts on migration decisions, or impacts on workers in specific industrial sectors.

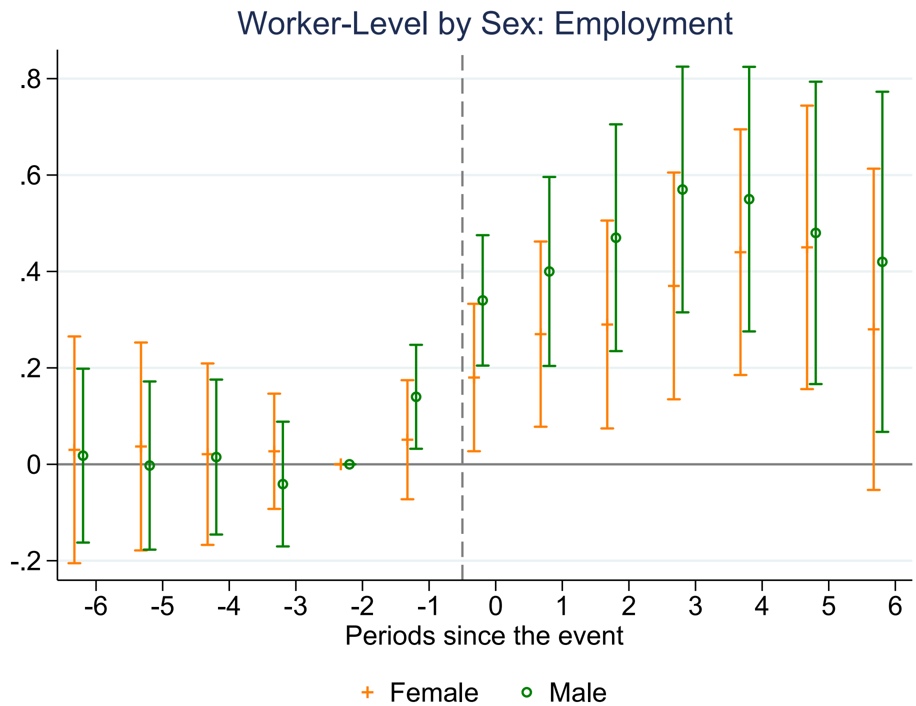

Figure 7: Impact of Wind Capacity on Employment: Worker-Level vs. County-Level

Notes: These are event study estimates from Equation (3) using worker-level data versus county-aggregate data. Treatment is a binary variable for whether at least 10 MW of wind has arrived within 20 miles of the worker's residence, or arrived within the county. Event year 0 equates to the year wind project operations began. Wind project construction may have occurred in the preceding year, so we normalize to event-year -2.

These are all important areas for future research. Similarly, understanding the "boundary" of a local community and its economy is also an interesting question for future research. While we have followed the common approach from the literature of using concentric circle distances, other methods such as isochrones for commuting time may be interesting to explore.[10]

[10]We thank Justin Kirkpatrick for this suggestion.

Figure 8: Impact of Wind Capacity on Log Earnings: Worker-Level vs. County-Level

Notes: These are event study estimates from Equation (3) using worker-level data versus county-aggregate data. Treatment is a binary variable for whether at least 10 MW of wind has arrived within 20 miles of the worker's residence, or arrived within the county. Event year 0 equates to the year wind project operations began. Wind project construction may have occurred in the preceding year, so we normalize to event-year -2. The dependent variable is logged earnings, so impact estimates are semi-elasticities, or approximately percentage changes in annual earnings.

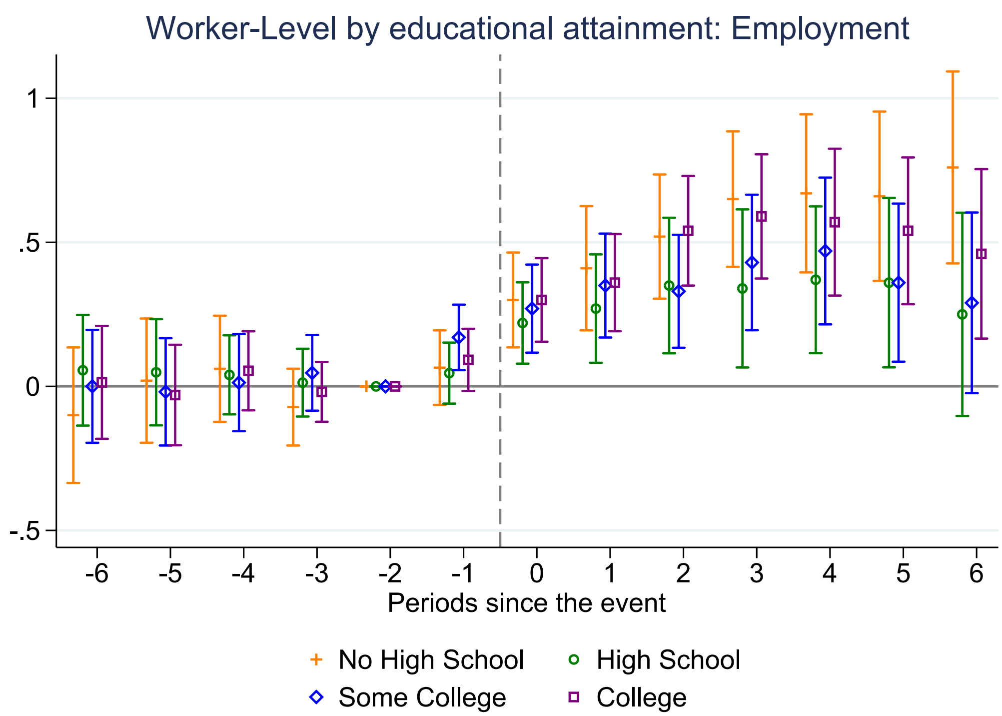

Figure 9: Impact of Wind Capacity on Employment: Worker-Level vs. County-Level

Notes: These are event study estimates from Equation (3) using worker-level data versus county-aggregate data. Treatment is a binary variable for whether at least 10 MW of wind has arrived within 20 miles of the worker's residence, or arrived within the county. Event year 0 equates to the year wind project operations began. Wind project construction may have occurred in the preceding year, so we normalize to event-year -2.

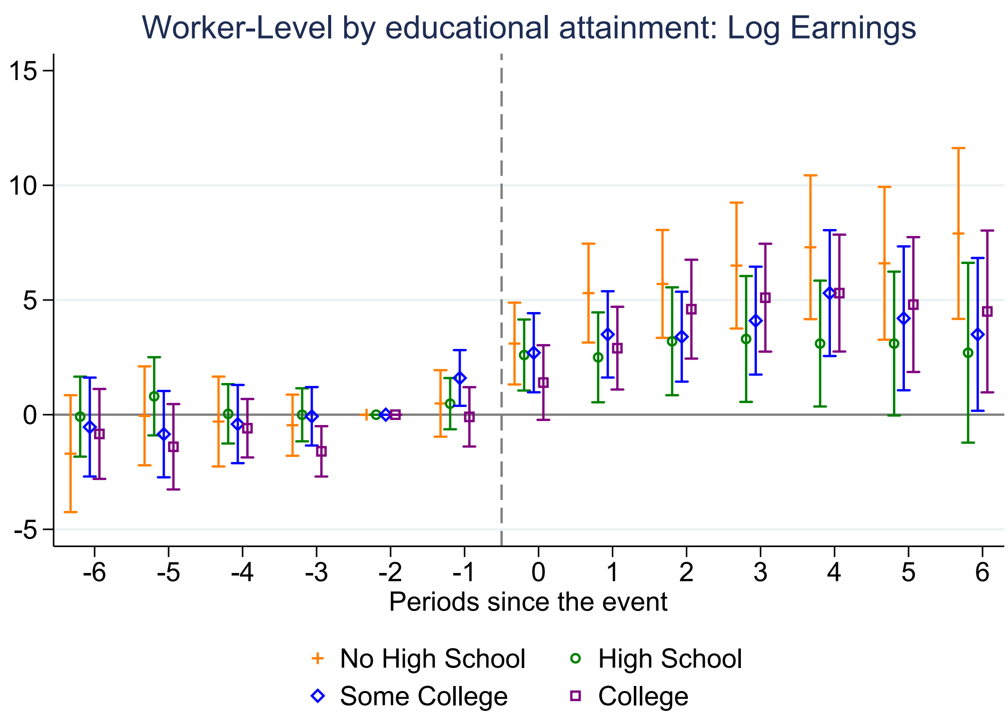

Figure 10: Impact of Wind Capacity on Log Earnings: Worker-Level vs. County-Level

Notes: These are event study estimates from Equation (3) using worker-level data versus county-aggregate data. Treatment is a binary variable for whether at least 10 MW of wind has arrived within 20 miles of the worker's residence, or arrived within the county. Event year 0 equates to the year wind project operations began. Wind project construction may have occurred in the preceding year, so we normalize to event-year -2. The dependent variable is logged earnings, so impact estimates are semi-elasticities, or approximately percentage changes in annual earnings.

References

Kim, Haeyeon (2019). “Three Essays on Mineral Economics and Economic Geography”. English. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2023-03-05. PhD thesis, p. 83. ISBN: 9781085771788. URL: http://mines.idm.oclc.org/login?url=https://www.proquest.com/dissertations-theses/three-essays-on-mineral-economics-economic/docview/2305191102/se-2.

Mueller, J Tom and Matthew M Brooks (2020). “Burdened by renewable energy? A multi-scalar analysis of distributional justice and wind energy in the United States”. In: Energy Research and Social Science 63 (December 2019), p. 101406. ISSN: 22146296. DOI: 10.1016/j.erss.2019.101406. URL: https://doi.org/10.1016/j.erss.2019.101406.

Mejía-Montero, Adolfo et al. (2021). “Grounding the energy justice lifecycle framework: An exploration of utility-scale wind power in Oaxaca, Mexico”. In: Energy Research and Social Science 75 (February), p. 102017. ISSN: 22146296. DOI: 10.1016/j.erss.2021.102017. URL: https://doi.org/10.1016/j.erss.2021.102017.

McCauley, Darren et al. (2019). “Energy justice in the transition to low carbon energy systems: Exploring key themes in interdisciplinary research”. In: Applied Energy 233-234 (November 2018), pp. 916–921. ISSN: 03062619. DOI: 10.1016/j.apenergy.2018.10.005.

Carley, Sanya and David M Konisky (2020). “The justice and equity implications of the clean energy transition”. In: Nature Energy 5 (8), pp. 569–577. ISSN: 20587546. DOI: 10.1038/s41560-020-0641-6. URL: http://dx.doi.org/10.1038/s41560-020-0641-6.

Dube, Arindrajit et al. (2023). A local projections approach to difference-in-differences event studies. Tech. rep. National Bureau of Economic Research.

Callaway, Brantly and Pedro HC Sant’Anna (2021). “Difference-in-differences with multiple time periods”. In: Journal of Econometrics 225.2, pp. 200–230.

Sun, Liyang and Sarah Abraham (2021). “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects”. In: Journal of Econometrics 225.2, pp. 175–199.

Cengiz, Doruk et al. (2019). “The effect of minimum wages on low-wage jobs”. In: The Quarterly Journal of Econom pp. 1405–1454. DOI: 10.1093/qje/qjz014.Advance.

Colmer, Jonathan, John Voorheis, and Brennan Williams (2021). “Air Pollution and Economic Opportunity in the United States”. In.

Gilbert, Ben, Hannah Gagarin, and Ben Hoen (2023). Geographic Spillovers of Wind Energy Development on Wag Tech. rep.

Brunner, Eric J and David Schwegman (2022). “Commercial Wind Energy Installations and Local Economic Development Evidence from U.S. Counties”. In: SSRN Electronic Journal 165.March, p. 112993. ISSN: 0301-4215. DOI: 10.2139/ssrn.4030617. URL: https://doi.org/10.1016/j.enpol.2022.112993.

Brown, Jason P et al. (2012a). “Ex post analysis of economic impacts from wind power development in U.S. counties”. In: Energy Economics 34 (6), pp. 1743–1754. ISSN: 01409883. DOI: 10.1016/j.eneco.2012.07.010. URL: http://dx.doi.org/10.1016/j.eneco.2012.07.010.

Brunner, Eric, Ben Hoen, and Joshua Hyman (2022). “School district revenue shocks, resource allocations, and student achievement: Evidence from the universe of U.S. wind energy installations”. In: Journal of Public Economics 206, p. 104586. ISSN: 00472727. DOI: 10.1016/j.jpubeco.2021.104586. URL: https://doi.org/10.1016/j.jpubeco.2021.104586.

Castleberry, Becca and J Scott Greene (2017). “Impacts of wind power development on Oklahoma’s public schools”. In: Energy, Sustainability and Society 7.1. ISSN: 21920567. DOI: 10.1186/s13705-017-0138-8.

Shoeib, Eman Ahmed Hamed, Elisabeth Hamin Infield, and Henry C. Renski (June 2021). “Measuring the impacts of wind energy projects on U.S. rural counties’ community services and cost of living”. In: Energy Policy 153. ISSN: 03014215. DOI: 10.1016/J.ENPOL.2021.112279.
Miller, Clark A, Jennifer Richter, and Jason O’Leary (2015). “Socio-energy systems design: A policy framework for energy transitions”. In: Energy Research and Social Science 6, pp. 29–40. ISSN: 22146296. DOI: 10.1016/j.erss.2014.11.004. URL: http://dx.doi.org/10.1016/j.erss.2014.11.004.
Borenstein, Severin and Lucas W Davis (2016). “The distributional effects of US clean energy tax credits”. In: Tax Policy and the Economy 30.1, pp. 191–234. ISSN: 08928649. DOI: 10.1086/685597.
Sunter, Deborah A, Sergio Castellanos, and Daniel M Kammen (2019). “Disparities in rooftop photovoltaics deployment in the United States by race and ethnicity”. In: Nature Sustainability 2.1, pp. 71–76. ISSN: 23989629. DOI: 10.1038/s41893-018-0204-z. URL: http://dx.doi.org/10.1038/s41893-018-0204-z.
Black, Dan, Terra McKinnish, and Seth Sanders (2005). “The economic impact of the coal boom and bust”. In: Economic Journal 115 (503), pp. 449–476. ISSN: 00130133. DOI: 10.1111/j.1468-0297.2005.00996.x.
Tsvetkova, Alexandra and Mark D Partridge (2016). “Economics of modern energy boomtowns: Do oil and gas shocks differ from shocks in the rest of the economy?” In: Energy Economics 59, pp. 81–95. ISSN: 01409883. DOI: 10.1016/j.eneco.2016.07.015. URL: http://dx.doi.org/10.1016/j.eneco.2016.07.015.
Weinstein, Amanda L, Mark D Partridge, and Alexandra Tsvetkova (2018). “Follow the money: Aggregate, sectoral and spatial effects of an energy boom on local earnings”. In: Resources Policy 55.December 2017, pp. 196–209. ISSN: 03014207. DOI: 10.1016/j.resourpol.2017.11.018.
Maniloff, Peter and Ralph Mastromonaco (2017). “The local employment impacts of fracking: A national study”. In: Resource and Energy Economics 49, pp. 62–85. ISSN: 09287655. DOI: 10.1016/j.reseneeco.2017.04.005. URL: http://dx.doi.org/10.1016/j.reseneeco.2017.04.005.
Varela-Vázquez, Pedro and María Del Carmen Sánchez-Carreira (2015). Socioeconomic impact of wind energy on p
DOI: 10.1016/j.rser.2015.05.045.
Brown, Jason P et al. (2012b). “Ex post analysis of economic impacts from wind power development in U.S. counties”. In: Energy Economics 34.6, pp. 1743–1754. ISSN: 01409883. DOI: 10.1016/j.eneco.2012.07.010. URL: http://dx.doi.org/10.1016/j.eneco.2012.07.010.
Acheampong, Alex O, Michael Odei Erdiaw-Kwasie, and Matthew Abunyewah (2021). “Does energy accessibility improve human development? Evidence from energy-poor regions”. In: Energy Economics 96, p. 105165. ISSN: 01409883. DOI: 10.1016/j.eneco.2021.105165. URL: https://doi.org/10.1016/j.eneco.2021.105165.
Fergen, Joshua and Jeffrey B. Jacquet (2016). “Beauty in motion: Expectations, attitudes, and values of wind energy development in the rural U.S”. In: Energy Research and Social Science 11, pp. 133–141. ISSN: 22146296. DOI: 10.1016/j.erss.2015.09.003.
Guo, Yue et al. (2015). “Not in my backyard, but not far away from me: Local acceptance of wind power in China”. In: Energy 82, pp. 722–733. ISSN: 03605442. DOI: 10.1016/j.energy.2015.01.082. URL: http://dx.doi.org/10.1016/j.energy.2015.01.082.
Kontogianni, A et al. (2014). “Planning globally, protesting locally: Patterns in community perceptions towards the installation of wind farms”. In: Renewable Energy 66, pp. 170–177. ISSN: 09601481. DOI: 10.1016/j.renene.2013.11.074. URL: http://dx.doi.org/10.1016/j.renene.2013.11.074.

Jones, Christopher R and J Richard Eiser (2010). "Understanding 'local' opposition to wind development in the UK: How big is a backyard?" In: Energy Policy 38 (6), pp. 3106–3117. ISSN: 03014215. DOI: 10.1016/j.enpol.2010.01.051. URL: http://dx.doi.org/10.1016/j.enpol.2010.01.051.

Shoeib, Eman Ahmed Hamed, Henry C Renski, and Elisabeth Hamin Infield (2022). "Who benefits from Renewable Electricity? The differential effect of wind power development on rural counties in the United States". In: Energy Research and Social Science 85.October 2021, p. 102398. ISSN: 22146296. DOI: 10.1016/j.erss.2021.102398. URL: https://doi.org/10.1016/j.erss.2021.102398.

Barry, John, Geraint Ellis, and Clive Robinson (2008). "Cool rationalities and hot air: A rhetorical approach to understanding debates on renewable energy". In: Global Environmental Politics 8 (2), pp. 67–98. ISSN: 15263800. DOI: 10.1162/glep.2008.8.2.67.

Zárate-Toledo, Ezequiel, Rodrigo Patiño, and Julia Fraga (2019). "Justice, social exclusion and indigenous opposition: A case study of wind energy development on the Isthmus of Tehuantepec, Mexico". In: Energy Research and Social Science 54 (March), pp. 1–11. ISSN: 22146296. DOI: 10.1016/j.erss.2019.03.004. URL: https://doi.org/10.1016/j.erss.2019.03.004.

Mills, Sarah Banas, Douglas Bessette, and Hannah Smith (2019). "Exploring landowners' post-construction changes in perceptions of wind energy in Michigan". In: Land Use Policy 82 (May 2018), pp. 754–762. ISSN: 02648377. DOI: 10.1016/j.landusepol.2019.01.010. URL: https://doi.org/10.1016/j.landusepol.2019.01.010.

Acheampong, Alex O, Janet Dzator, and Muhammad Shahbaz (2021). "Empowering the powerless: Does access to energy improve income inequality?" In: Energy Economics 99, p. 105288. ISSN: 01409883. DOI: 10.1016/j.eneco.2021.105288. URL: https://doi.org/10.1016/j.eneco.2021.105288.

Farrigan, Tracey (2021). "Data show U.S. poverty rates in 2019 higher in rural areas than in urban for racial/ethnic groups". In: USDA ERS. URL: https://www.ers.usda.gov/data-products/chart-gallery/gallery/chart-detail/?chartId=101903.

Shoeib, Eman Ahmed Hamed, Henry C Renski, and Elisabeth Hamin Infield (2022). "Who benefits from Renewable Electricity? The differential effect of wind power development on rural counties in the United States". In: Energy Research & Social Science 85, p. 102398.

Mauritzen, Johannes (2020). "Will the locals benefit?: The effect of wind power investments on rural wages". In: Energy Policy 142.April, p. 111489. ISSN: 03014215. DOI: 10.1016/j.enpol.2020.111489. URL: https://doi.org/10.1016/j.enpol.2020.111489.

Pedden, M. (2006). "Analysis: Economic Impacts of Wind Applications in Rural Communities". In: 11 (4). ISSN: 10991824. DOI: 10.1002/WE.273.

Hoen, Ben et al. (2015). "Spatial Hedonic Analysis of the Effects of US Wind Energy Facilities on Surrounding Property Values". In: Journal of Real Estate Finance and Economics 51.1, pp. 22–51. ISSN: 08955638. DOI: 10.1007/s11146-014-9477-9.

Abowd, John M et al. (2013). The LEHD Infrastructure Files and The Creation of The Quarterly Workforce Indic pp. 149–230. ISBN: 9780226172569. DOI: 10.7208/chicago/9780226172576.003.0006.

McKinney, Kevin L and John M Abowd (2020). "Male earnings volatility in LEHD before, during, and after the Great Recession". In: arXiv, pp. 1–54. ISSN: 23318422.

Card, David, Jesse Rothstein, and Moises Yi (2021). Location , Location , Location. Tech. rep. Center for Economic Studies.

Hoen, BD et al. (2021). "United States Wind Turbine Database (ver. 4.0)". In: US Geological Survey, American Wi

Rand, Joseph T et al. (2020). "A continuously updated, geospatially rectified database of utility-scale wind turbines in the United States". In: Scientific data 7.1, p. 15.

Feyrer, James, Erin T Mansur, and Bruce Sacerdote (2017). "Geographic dispersion of economic shocks: Evidence from the fracking revolution". In: American Economic Review 107.4, pp. 1905–1913. ISSN: 19447981. DOI: 10.1257/aer.20180888.
James, Alexander G and Brock Smith (2020). "Geographic dispersion of economic shocks: Evidence from the fracking revolution: Comment". In: American Economic Review 110.6, pp. 1905–1913. ISSN: 19447981. DOI: 10.1257/aer.20180888.
Jacobsen, Grant D, Dominic P Parker, and Justin B Winikoff (2023). "Are Resource Booms a Blessing or a Curse?: Evidence from People (Not Places)". In: Journal of Human Resources 58.2, pp. 393–420.
Chodorow-Reich, Gabriel (2019). "Geographic cross-sectional fiscal spending multipliers: What have we learned?" In: American Economic Journal: Economic Policy 11.2, pp. 1–34.
Ramey, Valerie A (2019). "Ten years after the financial crisis: What have we learned from the renaissance in fiscal research?" In: Journal of Economic Perspectives 33.2, pp. 89–114.

Appendix

Table 6: Average Treatment Effects (Continuous GW w/in 20 miles): Employment

	All	Black	Am. Ind	White	Hispanic	Female	Male	No High	High Sch.	Some Coll	College
						Panel A: Worker-Level					
ATE	1.2	2.2	0.25	1.0	0.89	0.74	1.5	1.3	0.81	1.2	1.1
	(0.60)	(0.87)	(0.90)	(0.59)	(0.66)	(0.52)	(0.57)	(0.56)	(0.55)	(0.50)	(0.49)
cumulative pretrend	-0.99	-2.6	-6.0	-1.3	-0.30	-0.47	-0.99	-1.1	-0.87	-0.056	-2.2
	(1.5)	(2.2)	(2.2)	(1.4)	(1.5)	(1.4)	(1.3)	(1.4)	(1.3)	(1.5)	(1.2)
						Panel B: County-Level					
ATE	0.59	-2.4	-2.0	0.63	-1.3	0.29	0.84	1.3	0.50	0.18	0.53
	(0.55)	(3.5)	(3.4)	(0.53)	(1.5)	(0.49)	(0.73)	(1.2)	(0.69)	(0.54)	(0.78)
cumulative pretrend	1.43	-10	2.7	0.83	3.2	2.2	0.72	4.4	0.57	1.0	2.4
	(1.2)	(10)	(12)	(1.1)	(4.3)	(1.4)	(1.4)	(2.7)	(1.5)	(1.5)	(1.8)

Notes: This table reports the average of event study coefficients (δ_h) in the post-treatment period as "Average Treatment Effects", for both worker-level and county-level regressions. These are estimates of γ from equation (2). The dependent variable is the percentage of the year in which a worker had non-zero earnings. The treatment is a continuous measure of the gigawatts (GW) of capacity within 20 miles of a worker's residence, or within the county in county-level regressions. Aggregating to 20 miles around county centroids produced similar results which are omitted here. Worker level estimates are parameter averages and standard deviations across 100 model estimates from repeated random draws from the near-population. County-level estimates use the full dataset aggregated to the county level, with standard errors clustered at the county level. The cumulative pre-trends test reports the sum of event study coefficients over the pre-treatment period, and its standard deviation across draws (worker level) or analytical standard error (county level).

Table 7: Average Treatment Effects (Continuous GW w/in 20 miles): Earnings

	All	Black	Am. Ind	White	Hispanic	Female	Male	No High	High Sch.	Some Coll	College
					Panel A: Worker-Level						
ATE	11	24	3.6	9.7	2.8	6.2	16	14	8.2	12	9.8
	(6.5)	(9.6)	(10)	(6.5)	(6.7)	(5.7)	(6.1)	(5.9)	(5.9)	(5.6)	(5.3)
cumulative	-19	-6.5	-71	-21	-9.9	-18	-15	-15	-8.1	-15	-36
pretrend	(17)	(24)	(26)	(15)	(16)	(15)	(14)	(15)	(14)	(15)	(12)
					Panel B: County-Level						
ATE	-1.4	0.66	-4.4	-1.5	-5.5	-0.87	-1.6	-4.6	-0.67	-1.3	-2.6
	(2.2)	(11.7)	(8.1)	(2.3)	(3.8)	(1.6)	(2.6)	(3.8)	(2.0)	(2.2)	(2.9)
cumulative	1.1	17	-3.1	0.88	-9.7	3.6	0.38	-9.7	1.9	0.17	5.4
pretrend	(3.6)	(37)	(24)	(3.6)	(11)	(3.7)	(4.2)	(7.1)	(4.5)	(4.9)	(4.8)

Notes: This table reports the average of event study coefficients (δ_h) in the post-treatment period as "Average Treatment Effects", for both worker-level and county-level regressions. These are estimates of γ from equation (2). The dependent variable is the log of earnings, with the sample limited to employed workers (those with at least two quarters of non-zero earnings in a year). The treatment is a continuous measure of the gigawatts (GW) of capacity within 20 miles of a worker's residence, or within the county in county-level regressions. Aggregating to 20 miles around county centroids produced similar results which are omitted here. Worker level estimates are parameter averages and standard deviations across 100 model estimates from repeated random draws from the near-population. County-level estimates use the full dataset aggregated to the county level, with standard errors clustered at the county level. The cumulative pre-trends test reports the sum of event study coefficients over the pre-treatment period, and its standard deviation across draws (worker level) or analytical standard error (county level).