A Place-Based Clean Electricity Subsidy and its Effect on Emissions

Matthew Ashenfarb

Yale University · School of the Environment

NBER Distributional Effects of New Clean Energy Policy Conference

September 2024

Abstract

Recent United States (US) clean energy policies incorporate distributional mechanisms, such as place-based or means-tested subsidies. One prominent place-based example is the Inflation Reduction Act (IRA) Energy Community Bonus. The Bonus is designed to steer clean electricity investment to Energy Communities, areas historically reliant on the fossil fuel industry for employment and wages. I first ask whether the Bonus is working as intended to shift the location of clean electricity generators. New electricity generators must go through an interconnection queue before connecting to the national grid, so I measure how the Bonus is affecting electricity project queuing. The Energy Community Bonus increases queue entry, lowers queue withdrawal, and increases the number of queued projects becoming operational in Energy Communities. I next assess the Bonus' emissions effect using projections of electricity emissions in the short and long run. In the near term, the Energy Community Bonus reduces the environmental benefit of new clean energy. This is because Energy Communities have cleaner average marginal electricity emissions, so clean energy in Energy Communities offsets less dirty energy than clean energy elsewhere in the US. However, this correlation reverses in the future, and after 2030 the place-based policy improves the environmental benefits of clean energy.

1 Introduction

The Energy Community Bonus is a place-based subsidy for clean electricity in the United States (US), enacted in 2023 as part of the Inflation Reduction Act (IRA). It offers tax credits for new clean electricity in "Energy Communities," areas historically reliant on fossil fuels for employment and tax revenues.

I ask whether the Bonus is working as intended to shift new clean electricity generators to Energy Communities. It is an open question whether the subsidy is sufficient to change the location of clean electricity generators. The answer also has implications for other US clean electricity subsidies. Many US clean electricity subsidies are federal-level, presenting a challenge to identify how subsidies affects clean electricity generation. How clean energy supply responds to the Energy Community Bonus offers insight into how it responds to subsidies generally.

The answer is also interesting because, should the Bonus have an effect on clean electricity generator locations, this new variation in subsidy levels presents a research opportunity.

Subsidy-induced variation in US clean electricity locations would offer a novel instrument to identify how clean energy affects other local economic, environmental, and political outcomes.

A challenge to measuring the Energy Community Bonus effect is that it may be too early to directly measure the effect of the Energy Community Bonus on existing electricity generation. There is a substantial time lag between planning a new electricity generator and connecting it the national electricity grid because projects must go through an "interconnection queue." The median lag time between interconnection request and commercial electricity generation for a project built in 2023 was five years (Rand et al., 2024). The IRA and Energy Community Bonus were announced in 2022, about two years ago, substantially less than the five years of queue time for a typical generator.

However, the Bonus may already be shifting the location of proposed generators on the interconnection queue. I therefore use interconnection queue data to measure the effect of the Energy Community Bonus. I use two methods to estimate whether plants requesting grid interconnection are shifting to Energy Communities, both relying on plausibly exogenous distinctions between Energy Communities and other US locations. The empirical analyses of the queue show that in Energy Communities projects are entering the queue at a faster rate, withdrawing from the queue at a slower rate, and completing the queue at a faster rate, all of which imply the Bonus will cause more clean electricity generation in Energy Communities in the future.

I then ask whether shifting clean energy to Energy Communities will increase or decrease the environmental benefits of clean energy subsidies relative to spatially constant subsidies. The answer is not obvious. It is possible that the Energy Community Bonus subsidizes clean electricity in the areas where clean generation has the greatest environmental benefits, but the reverse may also be true, such that the subsidy targets areas with worse than average environmental benefits. Note, this is a question of policy efficiency. Unlike carbon taxes, second-best subsidy mechanisms do not inherently equalize and minimize abatement costs. The subsidy creates deadweight loss (DWL) in the clean energy market, with knock-on effects of displacing dirty energy for a social gain. Quantifying the policy's environmental benefit per subsidy dollar is equivalent to quantifying (the inverse of) its DWL per social gain. In this sense, this analysis answers both a question of whether there is a trade-off or complementarity between efficiency, equity, and environmental goals.

To answer this question, I use the National Renewable Energy Laboratory (NREL) Cambium model to explore how electricity emissions vary over space, and how Energy Community locations correlate with electricity emissions. Electricity generation tends to be cleaner near Energy Communities, so in the near term the Bonus will shift clean generation to locations where it has less benefit, a trade-off between efficiency and environmental outcomes and equity. However, after about 2030, Cambium predicts that result will reverse. Electricity generation near Energy Communities is projected to be dirtier than elsewhere in the US, suggesting that steering clean generation to Energy Communities increases the subsidy's environmental benefit.

2 Background

2.1 The Energy Community Bonus

The Inflation Reduction Act of 2022 (IRA) offers two nationwide tax credits for new clean electricity, and the Energy Community Bonus increases the tax credits' values. The Investment Tax Credit (ITC) subsidizes 30 percent of construction costs, and the Production Tax Credit (PTC) subsidizes 10 years' worth of electricity generation at $27.50 per megawatt-hour (MWh) in 2022, indexed to rise with inflation.[1] The credits cannot be combined, so clean electricity developers must choose which to claim.

The Energy Community Bonus increases the ITC by 10 percentage points and the PTC by 10 percent ($2.75 in 2022) for projects located in eligible areas, so-called Energy Communities. Energy generation technologies that are eligible for the Energy Community Bonus are the same as those eligible for the ITC or the PTC. Technologies that are eligible for both the ITC and PTC include solar, wind, municipal solid waste, and geothermal electric. Only eligible for the ITC are energy storage, fuel cells, combined heat and power, microturbines, microgrid controllers, and interconnection property. Only eligible for the PTC are hydroelectric, biomass, and landfill gas.

Energy Communities are defined according to three criteria described in Table 1. The first is brownfield sites. These are parcels of land that are underutilized due to pollution, mine scars, or other residual effects of former industrial use. The second category is metropolitan statistical areas (MSAs) or non-metropolitan statistical areas (non-MSAs) that meet criteria based on fossil fuel industry employment or local tax revenues at any time after December 31, 2009, plus a criterion based on local unemployment rate in the previous year. The final Energy Community category is census tracts in which a coal mine closed after December 31, 1999, or a coal-fired electricity generating unit retired after December 31, 2009.[2]

The Energy Community Bonus was announced in August 2022 and available in January 2023. However, precise definitions of Energy Communities were not available until April 2023, when the US Treasury released guidance and public data that define Energy Communities. This is especially relevant for the MSA and non-MSA Energy Community definition, where it was ambiguous which datasets and NAICS codes would define fossil fuel employment before the Treasury guidance. Indeed, there is still no definition of a dataset that reports local tax revenues from fossil fuel industries, so in effect this provision is not enacted.

Figure 1 maps the known Energy Communities for tax year 2023, based on the fossil fuel employment, coal mine, and coal-fired generator closure definitions. The map does not show the parcel-level brownfields because there is no publicly available data that exhaustively defines them (IWG, 2024). Under these definitions, approximately 38 percent of the land area in the contiguous US is an Energy Community, with the largest concentrations in the Southwest, the Rocky Mountains, and Appalachia. The map omits brownfield sites because no single data source lists sites that qualify.[3] It also omits the areas that would qualify according to their

[1]The PTC value reported assumes the project meets wage and apprenticeship requirements for 10 years, and does not meet the domestic content requirement. The baseline credit value is $5.50 per megawatt-hour (MWh), multiplied by 5 for meeting wage and apprenticeship requirements, with an additional 10 percent for domestically produced construction materials and finished goods.

[2]MSAs and non-MSAs are groups of counties. Counties are groups of census tracts. So, the greatest common factor of non-brownfield Energy Communities are census tracts.

[3]The EPA produces data on brownfield sites eligible for federal Brownfield funding, but according to IWG

<table>
<tr><td rowspan="4">(1) Brownfield site</td><td rowspan="4">or</td><td colspan="3">(2) Metropolitan statistical area or non-metropolitan statistical area (MSA or non-MSA) with…</td><td rowspan="4">or</td><td colspan="3">(3) Census tract with…</td></tr>
<tr><td>0.17% or greater direct employment related to the extraction, processing, transport, or storage of coal, oil, or natural gas at any time after Dec. 31, 2009</td><td rowspan="3">and</td><td rowspan="3">Unemployment rate at or above national average unemployment rate for the previous year</td><td rowspan="3">A coal mine has closed after Dec. 31, 1999, or a census tract which is directly adjoining to such census tract</td><td rowspan="3">and</td><td rowspan="3">A coal-fired electric generating unit has been retired after Dec. 31, 2009, or a census tract which is directly adjoining to such census tract</td></tr>
<tr><td>or</td></tr>
<tr><td>25% or greater local tax revenues related to the extraction, processing, transport, or storage of coal, oil, or natural gas at any time after Dec. 31, 2009</td></tr>
</table>

Table 1: Definitions of an Energy Community. Although the tax revenue definition is in the IRA legislation, there is currently no guidance on what data to use to quantify tax revenues, so in effect this provision is not enacted. Table formatting and information sourced from IWG (2024).

local tax revenues because currently no Treasury guidance describes data that could identify locations meeting this qualification.

Table 2 sketches the value of the Energy Community Bonus for a typical solar or wind electricity plant. Columns (1) and (2) show the annual value of the Energy Community Bonus for a typical plant claiming the PTC.[4] Column (1) shows the mean and median monthly generation of solar and wind plants in December 2022. Column (2) multiplies this value by the PTC value (assuming the entire plant qualifies and meets wage and apprenticeship standards), by 12 months, and by 10 percent to get the Energy Community Bonus value. For the mean solar plant, the bonus could be worth an extra \$80 thousand per year, in addition to the standard production tax credit, and for the mean wind plant nearly \$1 million per year. The PTC lasts

(2024), "not all properties that qualify as brownfield sites eligible for federal Brownfield funding are included in the definition of brownfield site for purposes of the energy community bonus credit."

[4]The baseline PTC in 2022 is \$5.50 per megawatt-hour (MWh), multiplied by 5 for facilities meeting wage and apprenticeship employment standards, by 1.1 for facilities meeting domestic content requirements, and by 1.1 for locating in an Energy Community. Table 2 assumes the wage and apprenticeship standards are met, but not the domestic content requirements. So the Bonus value in the table is calculated as \$5.50 $\times$ 5 $\times$ 0.1=\$2.75 per MWh.

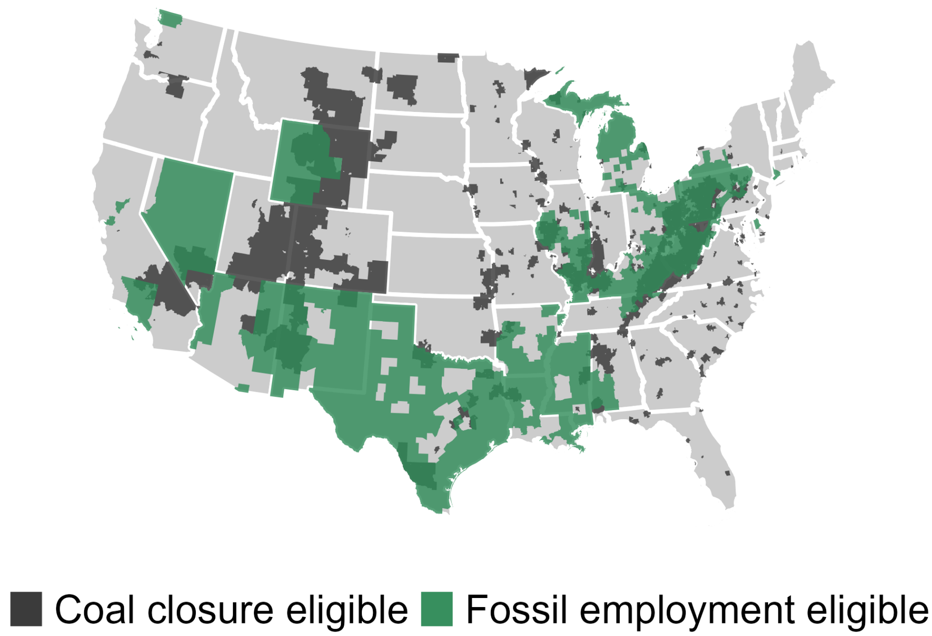

Figure 1: Energy Community census tracts, MSAs, and non-MSAs in tax year 2023. This map uses definitions based on fossil fuel direct employment, coal mine closures, and coal electricity generator closures. Appendix figure A.1 shows how the map changed in 2024.

for 10 years and adjusts with inflation, so the present value of locating in an Energy Community could be worth millions of dollars to prospective solar and wind developers. In 2022, the average wholesale price of electricity for a subset of price hubs available via EIA (2022) was $99.[5] So, the $2.75 per MWh Energy Community Bonus increases electricity generation revenue by about 3 percent.

Columns (3) and (4) of Table 2 show that for plants claiming the ITC, locating in an Energy Community is similarly valuable. Column (3) shows mean and median solar and wind plant capacities, and column (4) multiplies the capacity by the average construction cost per MW for solar and wind plants and by 10 percent. The mean solar plant would receive nearly $2.5 million in extra project subsidies for locating in an Energy Community, and an average wind plant an extra $16 million.[6]

2.2 Grid Interconnection Queues

Each balancing authority in the US manages its own grid interconnection requests, with its own interconnection queue. This has historically occurred on a rolling, first-come first-

[5] Authors' calculations of the volume-weighted average price of all wholesale electricity trades in 2022. The subset of price hubs is determined by those made available by the Intercontinental Exchange. This includes hubs located in CAISO, ERCOT, MISO, ISO-NE, PJM, and some western non-ISOs.

[6] A caveat in interpreting these magnitudes is that the PTC and ITC subsidize generators, not plants. Plants typically contain many generators owned by a common entity. Plant expansions to add new generators are eligible for the PTC and ITC.

		Production Tax Credit		Investment Tax Credit	
		Monthly Net Generation 2022 (MWh/Month)	***Annual Bonus***	***Plant Capacity 2022 (MW)***	***Total Bonus***
		(1)	(2)	(3)	(4)
Solar	Mean	2,442	$80,576	15	$2,408,000
	Median	382	$12,623	3	$497,000
Wind	Mean	28,393	$936,954	108	$16,126,000
	Median	17,692	$583,828	82	$12,239,000

Table 2: The value of the Energy Community Bonus for mean and median solar and wind plants. Monthly net generation is calculated from EIA 923 forms for December 2022. Net summer capacities of are calculated from EIA 860 forms for December 2022. The ITC estimates use construction cost estimates from EIA (2022).

serve basis. In recent years, to improve queue processing speed, balancing authorities have switched to a clustered, first-ready first-served basis. In 2023, the Federal Energy Regulatory Commission (FERC) mandated balancing authorities to implement these processes to reduce queue wait times (FERC, 2023).

The timing of this mandate coincides with IRA implementation. It would be problematic to compare pre- and post-IRA data in evaluating Energy Communities if the FERC order affects Energy Communities differently than it affects other locations. I argue that the changes to the queue process are substantial, but none should differently affect Energy Communities.

The life cycle of a new electricity generation project starts with a developer making a grid interconnection request to a balancing authority, which places the project in its interconnection queue. The balancing authority conducts a series of technical studies to assess the feasibility and impact of the new generator's connection to the grid, as well as its cost of interconnection. If a project is deemed feasible, then the balancing authority and project negotiate interconnection cost responsibility, to what extent it falls on the project developer and utility. This is formalized in an interconnection agreement, in which a project learns its interconnection costs. Should a developer choose to accept its interconnection costs, it signs the interconnection agreement, connects to the grid, and begins commercial operation (Rand et al., 2024). A project may begin construction at any stage of this process.

2.3 The Energy Community Bonus and Interconnection Queues

Note, the Energy Community definition based on fossil fuel employment and unemployment rates implies that Energy Community locations change year to year as local and national unemployment rates change. A project can claim permanent Energy Community status—all 10 years of of its PTC or whenever it is placed in service for the ITC—if it *begins construction* while located in an Energy Community (Treasury, 2023a). Beginning construction may mean physical work on a project or purchasing electricity generating equipment. More generally, beginning construction can also be defined as incurring 5 percent or more of a project's cost (Treasury, 2013).

High cost and uncertain interconnection costs cause a queue attrition rate of nearly 80 percent, so beginning construction before a project learns its interconnection costs is risky (Johnston et al., 2024; Rand et al., 2024). Therefore, the effect of MSA or non-MSA Energy Community designation on interconnection queue entry will be muted by the uncertainty of future Energy Community status. It should therefore be more likely to observe the effects of this Energy Community designation on queue withdrawal, interconnection agreement, and completion, as the Energy Community Bonus makes some marginal projects worthwhile for the developers as they learn their interconnection costs and wait on the queue.

On the other hand, Energy Communities defined by coal tracts and brownfields are permanently designated, as long as the Energy Community Bonus policy exists. These designations may offer similar incentives for queue entry, withdrawal, interconnection agreement signing, and completion.

3 Empirical Analysis of Energy Community Bonus Effected on Grid Interconnection Queues

I use two strategies to estimate the effect of the Energy Community Bonus on projects in the grid interconnection queue entry, withdrawal, and completion. The first strategy is a regression discontinuity (RD) design using the Energy Community definition at the MSA or non-MSA level, type (2) in Table 1, and the other is a difference-in-differences (DID) using the coal mine Energy Community definition, type (3) of Table 1. The remainder of this section describes more detail the data shared by both strategies, the RD empirical methods, results, and discussion, and the DID empirical method, results, and discussion.

3.1 Data

Interconnection queue data through December 2023 come from Queued Up project from the National Renewable Energy Laboratory (NREL) (Rand et al., 2024). They record projects that enter the interconnection queue, including their energy types, capacities, the date projects become operational if applicable, and county locations. The data track plants in the queue between 1995 and 2023. They cover all seven major Independent System Operators (ISOs) and Regional Transmission Organizations (RTOs) and 44 non-ISO utilities, which account for about 95 percent of US electricity generation. I use the full dataset to tally county queue entry. For queue withdrawal and queue completion, I restrict the sample to only the seven major ISOs and RTOs because many projects in the non-ISOs have a queue status of "withdrawn" or "operating" but do not have an associated withdrawal date, so it is impossible to tell if they withdrew before or after the Energy Community Bonus was active.

I supplement the NREL interconnection queue data with more recent data sourced directly from the major ISOs and RTOs, including the Electricity Reliability Council of Texas (ERCOT),[7] ISO-New England (ISO-NE),[8] Midwest ISO (MISO),[9] New York ISO (NYISO),[10] and

[7]https://www.ercot.com/mp/data-products/data-product-details?id=PG7-200-ER. ERCOT does not directly report a queue date, so queue date is assumed to be the first month where a project appears in its queue.

[8]https://irtt.iso-ne.com/reports/external

[9]https://www.misoenergy.org/planning/resource-utilization/GI_Queue/gi-interactive-queue/

[10]https://www.nyiso.com/interconnections

Southwest Power Pool (SPP).[11] California ISO (CAISO) is not available yet because it will not release its 2024 cluster requests (Cluster 15) until they are valiated, expected in 2025.[12] The Pennsylvania-New Jersey-Maryland (PJM) 2024 cluster study is not until late 2024. ERCOT data are released monthly without official dates of queue entry, withdrawal, and completion. I infer that the first month a project appears with one of those statuses is the month the project attained the status.

The outcome variables—queue entry, withdrawal, completion—only count projects eligible for the ITC and PTC, and thereby eligible for the Energy Community Bonus. This drops queued dirty technologies from the analysis. When measuring the count of projects, I include batteries, other energy storage, biofuel, biogas, biomass, fuel cells, geothermal, hydro, landfill energy, nuclear, wind, and solar. However, a challenge arises measuring the MW capacity of projects because many projects combine renewable generation with battery storage, and the queue data often does not record the MW capacity for these projects' batteries. Therefore, the MW capacity outcome variables only count renewable technologies, including geothermal, hydro, solar, and wind.

I also use data to identify Energy Communities and suitable controls based on the Energy Community criteria Table 1. Fossil fuel employment thresholds are defined using US Census County Business Patterns (CBP) for 2010 to 2021 (Census).[13] Local unemployment rates are calculated using the Bureau of Labor Statistics (BLS) Local Area Unemployment Statistics (LAUS) for years 2022 and 2023, to define Energy Communities in years 2023 and 2024, respectively (BLS). I identify census tracts with coal mine closures and coal generator closures using the official shapefiles from the US Department of Energy (DOE) for tax years 2023 and 2024 (DOE et al., 2023, 2024).[14]

For control variables, I use demographics from the 2021 Census American Community Survey and average wages from the 2022 Quarterly Census of Employment and Wages. I calculate county average and maximum direct normal irradiance (DNI, or sunniness) from the National Solar Radiation Database (NSRDB) and wind speeds at 100m from the WIND Toolkit (Draxl et al., 2015; Sengupta et al., 2018).

Further detail on data cleaning are in Appendix B.

3.2 Regression Discontinuity for MSA and Non-MSA Energy Communities

3.2.1 Regression Discontinuity Empirical Methods

An area is an Energy Community if it has 0.17% or greater direct employment related to fossil fuels at any time after 2009 *and* an unemployment rate at or above the national unemployment rate in the previous year. I use an RD to compare counties that are just over

[11] https://opsportal.spp.org/Studies/GIActive

[12] https://www.caiso.com/generation-transmission/generation/generator-interconnection

[13] Energy Communities are defined based on publicly available data, and 2021 was the most recent data at the time of analysis, so this is the data used in the official definitions.

[14] I use the official DOE dataset for coal tracts for convenience and to guarantee alignment official definitions. There are also publicly available datasets that identify the MSA and non-MSA Energy Communities, but the empirical strategies for these Energy Communities require the underlying data. I use the official dataset just to confirm I accurately characterize the MSA and non-MSA Energy Communities.

both thresholds to counties that just miss one or both thresholds. The assumption is that Energy Community assignment is as-if random near the threshold.

This setting deviates from a typical RD because there are two separate thresholds, and two separate running variables, and Energy Communities are defined as the intersection of clearing both thresholds. I use an estimating equation that accounts for both thresholds in one regression.[15] The estimating equation is specified as

$$
\begin{aligned}
Y_{ct} = & \alpha_0 + \beta_0 D_{1m}(1-D_{2m}) + \beta_1 D_{2m}(1-D_{1m}) + \beta_2 D_{1m} D_{2m} \\
& + MaxFFE_m(\beta_3 + \beta_4 D_{1m}(1-D_{2m}) + \beta_4 D_{2m}(1-D_{1m}) + \beta_6 D_{1m} D_{2m}) \\
& + Unemp_{my}(\beta_7 + \beta_8 D_{1m}(1-D_{2m}) + \beta_9 D_{2m}(1-D_{1m}) + \beta_{10} D_{1m} D_{2m}) + \varepsilon_{ct}
\end{aligned}
\tag{1}
$$

Y_{ct} is the outcome variable in county c and month t, where c is in MSA or non-MSA m. $MaxFFE_m$ is the maximum percent of the workforce employed in fossil fuels between 2010 and 2021 of the MSA or non-MSA m that contains c, denoted with an asterisk to indicate it is re-centered around 0.17 percent, and $D_{1m} = 1(MaxFFE_m \geq 0)$. $Unemp_{my}$ is the unemployment rate of MSA or non-MSA m, re-centered around the previous year's national unemployment rate, and $D_{2m} = 1(Unemp_{my} \geq 0)$.[16][17]

The equation estimates threshold jumps and trends for four groups of counties: The coefficients on $D_{1m}(1-D_{2m})$ estimate the jumps and trends for the group of counties that meet the fossil fuel employment cutoff but just miss the unemployment cutoff; coefficients on $D_{2m}(1-D_{1m})$ estimate the same for counties that just misses the fossil fuel employment cutoff and meets the unemployment cutoff; the uninteracted coefficients estimate for counties that just miss both cutoffs; and coefficients on $D_{1m}D_{2m}$ for Energy Community counties. The conditional local average treatment effect (LATE) of the Energy Community Bonus is β_2.

Standard errors are clustered at the county and balancing authority levels. I cluster at the county-level because queuing in counties in previous periods may affect queuing in counties in future periods. I cluster at the balancing-authority-level because there is likely correlated queue entry in balancing authorities as they process clusters.

I trim the data to a local sample within a bandwidth around the threshold using the data-driven procedure in Calonico et al. (2014, 2020), with a equal size on each side and clustered at the county level. The procedure requires only one running variable. I define a joint running variable, $MaxDist_{ct} = \arg\max_{FFemp_m, Unemp_{my}} |x|$. Recalling these variables are re-centered around the threshold, this will take the value farthest from the threshold. $MaxDist_{ct}$ is greater

[15] Other ways to account for the intersection of thresholds include estimating two separate regression equations, but this will either yield two different treatment effects. This could be a use case for a seemingly unrelated regression, but this is more complicated. I could combine the variables taking the maximum distance from either threshold, but the variables are measured in different units and any combination would be ad hoc. I do, however, use this maximum procedure to generate the data-driven bandwidths, as described below.

[16] Based on the official DOE Energy Community maps, the unemployment rates used to define this threshold seem to be rounded to the nearest tenth. For example, the national unemployment rate in 2022 was 3.66, but MSAs or non-MSAs qualify with unemployment rates of 3.65. Therefore, I set the effective cutoff to be 0.0005 less than the national unemployment rate rounded to the nearest tenth.

[17] Note that, due to these time-dependent Energy Community definitions, a county may exceed both thresholds and be a treatment unit in one time period, and miss a threshold and be a control unit in another time period. This is a concern if there is serial correlation in queuing at the county level. For example, if a county is an Energy Community in time period t and not in $t+1$, and the Bonus causes more queuing in the county in time period t, then serial correlation implies the county is not a suitable control in time period $t+1$. Future versions of this paper may consider an estimation procedure to account for this.

than zero for Energy Communities, less than zero for counties that miss both thresholds, and equal to whichever value is farthest from the threshold for counties that exceed only one threshold. I winnow the sample to only observations with the absolute value of both variables less than the bandwidth. The bandwidth-generating local linear regression procedure also depends on the outcome variable, so each outcome variable has a distinct bandwidth. The resulting bandwidth will be conservatively large and is necessarily *ad hoc*, so I run bandwidth sensitivity analyses.

The estimating sample is also trimmed to include only observations from after April 2023, when the Treasury guidance was released and the definition of the treatment areas was first clear. The sample excludes counties that contain Energy Community census tracts defined by coal closures, so the control group does not have any Energy Communities (other than brownfields), and the treatment group is similar to the controls in that respect.

3.2.2 Regression Discontinuity Results

Figure 2 shows the all of the possible Energy Community counties and control counties for analysis. The counties are colored based on their MSA and non-MSA's Energy Community status in 2023 and 2024. The gray counties are dropped from analysis because they include a census tract that is eligible for the Bonus for containing a coal mine or coal generator.

Table 3 compares descriptive statistics for the counties selected for analysis. The first three columns describe Energy Community counties in 2023, 2024, and both years, respectively, and the final column shows counties that may be used as controls. Over 600 counties were Energy Communities as defined by the MSA/non-MSA thresholds for at least one year in 2023 and 2024, and there are nearly 1,700 control counties. The Energy Community counties tend to be sunnier and windier than elsewhere in the US, but there is still substantial overlap between the groups. There is also overlap between the Energy Communities and non-Energy Communities for populations and average weekly wages. The Energy Community counties have higher fossil fuel employments, based on both the 2023 and 2024 definitions, and higher unemployment rates in all years.

The table also reports the mean monthly outcome variables for May 2023 to June 2024. The Energy Community counties tend to have more queued projects and MW since the Energy Community guidance passed than the full set of non-Energy Community counties. Energy Community counties have a similar number of withdrawn projects and MW. They have more newly operating projects and MW.

Turning to the regression discontinuity, Figure 3 shows the final set of counties selected for each of the optimal bandwidth regression specifications. Most counties are dropped by the bandwidth selection. The withdrawal and operating analyses also drop counties that are not in a major ISO or RTO. Appendix Figure A.2 shows which counties fall into balancing authorities based on location and interconnection queue data.

Table 4 shows the optimal bandwidth results and regression discontinuity coefficients for equation 1, where the six columns show six different outcome variables.[18] Projects are entering

[18]I also estimate specifications equations without the balancing authority fixed effect, with more control variables, as a regression difference-in-discontinuity (RDID) using pre-IRA data, and with a quadratic specification. Queue entry results are insignificant for all specifications. Withdrawal RD results are robustly negative and significant, except using a quadratic specification with covariates. Withdrawal RDID results are significant only under the linear, non-covariate specification. Completion rate results are insignificantly negative for all

	Energy Community 2023 only	Energy Community 2024 only	Energy Community 2023-24	Non-Energy Community
Population	18,292	332,343	175,396	126,664
	(35,102)	(690,461)	(988,085)	(314,741)
Avg. weekly wage (2021$)	1031	1078	944	972
	(240)	(354)	(190)	(226)
DNI (kWh/m^2)	6.26	5.35	5.07	4.95
	(0.76)	(0.79)	(0.86)	(0.66)
Wind speed at 100m (m/s)	7.17	6.55	6.64	6.76
	(0.8)	(0.88)	(0.67)	(0.83)
Fossil employment (%), 2023 definition	11.41	0.96	2.95	0.86
	(5.59)	(1.51)	(3.11)	(2.73)
Fossil employment (%), 2024 definition	12.67	1.32	3.66	1.1
	(6.09)	(1.69)	(3.63)	(3.05)
Unemployment 2022 (%)	3.75	3.88	4.51	3.23
	(0.16)	(0.67)	(0.71)	(0.94)
Unemployment 2023 (%)	3.44	4.12	4.37	3.26
	(0.19)	(0.66)	(0.63)	(1.03)
Queued projects	0.077	0.044	0.075	0.021
	(0.37)	(0.26)	(0.39)	(0.21)
Queued MW	14	3.2	8.2	2.2
	(97)	(35)	(65)	(35)
Withdrawn projects	0.012	0.024	0.014	0.02
	(0.11)	(0.19)	(0.15)	(0.29)
Withdrawn MW	1.3	2	1.6	1.5
	(21)	(21)	(26)	(28)
Operating projects	0.017	0.0085	0.0046	0.0035
	(0.14)	(0.11)	(0.075)	(0.088)
Operating MW	2.4	1.2	0.6	0.32
	(29)	(20)	(11)	(12)
Number of counties	104	109	404	1,697
County-month observations	3,120	3,270	12,120	50,910

Table 3: Descriptive statistics for counties in Figure 2. DNI is direct normal irradiance, or sunniness. Values are means, with standard deviations in parentheses. Measures of project counts include all projects eligible for the ITC and PTC. Queued, withdrawn, and operating MW include only renewable generators.

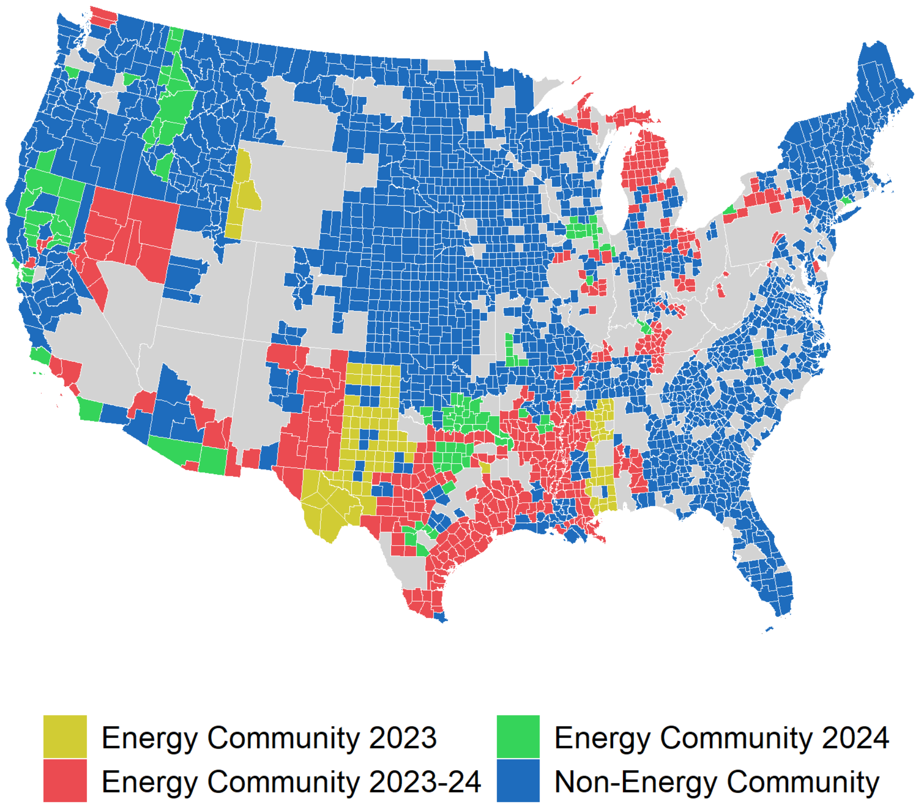

Figure 2: All possible treatment and control counties in each year, before trimming using bandwidth. The counties in gray are excluded from the analysis because they include coal census tracts.

the connection queue at an insignificantly slower rate in Energy Communities, measured by both the number of eligible projects and renewable capacity. Projects are withdrawing at a significantly slower rate in Energy Communities measured in MW, which aligns with the expected Bonus effect, but the withdrawal result measured in projects is only significant with a 90 percent confidence interval. Projects complete the queue and begin operating at an insignificantly lower rate in Energy Communities.

A concern is that the results will be driven by spurious correlations between Energy Community prevalence and the number of clusters the major ISOs and RTOs have processed in the post-period. I therefore also run specifications with balancing authority fixed effects, but the results are qualitatively unchanged. Including control variables both also do not change the qualitative result. Specifying the estimating equation as a quadratic equation makes all the estimated results null, including withdrawal. A concern is there may be some other variable, omitted or observed, that changes discontinuously at the thresholds. However, regression difference-in-discontinuity (RDID) specifications using pre-treatment data yield similar treatment effect results.

specifications.

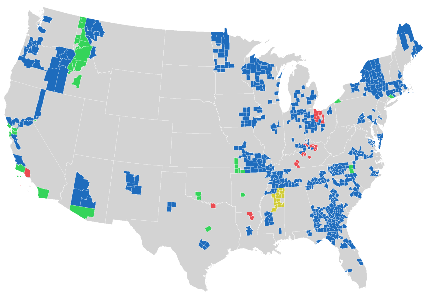

(a) Queue entry, number of projects

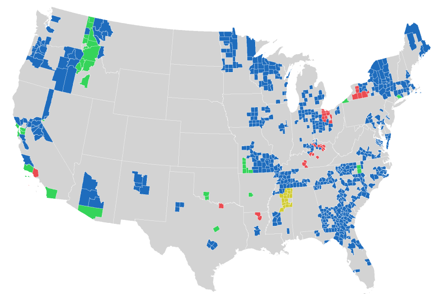

(b) Queue entry, MW

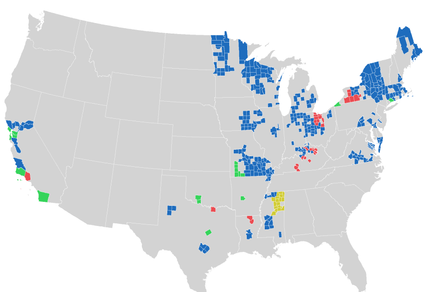

(c) Withdrawal, number of projects

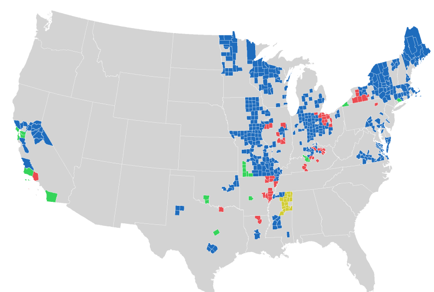

(d) Withdrawal, MW

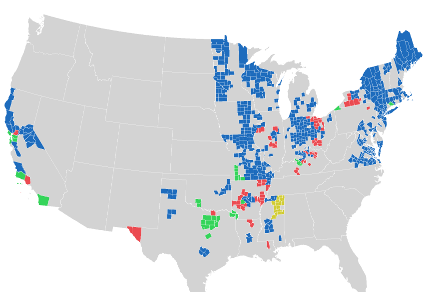

(e) Queue completion, number of projects

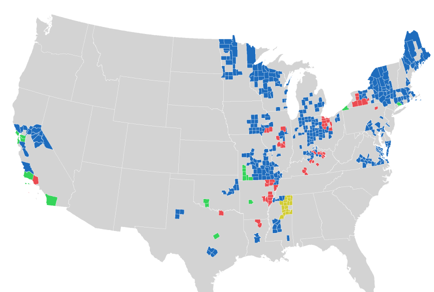

(f) Queue completion, MW

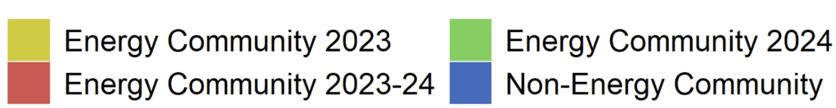

Figure 3: Counties included in the preferred regression specifications.

	Queued		*Withdrawn*		*Operating*	
	Projects	MW	Projects	MW	Projects	MW
RD Estimate	−0.012	0.276	−0.178	−8.498**	−0.011	−2.113
	(0.033)	(2.951)	(0.115)	(4.140)	(0.007)	(2.064)
Bandwidth	0.00406	0.00442	0.00422	0.00599	0.00692	0.00489
Observations	6,772	7,140	3,476	5,432	7,026	4,410
Note:						*p<0.1; **p<0.05; ***p<0.01

Table 4: RD results, which are estimates of β_2 in equation 1. The first two columns show the estimated LATE of the Bonus on queued projects and capacity, the next two show queue withdrawals, and the final two show queue completions. Bandwidths are the MSE-optimal bandwidths from Calonico et al. (2014), where the running variable is the maximum of the two running variables' distance from its respective threshold, clustered at the county level. Standard errors are clustered at the county and balancing authority levels.

Figure 4 shows how the results vary for each of the six regressions using different bandwidths. Whether the queue entry results is greater or less than zero is sensitive to the choice of the bandwidth, but none of the results are significant. The withdrawal rate result is robustly negative, but the significance depends on the bandwidth. The queue completion rate coefficients are robustly negative, but the results are insignificant.

The graphs in Figure 5 show the means of the six outcome variables, binned into quantiles. The quantiles bin the county-month observations by the minimum distance of the running variables above their respective thresholds. The vertical red line intersects the x-axis at the quantile with counties that have a minimum distance equal to zero, such that all Energy Communities are to the right of the line. There is a positive correlation between queue entry rate and the minimum distance from a threshold. Although it looks like there may be a jump at the thresholds, the RD results above show the jump net of trends is insignificant. Project withdrawal rates also seem to sharply decrease for Energy Communities, but the regressions show insignificance. As measured in MW, there is no clear change at the threshold. Project operating rates are noisy and show no clear patterns. Appendix Figure A.3 uses these same quantiles and vertical red lines, but plots only pre-treatment data.

3.2.3 Regression Discontinuity Discussion

This empirical analysis shows that the withdrawal rate is slower in Energy Communities compared to similar non-Energy Communities. The magnitude is large; lowering the withdrawal rate by 8 MW per county per month is four times the average withdrawal rate per county per month. To some extent, this is evidence that the Energy Community Bonus is working as intended. The graphs of queue entry in Figure 5 also tell a different story than the regression results. Panels (a) and (b) in that Figure show substantially higher queue entry in Energy Communities, but the local discontinuity analysis does not register this effect. Notably, the same trend on these quantiles is not observed in the pre-IRA data plotted in Appendix Figure

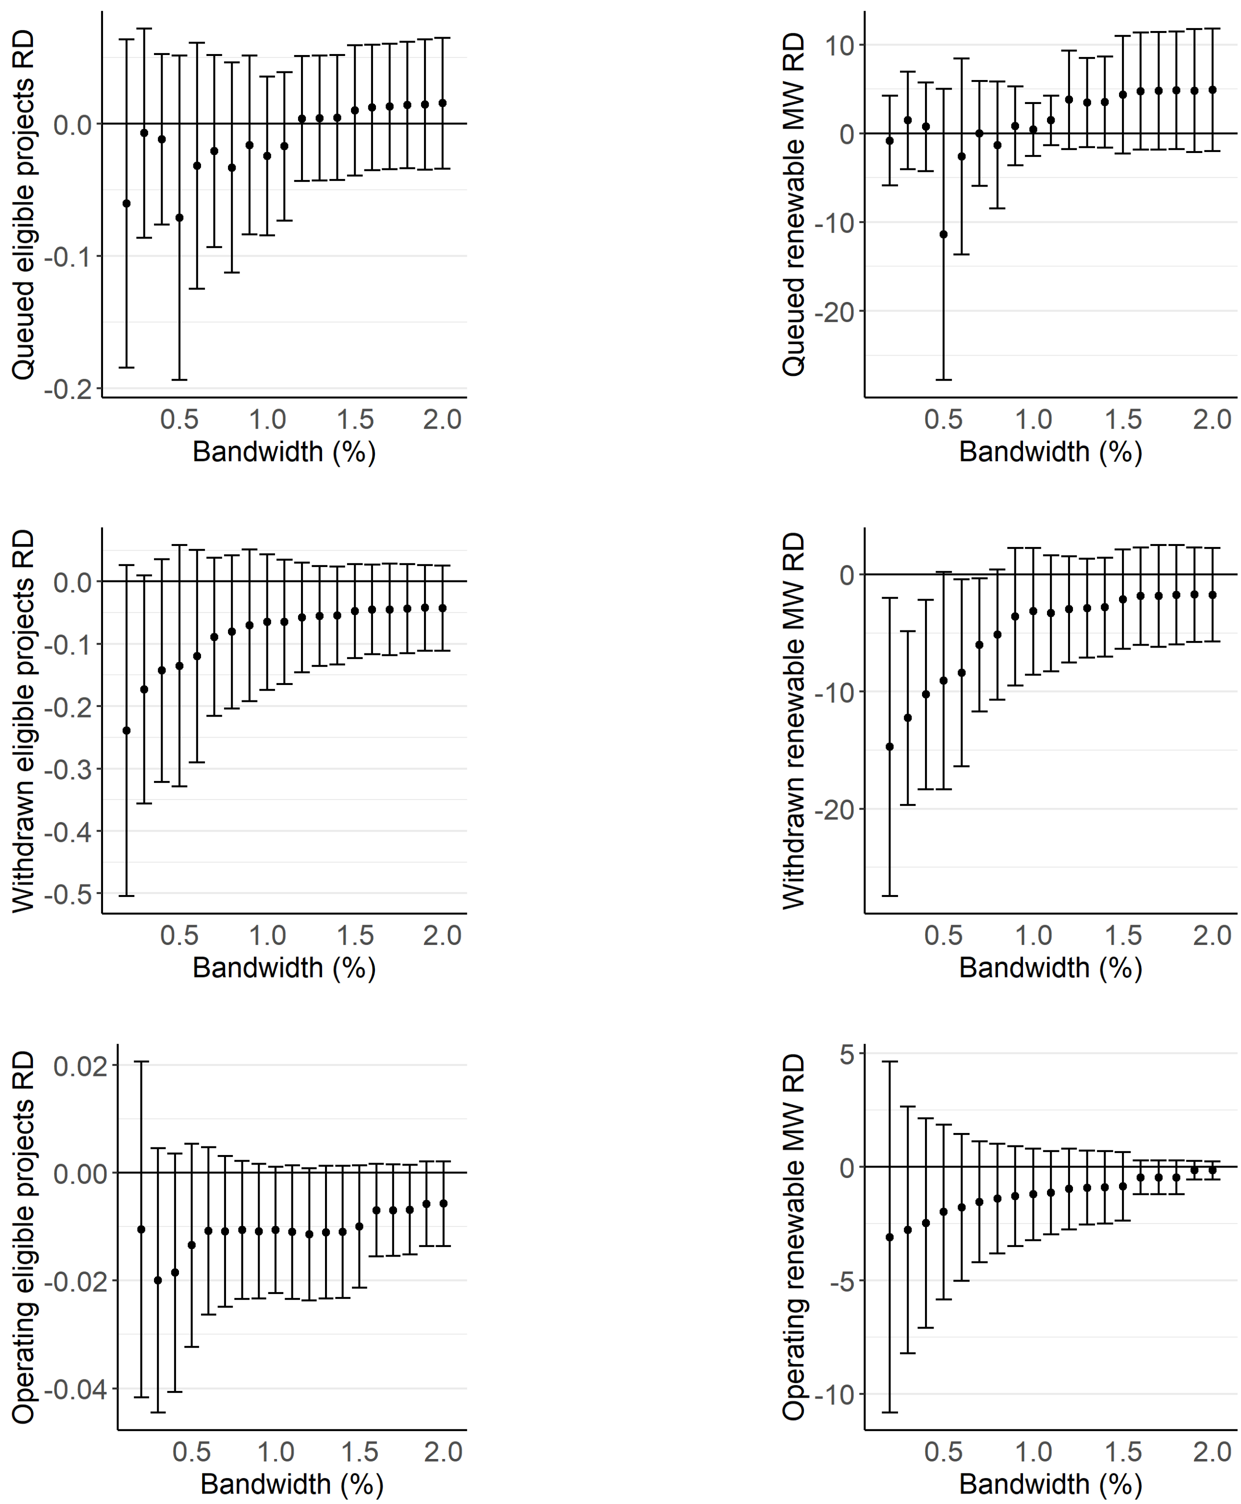

Figure 4: RD coefficients for the six regressions reported in Table 4, varying the bandwidths to select the samples. The included observations are those with both running variables within the bandwidth from their respective thresholds. The first row shows the change in eligible projects and renewable MW capacity entering the queue, respectively for the two columns. The second row shows the same measures of withdrawal, and the third row shows the same measures of new operations.

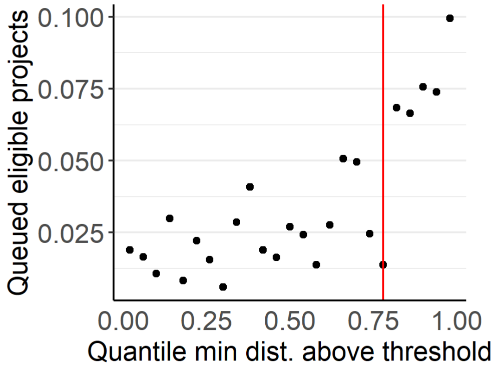

(a) Queue entry, number of projects

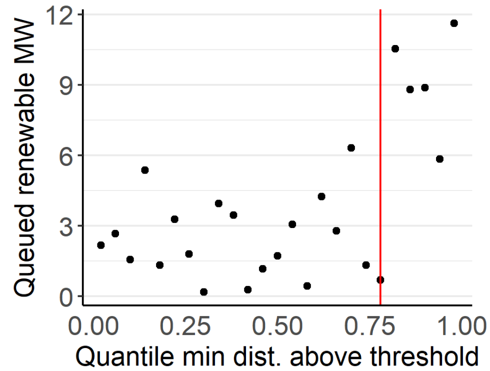

(b) Queue entry, MW

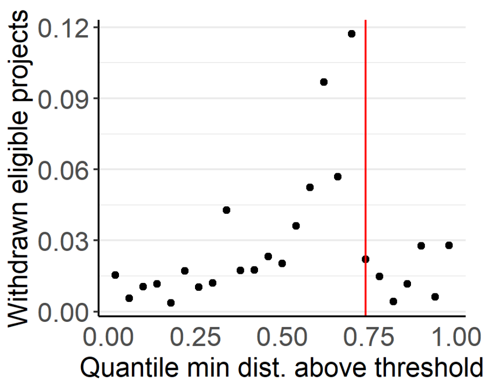

(c) Withdrawal, number of projects

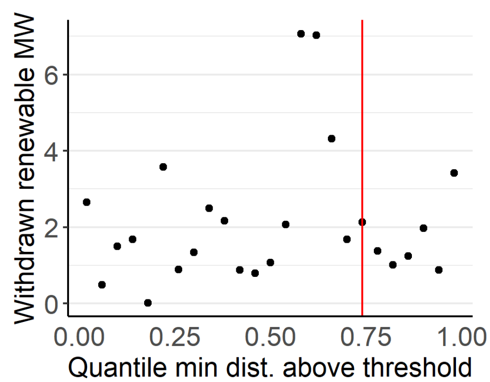

(d) Withdrawal, MW

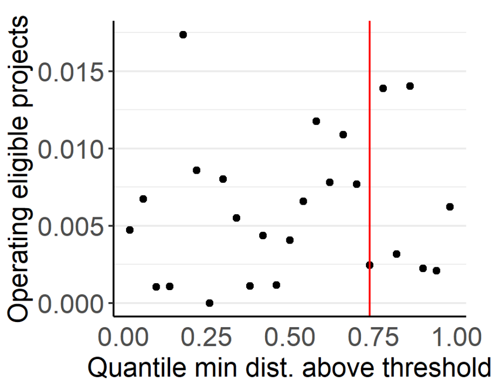

(e) Queue completion, number of projects

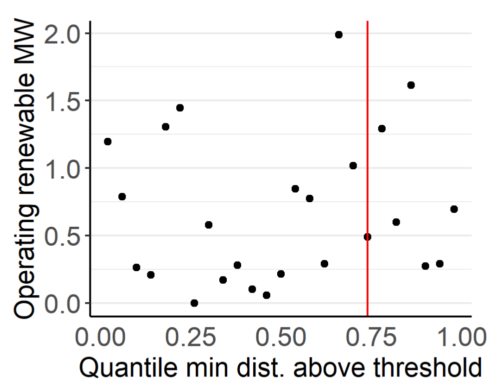

(f) Queue completion, MW

Figure 5: The x-axes show the quantile of the maximum difference between the fossil fuel employment or previous year unemployment and their respective thresholds. The vertical red line intersects the axis at the quantile that includes zero. All values below the red line are below one or both of the thresholds, and all values above the red line are above both thresholds. The figures differ by the outcome variables, where the outcomes are average monthly changes between May 2023 and June 2024.

A.3. However, any changes to queue entry and withdrawal have not yet converted into more projects completing the queue in these counties. Those results are null.

The results should be considered with their limitations. First, the guidance first defining the Energy Community MSAs and non-MSAs just released in April 2023, so the analysis only uses 14 months of post-treatment data through June 2024. Clean energy developers may still be adjusting to the new information. The analysis also does not recognize brownfield Energy Communities, and it is almost certain there are brownfield Energy Communities in the control group. The result is an unbiased LATE estimate of MSA or non-MSA Energy Community status, but it is biased toward a null result if interpreted as a general Energy Community Bonus effect. Finally, since the fossil fuel employment Energy Community definition changes year-to-year, project developers' responses to the policy are limited by the uncertainty that areas may lose Energy Community status in the future, before they begin construction. The next analysis uses the coal mine closure Energy Community definition, and is not subject to this limitation because it is a permanent Energy Community designation as long as the Bonus exists.

3.3 Difference-in-Differences for Coal Mine Energy Communities

3.3.1 Difference-in-Differences Empirical Methods

This identification strategy uses the Energy Community criterion Type (3) in Table 1, based on census tract coal mine or coal generator closures. Unlike the Energy Community criterion used for the RD, this eligibility criterion is not based on a threshold level in a continuous variable. Areas have either had a coal mine abandonment or they have not, a coal generator retirement or not. There is no possibility of identifying this third effect with an RD to separate the effect from confounding variables because there is no running variable.

Using this definition requires assuming the coal tract Energy Community designation is exogenous. It would be problem, then, if coal mine or coal generator closures were associated with local characteristics that were beneficial or adverse to development of clean energy plants. For this reason I do not use coal generator retirements. A coal-fired generator closure may itself cause new local wind and solar electricity generators to be built, confounding the Energy Community Bonus effect. Instead, I proceed only using the Energy Communities defined by coal mine abandonments, on the presumption that the presence or absence of local coal mines do not affect the profitability of nearby renewable energy generators. Evidence suggests this is a fair assumption.

Watson et al. (2023), show that coal mine closures in the U.S. have been predominantly caused by local characteristics unlikely to be correlated with suitability for wind or solar plants: increasing mining costs, declining natural gas prices, and stagnant US total electricity demand. Mine costs, by contrast, have risen due to decreasing productivity, and to a lesser extent wage increases. Labor costs are easy to observe and control for at the county level. And declining gas prices and stagnant electricity demand are national trends unlikely to affect county-level variation in wind and solar electricity development. Overall, there is good reason to believe that the Energy Community assignment due to coal mine abandonments can be treated as exogenous variation in Energy Community Bonus eligibility without worrying that the abandonments are correlated with clean electricity development.

However, using coal mine closures faces a data-related challenge in that these areas are defined at the census tract level, but the interconnection queue data are at the county level. Most counties are substantially larger than the census tracts they contain.[19] As a consequence, for this strategy I define treated counties as those with at least one census tract eligible for the Energy Community Bonus due to a coal mine closure.

My primary specification limits the estimating sample to these counties and their neighbors. I assume treated Energy Community counties are geographically similar to the neighboring counties with respect to suitability for renewable energy generation, land availability, and transmission availability. I also drop counties from the analysis that are eligible for the Energy Community Bonus by other criteria, either the fossil fuel employment or coal generator definition. I do so for the treatment group because of the concern that coal generator closures may affect an area's suitability for renewable generators, and for the control group to ensure they do not receive the Energy Community Bonus.[20]

The estimating equation is

$$Y_{ct} = \delta_{ry} + \delta_1 Post_t + \delta_2 CoalMineEC_c + \delta_3 Post_t \times CoalMineEC_c + \boldsymbol{X_c'}\boldsymbol{\zeta} + \varepsilon_{ct}, \qquad (2)$$

δ_{ry} is a region-by-year fixed effect to account for differences in cluster frequency in the pre- and post-periods, where $CoalMineEC_c$ is an indicator equal to one if county c contains a coal mine closure Energy Community and zero otherwise. $Post_t$ equals one if month t is in August 2022 or later. $\boldsymbol{X_c'}$ is a vector of controls, including local population, wages, incomes, total employment, and the maximum DNI and wind speed in the county.

Note the definition of post-policy differs between this analysis and the regression discontinuity above. Here I use post-IRA announcement in August 2022 as the treatment because it was clear where coal mines closed, and therefore clear where are coal mine Energy Communities. This clarity contrasts with the MSA and non-MSA fossil fuel employment definition, which depended on which datasets and NAICS codes would define fossil fuel employment.

Standard errors are clustered at the county level to account for serial correlation within counties. The primary specifications also use inverse propensity score weights to improve the treatment and control group balance. I estimate propensity scores with a logistic regression of a coal mine closure indicator on the number of households in the county, median and mean household income, average weekly wages, and average employment.

3.3.2 Difference-in-Differences Results

Figure 6 shows the counties selected for DID analyses. The queue entry analyses use the colored counties in the left panel, including counties from the non-ISO queues. The withdrawal and operating analyses use colored counties mapped in the right panel, which include only states manged by the major ISOs and RTOs. Dark blue counties contain a census tract that is an Energy Community due to a coal mine closure and have no other Energy Community designation. The green counties have no Energy Community designation and neighbor a dark blue coal mine Energy Community county.

[19] This is a non-issue for the RD because the eligibility criterion assigns an MSA or non-MSA Energy Community status, and these areas are larger than the counties they contain.

[20] I also consider using counties with coal mines that remain opened, or coal mines that closed before the 1999 cutoff date for eligibility. However, most of these counties have an Energy Community designation, and the resulting control sample is just a few counties.

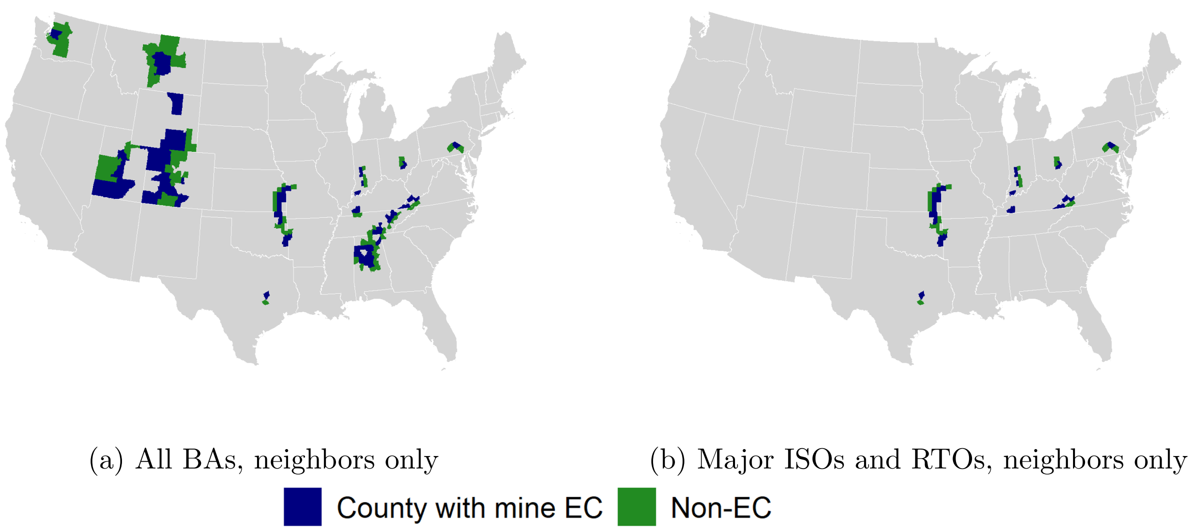

(a) All BAs, neighbors only

(b) Major ISOs and RTOs, neighbors only

Figure 6: Counties used for estimating equation 2. The dark blue counties contain census tracts that are Energy Communities due to a coal mine closure, do not contain census tracts that have a coal-fired generator closure, and did not have MSA or non-MSA Energy Community status in 2023 or 2024. The green counties have no Energy Community designation and neighbor the dark blue counties. The left panel shows all such counties in the US, used for the queue entry regressions. The right panel shows only counties in states managed by the major ISOs and RTOs, used for the withdrawal and completion analyses. BA stands for balancing authority.

Table 5 shows descriptive statistics for the treatment and control counties. These are the variables shown in this table are also the dependent variables used to generate inverse propensity score weights. The first two columns show statistics for the colored counties in the left panel of 6, and the second two columns show descriptive statistics for the right panel. The pre-IRA queue statistics are county monthly average changes. The Energy Communities and their neighbors have similar pre-IRA queue outcome variables, sunniness, and windiness. The Energy Communities have lower populations, employment, and incomes than the control counties, but have similar wages.

The propensity score logistic regression results are reported in Appendix Table A.1. There are separate regressions corresponding inverse propensity scores used to weight the regressions including and excluding counties outside states managed by major ISOs and RTOs. All the regressions show at least one variable is a significant correlate with likelihood of a mine closure. All variables in Table 5 are significant predictors of a mine closure.

Table 6 shows the DID estimates using the two control groups (the full result is in Appendix Table A.2). The regression uses balancing-authority-by-year fixed effects and covariates, with the full table of results including covariates reported in the Appendix. The DID coefficient is the change in queuing in coal mine counties before and after the IRA was announced, net of the change in control counties before and after the IRA, reported as monthly average changes. Projects are queuing in Energy Communities at a significantly higher rate at a 95 percent confidence interval, but the result measured in MW is only significant at a 90 percent confidence interval. Projects are withdrawing at a lower rate in Energy Communities based on a 95 percent confidence interval, but the result is insignificant measured in MW. Conversely, more MW of projects are completing the queue and becoming operational at a higher rate in Energy Communities, but the result is insignificant measured in the number of projects.

Figure 7 decomposes the DID estimate into annual effects, relative to 2021. Note that eight months of 2022 were before the IRA was public, and the 2024 data are incomplete and missing some major ISO clusters and the latter six months, so the 2023 data may be more reliable on its own. The queue results are insignificant, but the coefficients are all positive in the post-period with consistent magnitudes. The withdrawal graph measured in projects shows some pre-trend in 2021, since the pre-2021 DID coefficients are persistently positive. There is a substantial and significant dip in withdrawal in Energy Communities in 2023, but a return to trend in 2024. The withdrawal result measured in MW is null and shows some significant pre-trends, as well as an increase in withdrawal in 2022. The operating graphs show null DID estimates, with significant pre-trends in 2020.

3.3.3 Difference-in-Differences Discussion

The DID regression coefficients that are significant all suggest the Energy Community Bonus is working as intended by increasing the rate of queue entry, slowing queue withdrawal, and increasing queue completion in Energy Communities. The graphs of annual results are less clear, but also have less statistical power.[21]

This identification strategy has two main benefits over the RD. First, the results here are

[21] I also run the same regressions using all counties that are non-Energy Communities as controls, and corresponding propensity scores, instead of just neighbors. The coefficients have similar directions and significance levels, but are much larger in magnitude. The graphs equivalent to Figure 7 show substantial pre-trends, though, so the present analyses of neighbors is preferred.

	Coal Mine EC, All BA	*Neighbor non-EC, All BA*	*Coal Mine EC, Major ISO/RTO*	*Neighbor Non-EC, Major ISO/RTO*
Pre-IRA queued projects	0.01 (0.04)	0.01 (0.04)	0.04 (0.06)	0.05 (0.07)
Pre-IRA queued MW	1.15 (3.55)	0.83 (2.52)	3.22 (5.4)	2.88 (4.07)
Pre-IRA withdrawn projects	0.003 (0.01)	0.004 (0.01)	0.007 (0.02)	0.01 (0.02)
Pre-IRA withdrawn MW	0.16 (0.77)	0.17 (0.86)	0.45 (1.25)	0.59 (1.53)
Pre-IRA operating projects	0.002 (0.02)	0.001 (0.007)	0.005 (0.03)	0.005 (0.01)
Pre-IRA operating MW	0.08 (0.73)	0.21 (1.28)	0.22 (1.22)	0.73 (2.33)
DNI (kWh/m^2)	5.53 (0.98)	5.44 (0.85)	4.83 (0.35)	4.88 (0.37)
Wind speed at 100m (m/s)	8.48 (1.69)	8.68 (1.95)	7.77 (0.68)	7.91 (0.56)
Population	79,139 (257,822)	102,988 (174,134)	48,031 (61,416)	111,381 (157,746)
Avg. household income (2021$)	73,636 (19,091)	80,025 (21,536)	68,915 (13,837)	78,742 (17,221)
Avg. weekly wage (2021$)	936 (226)	953 (187)	924 (171)	921 (168)
Employment	39,396 (162,574)	43,222 (78,126)	20,495 (37,454)	49,163 (86,197)
Number of counties	87	83	31	24

Table 5: Descriptive statistics for counties with coal mine closures and those used as controls. Values are means, with standard deviations in parentheses. Queuing, withdrawals, and operating statistics are for monthly averages for January 2018 to July 2022. DNI is direct normal irradiance, or sunniness.

not contaminated by any uncertainty of future Energy Community status because the coal mine Energy Community assignment is permanent, unlike the MSA and non-MSA Energy Community assignment that can change year to year. Second, there is more data in the post-

	Queued		*Withdrawn*		*Operating*	
	Projects	MW	Projects	MW	Projects	MW
Post-IRA	−0.002	0.208	0.004	−0.500	−0.003	−0.636
	(0.005)	(0.696)	(0.006)	(0.408)	(0.002)	(0.415)
EC	−0.001	0.133	0.014*	−0.425	−0.001	−0.937*
	(0.004)	(0.360)	(0.007)	(0.471)	(0.005)	(0.568)
EC×Post	0.014**	1.288*	−0.045**	0.021	0.009	2.395**
	(0.007)	(0.747)	(0.020)	(0.530)	(0.007)	(1.194)
R^2	0.053	0.078	0.132	0.029	0.017	0.029
Obs.	13,260	13,260	4,290	4,290	4,290	4,290

Note: *p<0.1; **p<0.05; ***p<0.01

Table 6: Conditional average treatment effect estimates of Equation 2, δ_3. Regressions all include balancing-authority-by-year fixed effects and controls for county average income, average wages, population, total employment, DNI and windiness, and DNI and windiness squared. Standard errors in parentheses are clustered at the county level. The regressions are weighted with inverse propensity scores.

treatment period for this analysis because it was likely clear where abandoned coal mines were before the Treasury guidance named the official dataset to define mine closures. The RD analysis post-treatment period begins in May 2023, while the DID post-treatment period begins in September 2022.

On the other hand, a limitation is that the queue data are at the county level, while coal mine Energy Communities are defined at the smaller census tract level. I code counties with a coal mine census tract as fully an Energy Community, when in reality only a subset of the county may be an Energy Community. This assumption counts some counties that are only partially treated as fully treated, so the DID underestimates the effect of full county Energy Community status.

Moreover, like in the RD, brownfield Energy Community parcels are not identified. The control group is almost certainly contaminated with projects claiming the Energy Community Bonus as brownfields. So, like the RD estimate, the DID estimate biased toward zero and is most accurately interpreted as the effect of a county containing some coal mine Energy Community tract or tracts, not as the pure Energy Community Bonus effect.

County-level identification of tract-level treatment and brownfields are two reasons the estimator of the Energy Community Bonus effect is biased to null, but the DID shows significant effects on queue entry, withdrawal, and completion. The next section investigates how shifting future renewable generators to Energy Communities will affect emissions abatement.

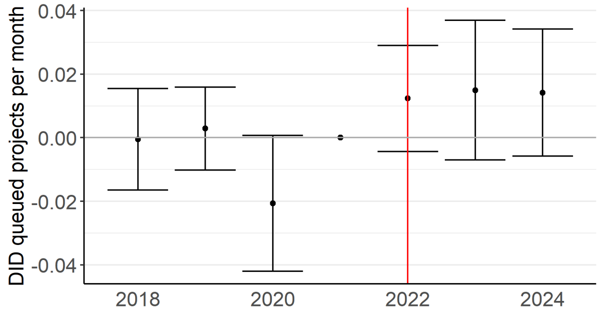

(a) Queue entry, number of projects

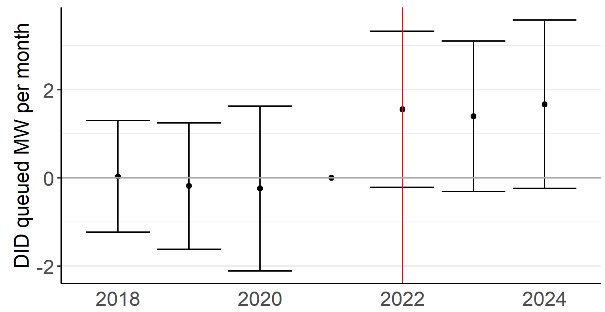

(b) Queue entry, MW

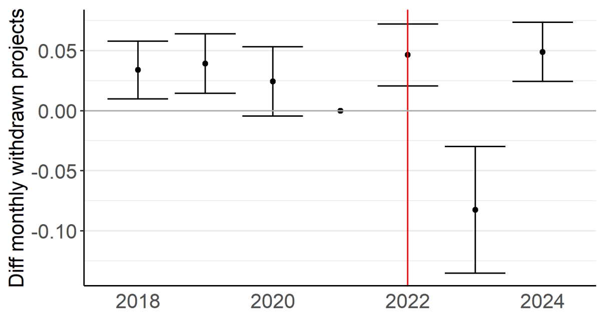

(c) Withdrawal, number of projects

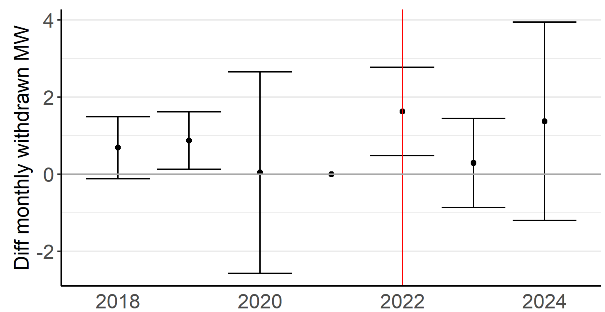

(d) Withdrawal, MW

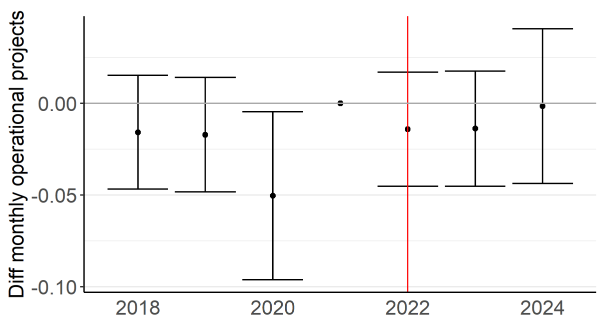

(e) Queue completion, number of projects

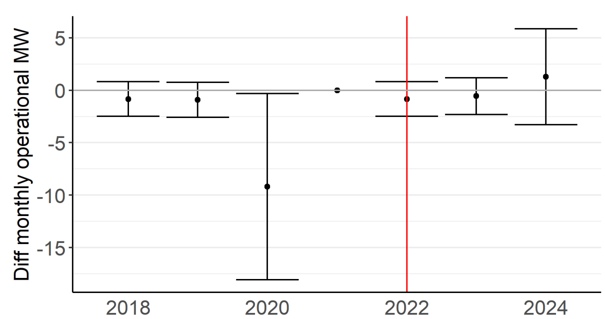

(f) Queue completion, MW

IRA Announced

Figure 7: Difference in outcome variables between Energy Communities and the neighbor control group by year, relative to 2021. Dots are point estimates and error bars are 95 percent confidence intervals. The regression used to generate the figure is the same as Equation 2, but replacing $Post_t$ with dummy variables for each year. The regressions are weighted by inverse propensity scores. Standard errors are clustered at the county level.

4 Emissions Effects of the Energy Community Bonus

Of course, clean electricity subsidized by the Energy Community Bonus will have no direct effect on emissions. Renewable electricity reduces emissions by displacing fossil fuel generated electricity. The marginal change in fossil fuel generation caused by a marginal change in renewable electricity is called the marginal emission. Marginal emissions vary by location, time of day, and congestion on the local grid.

As a first pass, to assess the emissions effect of the bonus, I compare the the marginal emissions in Energy Communities to marginal emissions in non-Energy Communities. If Energy Communities have higher marginal emissions than non-Energy Communities, then the Energy Community Bonus will improve emissions reductions per subsidy dollar relative to an equivalent subsidy without the place-based adjustment.

4.1 Marginal Emissions Data

I use the NREL Cambium 2023 dataset for marginal emissions estimates, which is based on NREL's Regional Energy Deployment System (ReEDS) model of the US electricity system (Gagnon et al., 2023). Specifically, I use their long-run marginal emissions rates (LRMER) of CO_2-equivalent (CO_2-e).[22] The data are hour-of-year-level projections for 2025 to 2050, at five-year intervals, for he 134 ReEDS balancing areas. I merge the balancing-authority-level data to the Energy Community county data. I use the "Mid Case" Cambium scenario, based on policies as they existed in September 2023 and NREL's central estimates for future electricity demand, generator costs, and fuel prices.

A downside to the Cambium dataset is that is not based on empirical estimates, despite a wide body of research that generates empirical estimates of marginal emissions. However, Cambium offers two important benefits for this project. First, the LRMER estimates are preferable to the short-run estimates generated by the empirical literature. The long-run estimates adjust adjust for how changes in generation or demand may affect transmission construction and generator retirement, while short-run marginal emissions keep capital assets fixed. The Energy Community Bonus may cause new generators to be constructed that are large enough to initiate new transmission or new generators, so LRMER are better suited to the use case.

Perhaps an even more importantly, another benefit using Cambium is its estimates are available at a relatively disaggregated spatial resolution. Empirical marginal emissions estimates are spatially coarser, dividing the US into either three grid interconnections (Holland et al., 2022; Zivin et al., 2014), six NERC regions (Siler-Evans et al., 2012), or six major balancing authorities (Callaway et al., 2018; Fell and Kaffine, 2018). More recently, Holland et al. (2024) estimate marginal emissions at the plant-level, but the regions that cause the plant-level emissions changes are still relatively aggregated, with just 13 covering the US.

Meanwhile, Energy Communities are much smaller than these regions, defined at the county, census tract, or parcel level. There are 3,143 counties, over 84,000 census tracts, and over 150 million parcels in the US. The Bonus may shift renewable development within regions, with different emissions effects depending on the local marginal emissions. Cambium includes 134 "balancing areas," each with distinct marginal emissions. The emissions effect of the Bonus

[22] The LRMER model incorporates transmission losses, and the default LRMER use transmission losses based on marginal changes in demand. I therefore use the adjustment factor to convert the transmission losses to approximate marginal changes in generation, following the Cambium documentation.

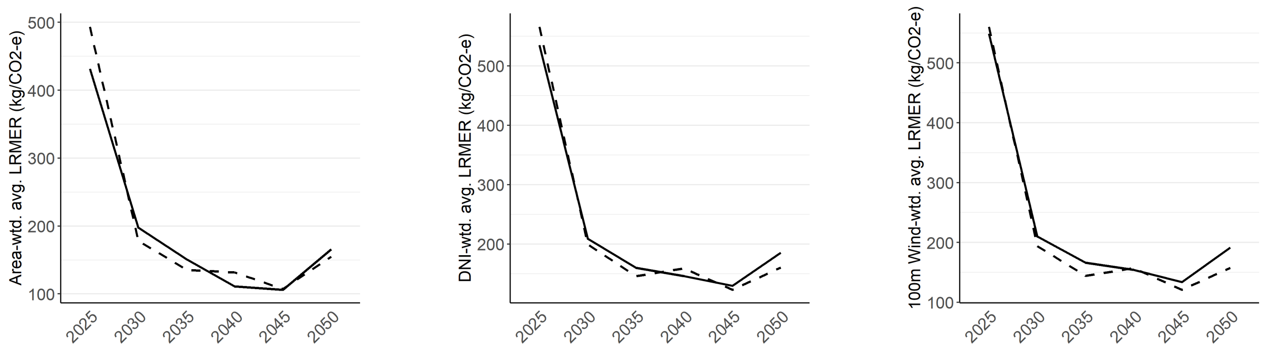

(a) Weighted by county area

(b) Weighted by county DNI

(c) Weighted by county wind speed

EC county   Non-EC county

Figure 8: Average marginal emissions by year

may depend on sub-regional variation in marginal emissions.[23] Future versions of this paper may explore using the empirical estimates from Holland et al. (2024) or developing county-level empirical estimates as described in Appendix C, but the Cambium estimates at least provide a first approximation.

4.2 Emissions Results

Figures 8 and 9 show how marginal emissions projections vary in Energy Communities and non-Energy Communities by year. The figures differ by how they define Energy Community counties. Figure 8 uses a broad definition, where Energy Community counties include those that were eligible as MSA or non-MSA Energy Communities in 2023 or 2024 or that contain a census tract that is an Energy Community for a coal mine or coal generator closure. The solid black line shows the projected average county LRMER in each year for Energy Communities, and the dashed line shows the same for non-Energy Communities. The three panels show three different weighted averages. The first shows the county area-weighted average marginal emissions, which is the relevant figure if renewable generation enters counties proportionate to its area. The second and third panel show the DNI- and wind-speed-weighted average marginal emissions, respectively. These are the relevant figures if renewable projects enter counties proportionately to their sunniness or wind speed.[24]

[23]This is true only to the extent there is actual sub-regional variation in marginal emissions, but there is reason to suspect sub-regional variation is important. Transmission lines often become congested, effectively splitting regional grids into many sub-regional grids. According to PJM (2024), "In a system as large as PJM, there is nearly always congestion somewhere on the system, and as a result there are nearly always multiple units on the margin." Congestion timing likely correlates with renewable generation (e.g., congested powerlines during hot, sunny hours), and since the Energy Community Bonus may shift renewable plant locations, it may be important to measure sub-regional marginal emissions under consistently congested grids.

[24]More specifically, the DNI figure uses county weights based on a county's maximum annual average DNI based on Sengupta et al. (2018), and the wind-speed-weighted figure uses county weights based on a county's maximum average annual wind speed based on Draxl et al. (2015).

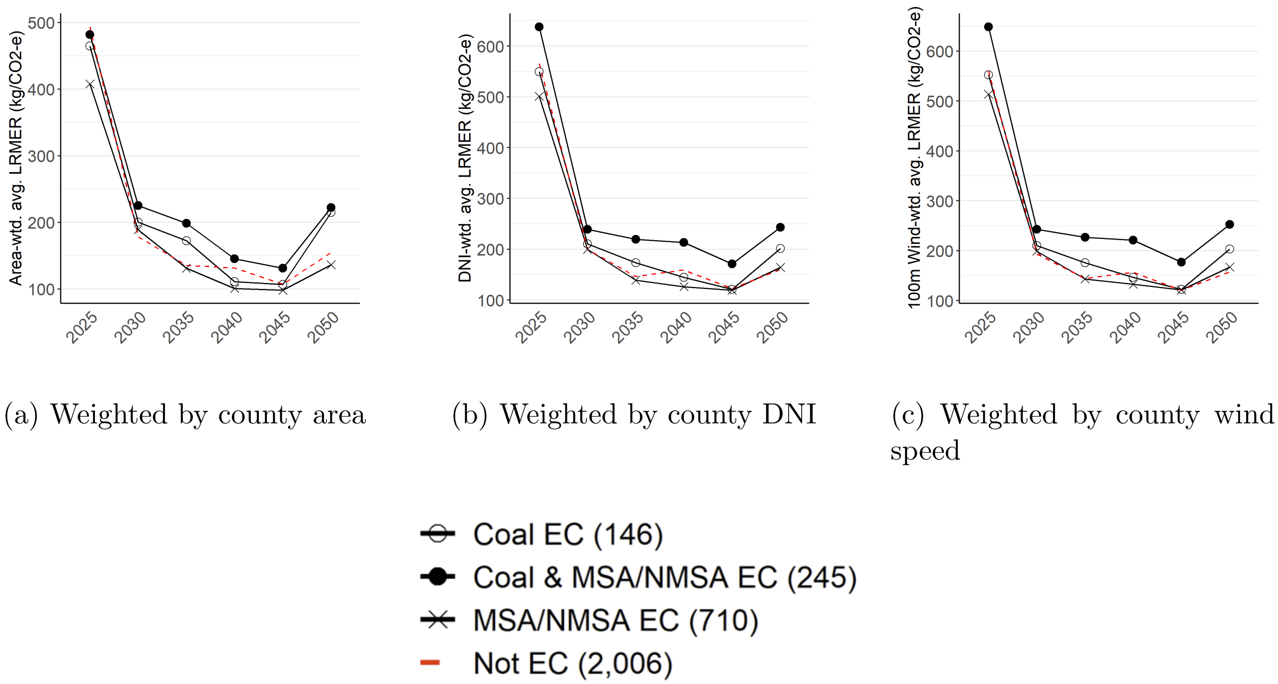

(a) Weighted by county area

(b) Weighted by county DNI

(c) Weighted by county wind speed

Figure 9: Average marginal emissions by year

Figure 8 shows average LRMERs are higher in non-Energy Community counties in 2025, meaning the Energy Community Bonus steers clean generation to counties with less environmental benefit. However, the trend reverses father in the future (with the exception of 2040), where the generation displaced in Energy Communities will be dirtier, and the Bonus shifts generators to counties with greater environmental benefit.

Figure 9 adds more nuance from different definitions Energy Communities. It shows whether coal tract Energy Communities or fossil fuel employment Energy Communities offer more environmental benefit. The solid black lines with different shaped points show the average LRMER by year for different Energy Community definitions, and the dashed red line shows average LRMER by year for non-Energy Communities. The number of counties in each group are in parentheses in the legend. The same as Figure 8, the three panels weight the averages by area, DNI, and wind speed, respectively. Energy Communities defined by only coal generator or coal mine closures offer more environmental benefit than Energy Communities defined by fossil fuel employment, on average. But the highest average environmental benefit is in the counties that meet both Energy Community definitions, with both high fossil fuel employment and a coal closure, offer the highest benefit.

Finally, Figure 10 shows how the results vary by hour of the day. Since solar and wind energy generation is intermittent, assuming imperfect energy storage capacity, the relevant LRMER for solar generation is during daylight hours, and for wind generation is during windier hours. The central hours of Figure 10 panel (b) is therefore a better evaluation of the expected LRMER that will be displaced by solar generation. The Energy Community Bonus is projected to steer solar generation to counties with less emissions benefit in 2025, and although the trend reverses in 2030 and 2035, the difference is very small. The difference is also small for all hours comparing the counties most suitable for wind generation. But farther into the future, during

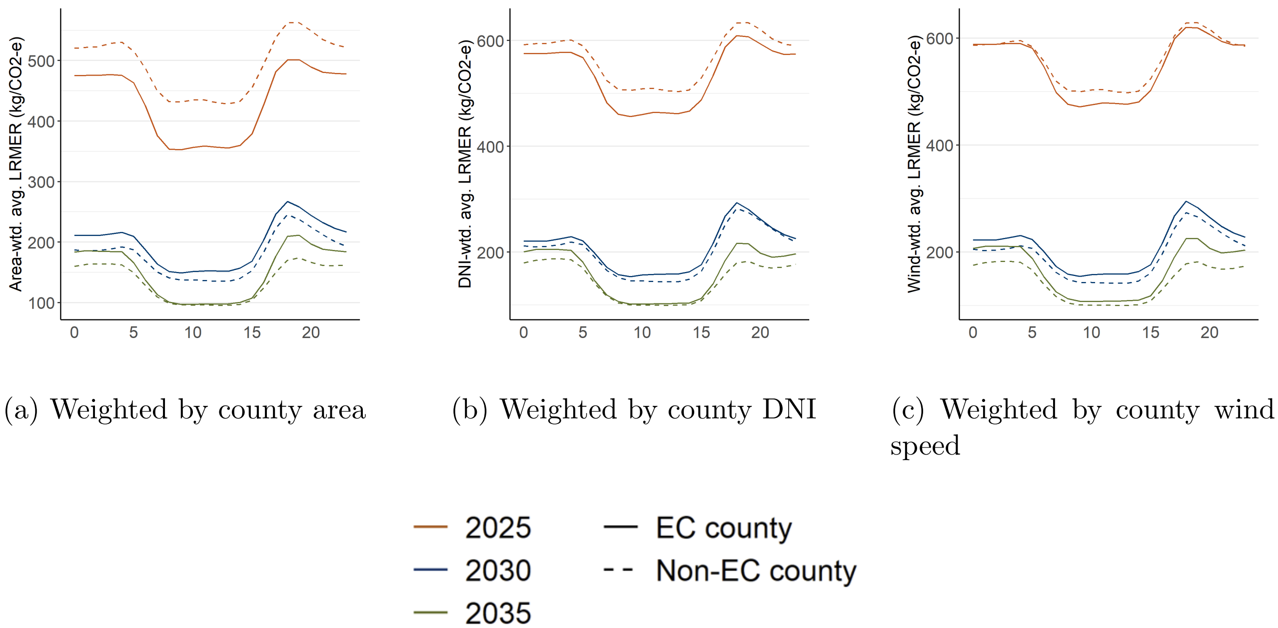

(a) Weighted by county area

(b) Weighted by county DNI

(c) Weighted by county wind speed

Figure 10: Average county marginal emissions by hour of day for different years

nighttime hours, Energy Community counties suitable for wind generation offer substantially more environmental benefit than comparable non-Energy Community counties.

4.3 Emissions Discussion

The Energy Community Bonus will shift generation to areas with higher marginal emissions in 2025, but lower marginal emissions in 2030 and 2035. The post-2030 results may be more relevant to evaluate the Bonus effect because renewable generators that are built in locations due to the Bonus will exist for many years. However, the forecasts farther in the future are more uncertain. The difference between the marginal emissions in Energy Communities and non-Energy Communities seems small, but may still result in substantial total emissions effects. Persistent small differences can add up to substantial differences.

A limitation of this analysis is that it compares all Energy Communities and all non-Energy Communities, but many of these areas are not be suitable for renewable generation, and are therefore unrealistic to include. The next section describes an analysis to generate a more realistic counterfactual emissions scenarios. The result is a measure of the overall Energy Community Bonus emissions abatement relative to counterfactuals.

5 Next Steps

5.1 Simulating the Energy Community Bonus Effect on Emissions

With reliable estimates on marginal emissions and of the Bonus effect on clean energy capacity in hand, future analysis will combine the estimates and calculate the Bonus effect on emissions relative to counterfactuals. I model a baseline scenario of emissions abatement from clean generation caused by the existing Energy Community Bonus. I compare the baseline to three counterfactuals: zero additional abatement with no bonus anywhere, abatement if the Bonus is available everywhere, and an in-between scenario equally distributing the observed capacity caused by the Bonus.

Since the estimates of the Bonus effect in Sections 3 and 5.2 are measured in capacity, not generation, I convert the estimate of the Bonus effect to hourly generation. Define Gen_c as county-level generation caused by one year of being treated by the Energy Community Bonus, a vector the length of the hours in the year (8,760 hours). Annual generation caused by the Bonus in county c can be modeled as

$$Gen_c^{EC} = EC_c(\beta_c \times CF_c), \tag{3}$$

where EC_c is a constant equal to one if c is an Energy Community and zero otherwise. β is the Bonus effect on annual county clean energy capacity additions, a vector equal to the length of the number of clean energy technologies that may be treated by the Bonus. The elements of β add up to the Energy Community Bonus effect measured in MW estimated in Section 3, and the technology-specific quantities in β are proportionate to the type of technologies in similar counties. For example, if county c is very sunny, β_x will be high for solar generation and lower for wind, or county c in the Midwest, β_x will be high for wind generation and lower for solar. CF_c is a matrix of technology-hour-specific capacity factors for county c. $CF_{xs\tau c}$ can be estimated using simulated datasets available from NREL.[25,26]

The baseline scenario of emissions abatement caused by the Bonus in county c in year y is

$$A_{cy} = Gen_c^{EC} \cdot \sum_{z=y}^{Y} \phi_{cz}, \tag{4}$$

where ϕ_{cz} is a vector of the hourly marginal emissions in county c and year z, and the sum is an element-wise sum for each hour of each year. This is the sum of all present and future emissions abated from the causal effect of the Bonus year y to year Y. The total effect of the Bonus over that time period is $A_c = \sum_{y=2023}^{Y} A_{cy}$.

The counterfactual to A_{cy} when no counties get the Bonus is simply zero for all counties and years. The counterfactual annual emissions in county c and year y when all counties get the Bonus is obtained by amending Equation 3 to have EC_c equal to one for all c. Finally, the counterfactual for distributing the Bonus equally throughout the US recalculates generation as

$$Gen_c^{*} = \tilde{\beta}_c \times CF_c,$$

where $\tilde{\beta} = C^{-1} \sum_c \beta EC_c$, and C is the total number of counties. Note, this makes a key assumption that the additional capacity in Energy Communities due to the Bonus would be

[25] If β is measured as effects on the queue, there are at least two options to convert the result to generation. One options is to set the vector sum of β equal to the effect on new operating capacity. If using queue entry, another option is to amend Equation 3 to

$$Gen_c^{EC} = EC_c p(\beta_c \times CF_c),$$

where p is the probability new queued capacity becomes operating, estimable from the data.

[26] Specifically, PLUSWIND are modeled, hourly wind capacity factor data covering all wind plants in the US from 2018 to 2021 (Millstein et al., 2023), and the NREL Utility-Scale Solar (USS) project offers hourly plant-level solar generation data (Bolinger et al., 2023).

distributed equally to all counties under an equal-subsidy policy. Moreover, all the counterfactuals assume the estimated Bonus effect β is externally valid to counties that were not used in estimation (i.e., those farther from the threshold in the RD and non-neighbor counties in the DID).

5.2 Energy Community Bonus Effect on Existing Capacity

The interconnection queue analysis is an interesting first estimate of the Energy Community Bonus effect. However, it will be possible to estimate the Energy Community Bonus effect on existing generators that have completed the interconnection queue as more time passes since the IRA was enacted. The DID estimates even show that the Energy Community Bonus is affecting the rate at which new generators complete interconnection queues. The affected generators will soon be generating electricity if they are not already.

Future versions of this paper may therefore switch from an analysis of queue data to an analysis of new generators. Energy Information Administration (EIA) Form 860 data are administrative data on new electricity generator capacities and locations. The location data in EIA 860 Forms are also coordinate point data, which would enable census tract-level identification for the DID analysis. It could also enable identifying some or all Energy Community projects in brownfields.

6 Conclusion

The empirical analysis of the Bonus effect shows the Energy Community Bonus may already be affecting generators on the grid interconnection queue. This is notable for a few reasons. First, a longstanding question has been to what extent the rapidly growing clean energy industry in the US is a result of federal subsidies. The very existence of interconnection queues could suggest there is more supply than the grid can accommodate and supply is inelastic. This paper offers empirical evidence that the price elasticity of supply is nonzero, suggesting federal policies may play a role in clean energy growth. It remains uncertain whether the observed Energy Community Bonus effect is an extensive margin supply expansion or an intensive margin change in supply locations. But given that subsidies in non-Energy Communities have not changed, there is reason to believe there is some extensive margin quantity increase.

Moreover, if Energy Community Bonus has any effect, whether on total quantities or locations, the policy offers promise for future research on the local effects of clean energy. This is new variation in US clean energy plant locations that can be exploited in empirical analyses. The Energy Community Bonus could serve as an instrumental variable in analyses of how clean energy plants affect local employment and wages, electricity prices, environmental conservation, home prices, voting, and other outcomes of interest.

The results here highlight that, unlike a first-best tax policy that necessarily entails an efficiency-equity trade-off, there is not necessarily an efficiency-equity trade-off with a second-best subsidy policy. Recent research has examined the income distributional effects of home energy subsidies (Borenstein and Davis, 2024), and further research may investigate whether distributional mechanisms could improve other policies' efficiency. In the case of the Energy Community Bonus, it is an empirical question whether a place-based clean electricity subsidy will increase or decrease emissions relative to a spatially constant subsidy. I find that the Energy Community Bonus is an equity mechanism that also improves long-run policy efficiency.

References

Energy community tax credit bonus faqs, 2024. URL https://energycommunities.gov/energy-community-tax-credit-bonus-faqs/. Accessed August, 2024.

BLS. Local area unemployment statistics. URL https://www.bls.gov/lau/tables.htm.

Mark Bolinger, Joachim Seel, Julie Mulvaney Kemp, Cody Warner, Anjali Katta, and Dana Robson. Utility-scale solar, 2023 edition: Empirical trends in deployment, technology, cost, performance, ppa pricing, and value in the united states. 2023.

Severin Borenstein and Lucas W Davis. The distributional effects of us tax credits for heat pumps, solar panels, and electric vehicles. Technical report, National Bureau of Economic Research, 2024.

Duncan S Callaway, Meredith Fowlie, and Gavin McCormick. Location, location, location: The variable value of renewable energy and demand-side efficiency resources. *Journal of the Association of Environmental and Resource Economists*, 5(1):39–75, 2018.

Sebastian Calonico, Matias D Cattaneo, and Rocio Titiunik. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326, 2014.

Sebastian Calonico, Matias D Cattaneo, and Max H Farrell. Optimal bandwidth choice for robust bias-corrected inference in regression discontinuity designs. *The Econometrics Journal*, 23(2):192–210, 2020.

US Census. County business patterns. URL https://www.census.gov/programs-surveys/cbp.html.

DOE, NETL, Interagency Working Group on Coal & Power Plant Communities & Economic Revitalization, and US Department of the Treasury. Ira energy community data layers. 2023. URL https://edx.netl.doe.gov/dataset/ira-energy-community-data-layers.

DOE, NETL, Interagency Working Group on Coal Power Plant Communities Economic Revitalization, and US Department of the Treasury. Ira energy community data layers. 2024. URL https://edx.netl.doe.gov/dataset/ira-energy-community-data-layers.

Caroline Draxl, Andrew Clifton, Bri-Mathias Hodge, and Jim McCaa. The wind integration national dataset (wind) toolkit. *Applied Energy*, 151:355–366, 2015.

EIA. Wholesale electricity and natural gas market data, 2022. URL https://www.eia.gov/electricity/wholesale/.

EIA. Continuous emissions monitoring system (cems), 2024. URL https://campd.epa.gov/.

Harrison Fell and Daniel T Kaffine. The fall of coal: Joint impacts of fuel prices and renewables on generation and emissions. *American Economic Journal: Economic Policy*, 10(2):90–116, 2018.

FERC. Improvements to generator interconnection procedures and agreements, 2023. URL https://www.ferc.gov/media/e-1-order-2023-rm22-14-000.

Pieter Gagnon, Pedro Andres Sanchez Perez, Kodi Obika, Marty Schwarz, James Morris, Jianli Gu, and Jordan Eisenman. Cambium 2023 scenario descriptions and documentation. Technical report, National Renewable Energy Laboratory (NREL), Golden, CO (United States), 2023.

Stephen P Holland, Matthew J Kotchen, Erin T Mansur, and Andrew J Yates. Why marginal co2 emissions are not decreasing for us electricity: estimates and implications for climate policy. *Proceedings of the National Academy of Sciences*, 119(8):e2116632119, 2022.

Stephen P Holland, Erin T Mansur, Valentin Verdier, and Andrew J Yates. Regularization from economic constraints: A new estimator for marginal emissions. Technical report, National Bureau of Economic Research, 2024.

Sarah Johnston, Yifei Liu, and Chenyu Yang. An empirical analysis of the interconnection queue. *Working Paper*, 2024.

Dev Millstein, Seongeun Jeong, Amos Ancell, and Ryan Wiser. A database of hourly wind speed and modeled generation for us wind plants based on three meteorological models. *Scientific Data*, 10(1):883, 2023.

PJM. Marginal emissions rates: A primer. 2024. URL https://www.pjm.com/-/media/etools/data-miner-2/marginal-emissions-primer.ashx.

Joseph Rand, Rose Strauss, Will Gorman, Joachim Seel, Julie Mulvaney Kemp, Seongeun Jeong, Dana Robson, and Ryan H Wiser. Queued up: Characteristics of power plants seeking transmission interconnection as of the end of 2023. 2024. URL https://emp.lbl.gov/queues.

Manajit Sengupta, Yu Xie, Anthony Lopez, Aron Habte, Galen Maclaurin, and James Shelby. The national solar radiation data base (nsrdb). *Renewable and sustainable energy reviews*, 89:51–60, 2018.

Kyle Siler-Evans, Ines Lima Azevedo, and M Granger Morgan. Marginal emissions factors for the us electricity system. *Environmental science & technology*, 46(9):4742–4748, 2012.

US Treasury. Irs notice 2013-29. 2013. URL https://www.irs.gov/pub/irs-drop/n-13-29.pdf.

US Treasury. Irs notice 2023-29. 2023a. URL https://www.irs.gov/pub/irs-drop/n-23-29.pdf.

US Treasury. Irs notice 2023-29 appendix a. 2023b. URL https://www.irs.gov/pub/irs-drop/n-23-29-appendix-a.pdf.

US Treasury. Irs notice 2023-45. 2023c. URL https://www.irs.gov/pub/irs-drop/n-23-45.pdf.

US Treasury. Irs notice 2023-47. 2023d. URL https://www.irs.gov/pub/irs-drop/n-23-47.pdf.

US Treasury. Irs notice 2024-30. 2023e. URL https://www.irs.gov/pub/irs-drop/n-24-30.pdf.

Brett Watson, Ian Lange, and Joshua Linn. Coal demand, market forces, and us coal mine closures. *Economic Inquiry*, 61(1):35–57, 2023.

Joshua S Graff Zivin, Matthew J Kotchen, and Erin T Mansur. Spatial and temporal heterogeneity of marginal emissions: Implications for electric cars and other electricity-shifting policies. *Journal of Economic Behavior & Organization*, 107:248–268, 2014.

A Additional Figures and Tables

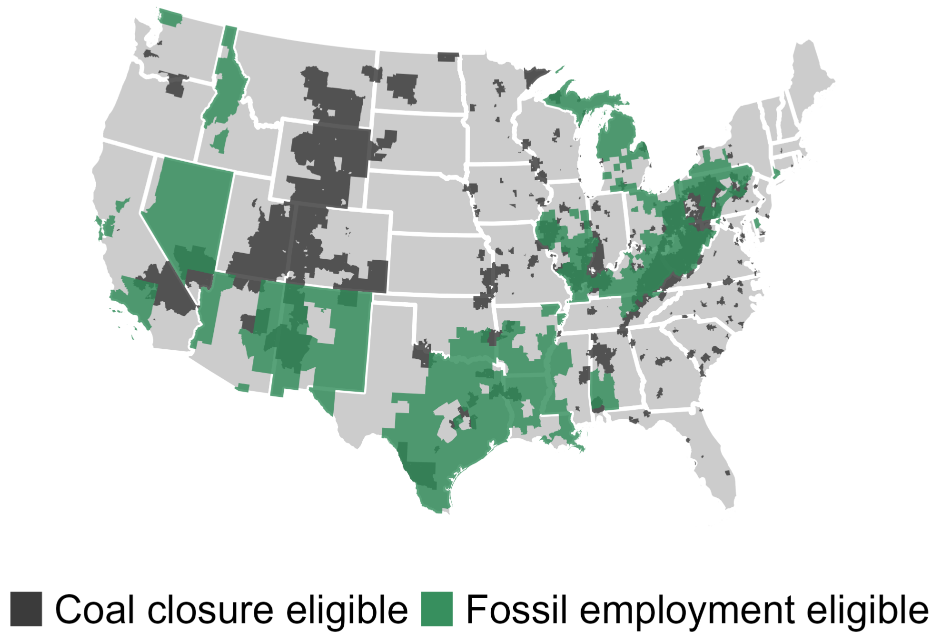

Figure A.1: Energy Community census tracts, MSAs, and non-MSAs in tax year 2024. This map uses definitions based on fossil fuel direct employment, coal mine closures, and coal electricity generator closures. Note, the figure uses fossil fuel employment NAICS codes from 2023, and udpated Treasury guidance in 2024 added more NAICS codes. However, this only results in one additional MSA in Ohio added as an Energy Community. Future iterations of this paper will include the updated NAICS code for 2024.

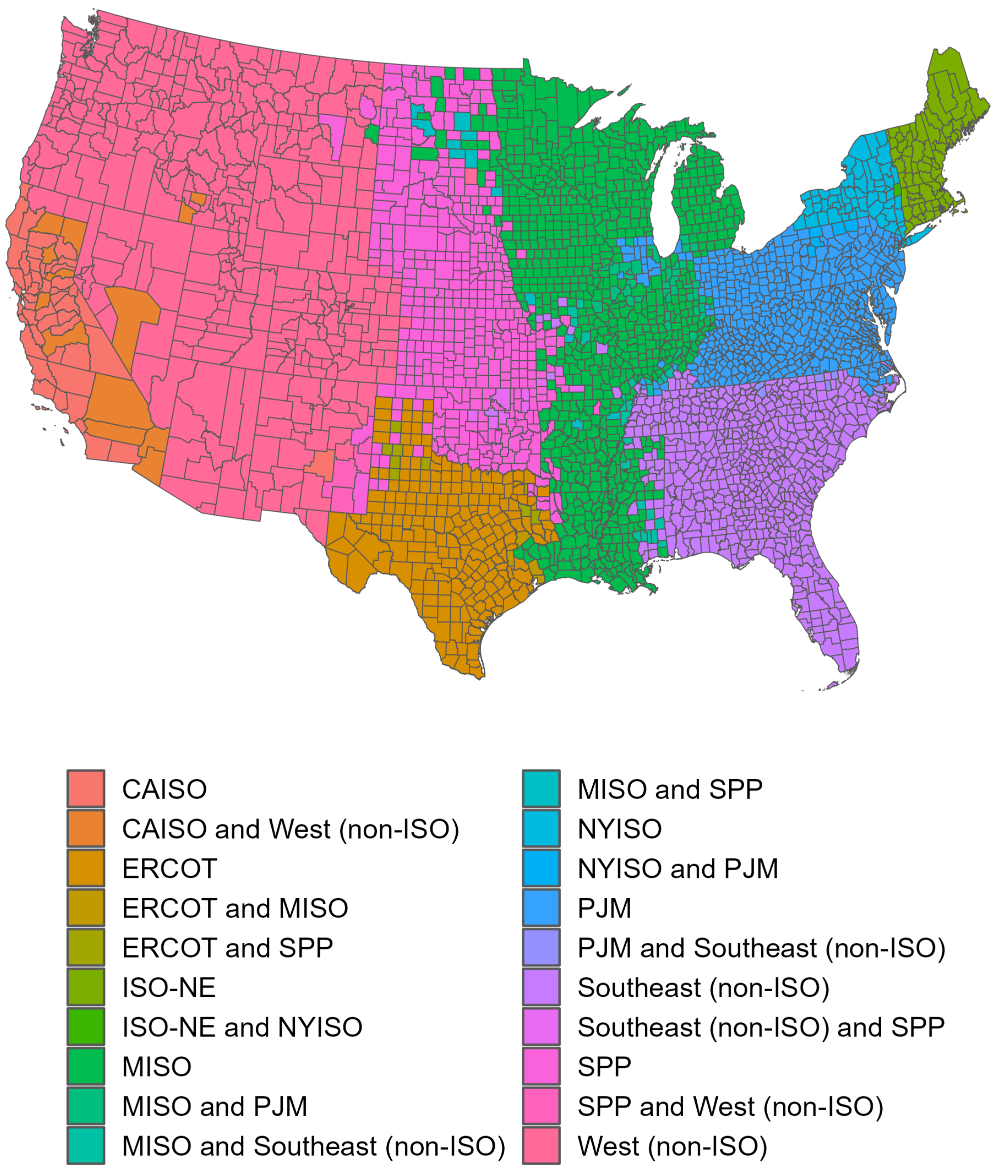

Figure A.2: Balancing authorities in counties. Map is based on counties appearing in queued generation data when available, and state when there have been no queued projects since 2017.

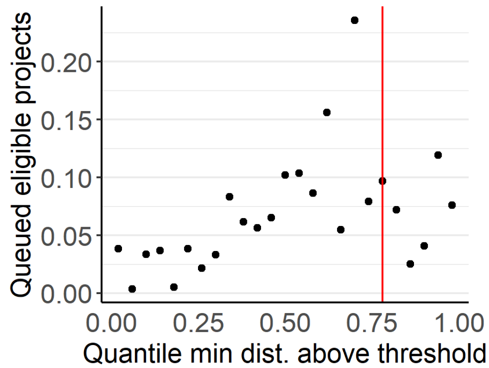

(a) Queue entry, number of projects

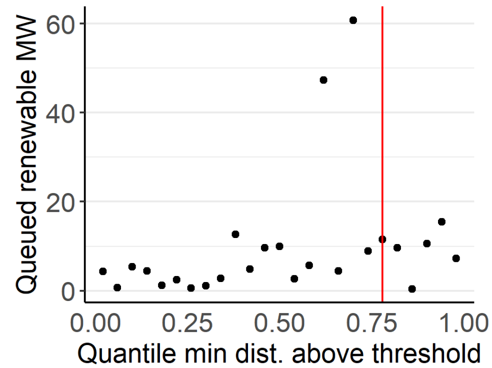

(b) Queue entry, MW

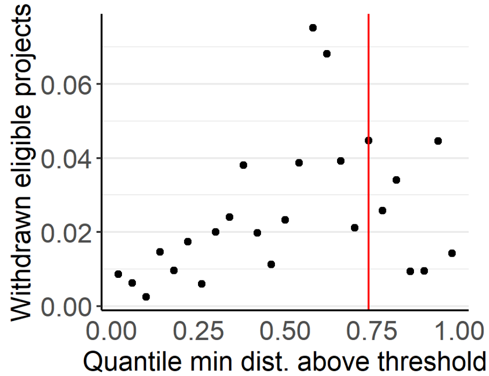

(c) Withdrawal, number of projects

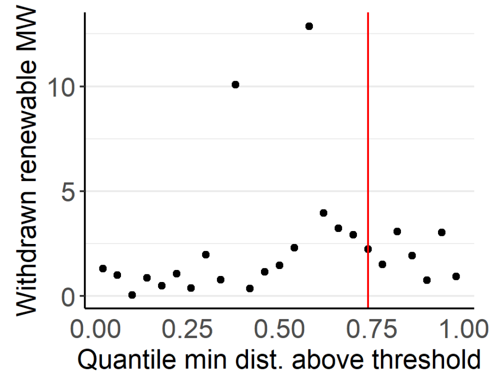

(d) Withdrawal, MW

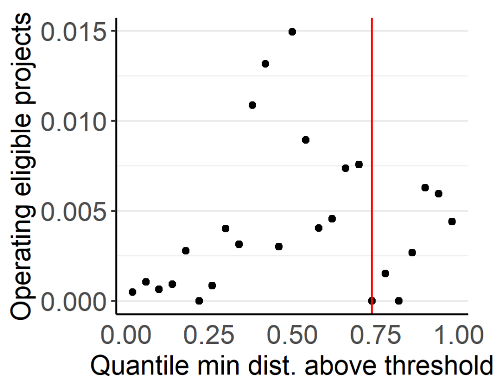

(e) Queue completion, number of projects

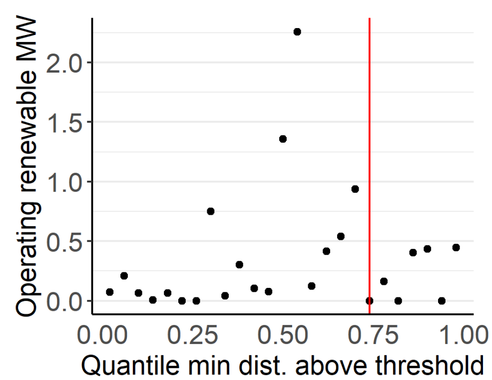

(f) Queue completion, MW

Figure A.3: The x-axes show the quantile of the maximum difference between the fossil fuel employment or previous year unemployment and their respective thresholds. The quantiles are generated using post-period data such that the bins are exactly the same as Figure 5. The vertical red line intersects the axis at the quantile that includes zero. The y-axes are the bin means of the outcome variables pre-IRA. These graphs are useful to compare trends on these quantile bins before and after the IRA.

	Mine Closure:	
	All BAs	*Major ISOs/RTOs*
Queues MW	0.055***	0.034***
	(0.002)	(−0.008)
Withdrawals MW	−0.138***	−0.222***
	(−0.021)	(0.053)
Operational MW	−0.121***	−0.092***
	(−0.002)	(0.017)
Population/1000	−0.011***	−0.026***
	(−0.00001)	(0.001)
Mean Income/1000	−0.018***	−0.025***
	(−0.004)	(−0.006)
Avg. weekly wage	0.001**	0.003***
	(−0.0003)	(−0.001)
Employment/1000	0.020***	0.036***
	(0.001)	(0.0004)
Constant	0.986*	−0.140
	(0.532)	(1.326)
Observations	170	55
Log Likelihood	−111.832	−32.211
Akaike Inf. Crit.	239.664	80.421
Note:		*p<0.1; **p<0.05; ***p<0.01

Table A.1: Results for a logistic regression of an indicator variable equal to one if it had a mine closure. The counties included are mapped in Figure 6, with descriptive statistics in Table 5.

	Queued		*Withdrawn*		*Operating*	
	Projects	MW	Projects	MW	Projects	MW
Post-IRA	−0.002	0.208	0.004	−0.500	−0.003	−0.636
	(0.005)	(0.696)	(0.006)	(0.408)	(0.002)	(0.415)
EC	−0.001	0.133	0.014*	−0.425	−0.001	−0.937*
	(0.004)	(0.360)	(0.007)	(0.471)	(0.005)	(0.568)
EC×Post-IRA	0.014**	1.288*	−0.045**	0.021	0.009	2.395**
	(0.007)	(0.747)	(0.020)	(0.530)	(0.007)	(1.194)
DNI	−0.071	−3.922	−0.029	1.046	−0.248	−34.651
	(0.048)	(3.076)	(0.259)	(17.770)	(0.154)	(28.529)
DNI^2	0.006	0.310	0.001	−0.338	0.027	3.901
	(0.004)	(0.246)	(0.027)	(1.946)	(0.017)	(3.189)
Wind speed	0.017	0.317	0.142*	0.570	0.015	−3.888
	(0.013)	(0.722)	(0.074)	(3.096)	(0.034)	(4.077)
Wind $speed^2$	−0.001	−0.011	−0.008*	−0.034	−0.001	0.214
	(0.001)	(0.034)	(0.005)	(0.183)	(0.002)	(0.252)
Population	−0.00000	0.000	0.00000***	0.00000	−0.00000	−0.00002***
	(0.00000)	(0.00000)	(0.00000)	(0.00000)	(0.00000)	(0.00001)
Avg. Income	0.00000	−0.00001	−0.00000*	−0.00002*	0.00000	0.00003**
	(0.00000)	(0.00001)	(0.00000)	(0.00001)	(0.00000)	(0.00002)
Avg. Wage	0.00001	0.002	0.00000	0.001	0.00002	0.001
	(0.00001)	(0.002)	(0.00002)	(0.001)	(0.00002)	(0.002)
Employment	0.00000	−0.00000	−0.00000***	0.00000	0.00000	0.00004***
	(0.00000)	(0.00000)	(0.00000)	(0.00000)	(0.00000)	(0.00001)
BA×year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	13,260	13,260	4,290	4,290	4,290	4,290
R^2	0.053	0.078	0.132	0.029	0.017	0.029
Note:						*p<0.1; **p<0.05; ***p<0.01

Table A.2: Full estimation results for Equation 2. Standard errors in parentheses are clustered at the county level. The regressions are weighted with inverse propensity scores.

B Additional Data Cleaning

.

Hawaii and Alaska are omitted from the queue data, so I drop them from all datasets. About 3 percent of projects recorded by NREL Queued Up since 2018 have null county names and are dropped from the analysis. Some counties in the queue datasets do not match census data because of typos. I manually clean some common and obvious typos in county names. Still, about 3 percent of queued entrants, 5 percent of queued withdrawals, and 5 percent of newly operating queued projects are dropped from the analysis.

Treasury (2023a,b,c,d,e) define the Energy Community definition data. The NAICS codes used to define Energy Communities were expanded in March 2024, over a year after the Bonus had been made available. The analyses below use the 2023 definition of fossil fuel employment for queue effects in 2023 and before and the updated definition of fossil fuel employment thereafter.

C County-level Marginal Emissions

This section describes an approach to estimating marginal emissions from electricity generation at the county-level. Like previous literature, it uses hourly generator-level emissions, regressed on an exogenous shifter. However, most literature uses small changes in demand (load) as that exogenous shifter (Holland et al., 2022; Fell and Kaffine, 2018; Siler-Evans et al., 2012; Zivin et al., 2014). Instead, like Callaway et al. (2018), the proposed analysis treats small changes in wind and solar electricity generation as exogenous. The next section describes the data, and the following section describes the method.

C.1 County-level Marginal Emissions Data

Hourly, generator-level CO2 emissions data come from the EPA Continuous Emissions Monitoring System (CEMS) (EIA, 2024). The CEMS data track fossil fuel generators with at least 25 MW capacity, covering about 97 percent of generation. I aggregate the unit-level data to obtain hourly balancing-authority-level emissions. I collect emissions and generation data covering 2018 to 2021.

To estimate county-level marginal emissions, I use exogenous variation in solar and wind generation, as explained below. This requires nationwide data on hourly, plant-level solar and wind electricity generation. Unfortunately, no such data exists in the US. Instead, I use two datasets that model historic wind and solar generation based on observed meteorological conditions, both produced by the Lawrence Berkeley National Laboratory.

The modeled plant-level, hourly wind generation data come from PLUSWIND (Millstein et al., 2023). PLUSWIND converts three meteorological models of hourly wind speed into three corresponding estimates of plant-level wind electricity generation. The data are available for all wind plants in the contiguous US from 2018 to 2021.

The wind generation data are alone sufficient to identify marginal emissions in counties with wind generators, but most counties do not have wind generators. I therefore supplement the wind generation data with another source of identifying variation: plant-level, hourly solar generation data from the Utility-Scale Solar (USS) initiative (Bolinger et al., 2023). The USS data rely on the National Solar Radiation Database (NSRDB) based on satellite imagery from the GOES satellites and NOAA High Resolution Rapid Refresh (HRRR) data, as well as measured data from ERCOT and CAISO where available.

C.2 County-level Marginal Emissions Method

I estimate county-by-season-by-hour marginal emissions. The time of generation is important because the county-specific marginal emissions may vary with load and temperature, which affect grid congestion. Hourly marginal emissions are especially important because solar and wind generation varies systematically by hour of day, so this level of temporal resolution is necessary to estimate the emissions effects of future solar and wind plants.

I estimate separate regressions for each balancing authority, each obtaining county-level marginal emissions for the counties within the balancing authority. I assume the fossil fuel emissions are only displaced in the same balancing authority as the generation. Although electricity sometimes moves between balancing authorities ("interchange"), most electricity demand in a balancing authority is satisfied by generation in the same balancing authority.

The estimating equation is

$$CO2e_{bt} = \alpha_{bst} + \beta_L Load_{bt} + \phi_{cs\tau} RenewGen_{ct} + \varepsilon_{bt}$$

where $CO2e_{bt}$ is CO2-equivalent emissions released in balancing authority b in day t (e.g., January 1st 2021 at 12:00am to 12:59am), $RenewGen_{ct}$ is the sum of solar and wind generation in county c in balancing authority b in hour t, the τ subscript denotes the hour of the day (e.g., 12:00am to 12:59am) and the s subscript indexes season (e.g., winter). $Load_{bt}$ is electricity demand, treated as exogenous because electricity demand is price inelastic and dictated by weather or demand shocks orthogonal to emissions. I control for load because wind generation may be correlated with load (e.g., fewer conditioners running when it is windy), and if left an omitted variable, changes to renewable generation may seem to cause emissions changes that are in fact caused by higher load. The balancing-authority-by-season-by-hour fixed effect $\alpha_{bs\tau}$ absorbs typical seasonal variation. The estimate of marginal emissions ϕ_{cs} is at the county by hour-of-day by season level.

One challenge to estimating this equation is collinearity across nearby counties, since variation in solar and wind variation may be similar in nearby counties. Pooling neighboring counties based on the degree of collinearity or geographical features may improve this outcome. Another challenge is that the solar and wind generation data will not cover all hours of the day in all counties. The counties without any marginal emissions estimates are likely less suitable for wind and solar generation, and therefore it is less important to model their marginal emissions for the purpose of predicting the effect of new renewable generators. Nonetheless, pooling neighboring counties can also help get marginal emissions estimates for counties with no renewable generation in the data.