

Where the Sun Don't Shine: The Null Effect of IRA Energy Community Bonus Credits on Clean Energy Investment

APEP Autonomous Research* @olafdrw

February 19, 2026

Abstract

The Inflation Reduction Act offers a 10 percentage point bonus tax credit for clean energy projects in “energy communities” with significant fossil fuel employment. I exploit the statutory 0.17% fossil fuel employment threshold using a sharp regression discontinuity design across 779 MSAs and non-MSAs. I find no positive effect and suggestive evidence of a *negative* effect: the baseline estimate is -5.28 MW per 1,000 employees ($p = 0.198$); with covariates, -8.14 MW ($p = 0.015$). Nearly all specifications produce negative estimates, robust across bandwidths and polynomial orders. Total clean energy capacity shows a significant negative discontinuity ($p = 0.011$), and a pre-IRA placebo confirms the pattern predates the policy. The IRA’s largest place-based provision targets areas where renewable energy fundamentals are weakest.

JEL Codes: H23, Q42, Q48, R11

Keywords: Inflation Reduction Act, energy communities, clean energy investment, place-based policy, regression discontinuity

*Autonomous Policy Evaluation Project. This paper was generated autonomously. Total execution time: 1h 25m (cumulative: 1h 25m). Correspondence: scl@econ.uzh.ch

1. Introduction

In August 2022, the United States enacted the largest climate investment in its history. The Inflation Reduction Act committed an estimated \$369 billion to clean energy, with a central promise: the green transition would not leave fossil fuel communities behind. To operationalize this promise, the law created “energy communities”—areas with significant fossil fuel employment or recent coal closures—and offered a 10 percentage point bonus on the Investment Tax Credit (ITC) and Production Tax Credit (PTC) for clean energy projects sited in these areas. The bonus increases the ITC from 30% to 40%, a 33% increase in the effective subsidy rate. This provision is the IRA’s primary mechanism for ensuring that the benefits of decarbonization flow to the communities most disrupted by it.

But does it work? Solar panels need sunlight. Wind turbines need wind. The optimal locations for renewable energy are determined substantially by geography and physics—the solar resource in the Southwest, the wind corridor of the Great Plains. A tax incentive can alter the economics at the margin, but whether a 10 percentage point bonus is sufficient to overcome the fundamental pull of natural resource endowments is an open empirical question. If developers build where the sun shines and the wind blows regardless of bonus credits, then the IRA’s most prominent place-based provision is a transfer to inframarginal projects rather than a tool for spatial redistribution.

This paper provides the first quasi-experimental estimate of whether IRA energy community designation causally increases clean energy investment. I exploit the statutory threshold that determines energy community status: areas qualify if fossil fuel employment exceeds 0.17% of total employment (Internal Revenue Code §45(b)(11)(B)(ii)) and local unemployment exceeds the national average. The 0.17% threshold was set by Congress and is applied mechanically to pre-existing employment data, creating a sharp discontinuity in investment incentives that is plausibly exogenous to local conditions.

My regression discontinuity design compares MSAs and non-MSAs just above and below the 0.17% fossil fuel employment threshold, among areas where unemployment already exceeds the national average (so the second qualification criterion is mechanically satisfied). This isolates the effect of energy community designation through the fossil fuel employment pathway alone. I construct the running variable from Census County Business Patterns data and measure outcomes using the universe of electricity generators from EIA Form 860.

The results are striking—but not in the direction one might expect. The baseline robust bias-corrected RD estimate for post-IRA (2023+) clean energy capacity is -5.28 MW per 1,000 employees ($SE = 4.10$, $p = 0.198$), negative but imprecise. With pre-determined covariates, the estimate sharpens to -8.14 MW per 1,000 employees ($SE = 3.33$, $p = 0.015$),

statistically significant at the 5% level. Against an outcome mean of just 0.84 MW per 1,000 employees, these estimates imply that energy community designation is associated with substantially *less* clean energy investment, not more. The negative pattern is robust across six alternative bandwidths (0.5–2× the CCT optimal) and under quadratic polynomials. All bandwidth multipliers produce negative point estimates. Placebo cutoffs at nine alternative thresholds show no systematic pattern, and a McCrary density test confirms no manipulation of the running variable ($p = 0.33$).

The negative pattern extends beyond post-IRA generators. Total clean energy capacity (all years) shows a significant negative discontinuity at the threshold ($p = 0.011$), indicating that the deficit in energy community areas predates the IRA. Proposed generators show a negative but insignificant effect ($p = 0.216$). Both MSA and non-MSA areas independently exhibit negative point estimates. The bivariate RDD exploiting both qualification thresholds (fossil fuel employment and unemployment) jointly confirms the finding.

These results contribute to three literatures. First, they provide the earliest causal evidence on the real-world effects of the IRA’s energy community provision, complementing simulation studies by [Bistline et al. \(2023\)](#) and [Holland et al. \(2025\)](#) that model projected impacts. The suggestive negative effect challenges optimistic projections more sharply than a simple null would: the bonus credit does not merely fail to redirect investment, it may target areas that are inherently less attractive for renewables. Second, this paper speaks to the broader place-based policy literature ([Kline and Moretti, 2014](#); [Busso et al., 2013](#); [Neumark and Simpson, 2015](#); [Greenstone et al., 2010](#)), providing a case where place-based incentives confront—and fail to overcome—the economic fundamentals of the targeted industry. Third, it contributes to environmental economics by documenting a tension between climate policy and spatial equity: the locations optimal for clean energy production are not the locations most in need of economic transition support, and the IRA’s mechanism for bridging this gap appears insufficient.

The paper proceeds as follows. Section 2 describes the institutional setting and the IRA’s energy community provisions. Section 3 presents the data and defines the analysis sample. Section 4 details the regression discontinuity design and identification strategy. Section 5 reports the main results. Section 6 presents extensive robustness checks. Section 7 discusses mechanisms and implications. Section 8 concludes.

2. Institutional Background

2.1 The Inflation Reduction Act and Energy Communities

The Inflation Reduction Act of 2022 (P.L. 117-169) represents the largest federal investment in clean energy in U.S. history. Among its many provisions, the law created bonus tax credits for clean energy projects located in designated “energy communities.” These bonuses add 10 percentage points to the base ITC (bringing it from 30% to 40% for qualifying projects) and an additional \$0.03/kWh to the PTC.

The IRA defines three pathways to energy community status. First, brownfield sites as defined under CERCLA. Second, census tracts where a coal mine closed after 1999 or a coal-fired generator retired after 2009, plus directly adjacent tracts. Third—and the focus of this paper—metropolitan and non-metropolitan statistical areas that satisfy two conditions simultaneously:

1. **Fossil fuel employment:** At any point after December 31, 2009, direct employment in extraction, processing, transport, or storage of coal, oil, or natural gas must equal or exceed 0.17% of total area employment.
2. **Unemployment:** The area’s unemployment rate must equal or exceed the national average for the prior year.

The 0.17% fossil fuel employment threshold is codified in IRC §45(b)(11)(B)(ii) and is applied using data from the Census Bureau’s County Business Patterns (CBP) survey. The Treasury Department publishes annual lists of qualifying MSAs and non-MSAs based on CBP employment data and Bureau of Labor Statistics unemployment figures.

2.2 Treasury Implementation and Timeline

The IRA’s energy community provisions required substantial regulatory implementation by the Department of the Treasury and the Internal Revenue Service. Initial guidance arrived in Notice 2023-29, issued in April 2023, which clarified the three designation pathways and announced the first list of qualifying areas based on CBP and BLS data. The Treasury subsequently published updated lists in June 2023 and annually thereafter, incorporating the most recent available employment and unemployment data.

The implementation timeline is important for interpreting the results. Although the IRA was signed in August 2022, developers could not know with certainty which areas would qualify under the statistical area pathway until Treasury published the first designations eight months later. This regulatory lag may have dampened the early investment response, as

developers faced uncertainty about whether specific areas would receive the bonus credit. By mid-2023, however, the designation list was public and the regulatory framework was clear.

The bonus credit structure itself was designed to be straightforward to claim. Projects qualify for the energy community bonus based on their physical location—specifically, the census tract or statistical area in which the project is sited. Developers do not need to apply for a separate designation or meet additional criteria beyond the geographic requirement. This simplicity was intentional: the legislative history suggests Congress wanted the bonus to function as an automatic geographic overlay on the existing ITC and PTC, minimizing administrative burden and reducing the transaction costs that have hindered uptake of other place-based incentive programs.

2.3 The 0.17% Threshold as a Natural Experiment

The statutory threshold creates a sharp discontinuity in the investment incentive landscape. For an area where unemployment already exceeds the national average, crossing 0.17% fossil fuel employment flips the energy community switch: clean energy projects in that area receive a substantially larger tax credit. The discontinuity in incentives is immediate and binary—there is no partial eligibility or phase-in.

The magnitude of the incentive discontinuity is economically significant. For a utility-scale solar project claiming the ITC, energy community designation increases the credit from 30% to 40% of eligible project costs. On a \$200 million solar installation—typical for a 100 MW facility—this translates to an additional \$20 million in tax credits. For wind projects claiming the PTC, the bonus adds \$0.03/kWh over the project’s 10-year credit period, potentially representing tens of millions of dollars in additional value. These are not trivial incentives: they exceed the profit margins on many renewable energy projects and, in principle, should meaningfully alter siting decisions at the margin.

Several features of this institutional setting support a credible regression discontinuity design. First, the 0.17% threshold was determined by Congress in 2022 and applied to pre-existing employment data. Areas could not have anticipated or manipulated their fossil fuel employment share in response to a threshold that did not exist when the employment was measured. The CBP data used to determine fossil fuel employment shares are derived from administrative records—specifically, the Census Bureau’s Business Register—and reflect actual firm-level payroll reports rather than self-reported survey data subject to strategic misreporting.

Second, the threshold is applied mechanically by the Treasury Department using administrative data (CBP), leaving no room for discretionary interpretation. Unlike many place-based programs where designation involves an application process, political lobbying,

or bureaucratic discretion (as with Empowerment Zones or Opportunity Zone nominations by state governors), the energy community determination is a purely algorithmic exercise. An area either crosses the 0.17% threshold or it does not.

Third, the threshold is arbitrary in the sense relevant for RDD—there is no economic theory suggesting that 0.17% is a meaningful boundary for any outcome other than energy community designation itself. The 0.17% figure emerged from the legislative process and was likely calibrated to produce a desired number of qualifying areas rather than to identify a theoretically meaningful discontinuity in economic conditions.

2.4 The Place-Based Policy Debate

The energy community provision embodies a broader tension in U.S. industrial policy between efficiency and equity. Place-based policies—from Enterprise Zones (Busso et al., 2013) to Opportunity Zones (Freedman et al., 2021)—aim to direct economic activity to disadvantaged areas. The evidence on their effectiveness is mixed. Kline and Moretti (2014) find positive effects of the Tennessee Valley Authority, while Neumark and Simpson (2015) survey a more ambiguous literature on enterprise zones.

The theoretical case for place-based subsidies rests on agglomeration externalities, labor market frictions, or distributional objectives. If workers are spatially immobile and local labor markets are thick, then place-based policies can improve welfare by bringing jobs to workers rather than expecting workers to relocate. The empirical evidence is more ambiguous. Enterprise Zones have generally produced small or insignificant effects on employment, with some notable exceptions where program design included strong compliance mechanisms. Opportunity Zones appear to have increased investment in designated tracts, though the distributional incidence of those investments remains debated.

Clean energy introduces a distinctive challenge for place-based policy that has no close parallel in the existing literature: the economic viability of renewable generation depends heavily on natural resource endowments—solar irradiance, wind speed, land availability—that are geographically determined and largely immutable. A factory can be built almost anywhere with adequate infrastructure; a solar farm’s output varies by a factor of two depending on latitude and cloud cover. This creates a potential tension between siting projects where they are most productive and siting them where they would most benefit communities affected by the energy transition. The energy community provision is, to my knowledge, the first large-scale attempt to use tax incentives to redirect an industry whose output is fundamentally tied to geographic endowments.

3. Data

3.1 County Business Patterns

I construct the running variable—fossil fuel employment as a share of total employment—from the 2021 County Business Patterns (CBP) survey administered by the Census Bureau. CBP provides annual employment by industry at the county level using NAICS codes. I aggregate employment across four fossil fuel industry categories: oil and gas extraction (NAICS 211), coal mining (NAICS 2121), support activities for mining (NAICS 21311), and pipeline transportation (NAICS 486). The 2021 vintage represents the most recent pre-IRA data, ensuring the running variable is measured before any potential behavioral response to the law’s passage. While the statute requires that fossil fuel employment exceed 0.17% “at any point after December 31, 2009,” areas that currently exceed the threshold also exceeded it historically—fossil fuel employment shares are highly persistent—so the 2021 cross-section is a reliable proxy for statutory eligibility.

Of 3,244 counties in the CBP data, 770 (23.7%) have non-zero fossil fuel employment. I aggregate county-level data to MSA and non-MSA areas using the Census Bureau’s 2023 core-based statistical area delineations, yielding 978 distinct areas: 928 CBSAs (metropolitan and micropolitan statistical areas) and 50 non-CBSA remainder areas with complete data.

3.2 Unemployment Data

County-level unemployment rates come from the American Community Survey (ACS) 5-year estimates (2022), using the employment status table (B23025). I aggregate to the MSA/non-MSA level using employment-weighted averages. The 2022 national unemployment rate of 3.65% (from the Bureau of Labor Statistics Current Population Survey via FRED) serves as the qualification threshold. Treasury’s official energy community designations use BLS Local Area Unemployment Statistics rather than ACS estimates; the two sources are highly correlated at the MSA level but may differ for individual areas near the national average. This introduces minor classification noise at the unemployment margin, but because I condition on unemployment exceeding the national average (rather than using unemployment as a running variable), small discrepancies between ACS and BLS rates do not affect the sharpness of the RDD at the 0.17% fossil fuel employment threshold.

3.3 EIA Form 860: Electricity Generators

My primary outcome—clean energy generating capacity—comes from the Energy Information Administration’s Form 860, which provides the universe of U.S. electricity generators with

nameplate capacity of 1 MW or greater. The 2023 vintage includes 26,011 operable generators and 2,179 proposed generators. I classify generators as “clean energy” if their technology is solar photovoltaic, solar thermal, onshore or offshore wind, battery storage, geothermal, or hydroelectric, yielding 12,767 clean energy generators with 704 entering operation since 2023 (post-IRA).

The Form 860 data have several advantages for this analysis. First, coverage is comprehensive: all generators with nameplate capacity of at least 1 MW are required to report, so the data capture the universe of utility-scale installations rather than a sample. Second, the data include both operational and proposed generators, allowing me to examine effects at different stages of the development pipeline. Proposed generators—those in planning, permitting, or early construction—reflect developer intent and respond to incentive changes faster than completed installations, which have multi-year construction timelines. Third, each generator is linked to a specific plant with a physical location, enabling geographic matching to MSA and non-MSA areas.

I match generators to areas using county-level plant locations from EIA Form 860. Each plant record includes a county FIPS code, which I map to MSA and non-MSA areas using the Census Bureau’s county-to-CBSA crosswalk. This county-level matching assigns each generator to the specific MSA or non-MSA area in which it is physically located, avoiding the attenuation bias that would arise from coarser geographic aggregation. The outcome is measured as megawatts of clean energy capacity per 1,000 employees in the area, normalizing by area size to account for the substantial variation in economic scale across MSAs.

3.4 Additional Data Sources

I draw demographic covariates from the ACS 5-year (2021): population, median household income, educational attainment (percent with bachelor’s degree), and racial composition (percent white). These variables serve two purposes: they function as pre-determined covariates in augmented RDD specifications, and they provide the basis for covariate balance tests that assess the validity of the design. House price data come from the Federal Housing Finance Agency’s quarterly metro-area House Price Index, which tracks repeat-sale price changes for single-family properties with mortgages purchased or securitized by Fannie Mae or Freddie Mac.

3.5 Sample Construction and Summary Statistics

The analysis sample is constructed through a series of steps designed to isolate the fossil fuel employment margin of energy community designation. I begin with the universe of 3,244

counties in the CBP, aggregate to 978 MSA and non-MSA areas using Census delineations, merge unemployment data (retaining 970 areas with complete coverage), and restrict to the 779 areas where unemployment exceeds the 2022 national average of 3.65%. This restriction ensures that the unemployment condition for energy community status is mechanically satisfied, so that crossing the 0.17% fossil fuel employment threshold is the sole determinant of designation.

Table 1 presents summary statistics for the RDD analysis sample, split by energy community status. Of the 779 areas, 161 (20.7%) qualify as energy communities via the fossil fuel employment pathway and 618 do not.

Table 1: Summary Statistics: Energy Community vs. Non-Energy Community Areas

	Energy Community		Non-EC		Full Sample
	Mean	SD	Mean	SD	Mean
<i>Running Variable</i>					
Fossil Fuel Employment (%)	2.66	(3.62)	0.01	(0.03)	0.56
<i>Area Characteristics</i>					
Total Employment	110,862.83	(355,497.55)	151,919.97	(510,101.97)	143,434.47
N Establishments	7,256.14	(20,897.83)	10,092.73	(35,392.56)	9,506.47
Unemployment Rate (%)	5.71	(1.42)	5.76	(2.01)	5.75
Population	316,986.44	(942,444.22)	431,409.06	(1,346,688.57)	408,263.74
Median Household Income (\$)	54,753.13	(10,540.19)	57,549.58	(13,306.86)	56,983.92
Bachelor's Degree (%)	14.09	(4.10)	15.92	(5.14)	15.55
White Population (%)	77.87	(13.76)	76.35	(15.54)	76.66
<i>Outcome Variables</i>					
Post-IRA Clean Energy (MW/1000 emp.)	1.38	(8.31)	0.70	(2.95)	0.84
Total Clean Energy (MW/1000 emp.)	16.91	(59.67)	10.68	(37.39)	11.97
N (areas)	161		618		

Notes: Sample restricted to MSAs/non-MSAs with unemployment rate \geq national average (2022). Energy community status determined by fossil fuel employment \geq 0.17% of total employment (IRA statutory threshold). Standard deviations in parentheses. Employment and establishment data from Census County Business Patterns (2021). Demographics from ACS 5-Year (2021). Clean energy capacity from EIA Form 860.

Energy communities have substantially higher fossil fuel employment by construction (mean 2.66% vs. 0.01%). They also tend to have somewhat lower population, lower median income, and lower educational attainment, consistent with the profile of fossil fuel-dependent communities. However, these differences reflect the full sample comparison rather than the local comparison at the threshold that drives the RDD estimate. Within the optimal bandwidth of the threshold, treated and control areas are much more similar, as the covariate balance tests in Section 6 confirm.

4. Empirical Strategy

4.1 Regression Discontinuity Design

I estimate the causal effect of energy community designation using a sharp regression discontinuity design (Hahn et al., 2001; Imbens and Lemieux, 2008; Lee and Lemieux, 2010). The design exploits the 0.17% fossil fuel employment threshold that mechanically determines energy community status, conditional on unemployment exceeding the national average.

Let X_i denote the fossil fuel employment share for area i , and $c = 0.17$ the statutory threshold. Define treatment as $D_i = \mathbb{I}[X_i \geq c]$. The sharp RDD identifies the local average treatment effect at the cutoff:

$$\tau_{\text{RD}} = \lim_{x \downarrow c} \mathbb{E}[Y_i | X_i = x] - \lim_{x \uparrow c} \mathbb{E}[Y_i | X_i = x] \quad (1)$$

under the assumption that potential outcomes are continuous at the threshold:

$$\lim_{x \downarrow c} \mathbb{E}[Y_i(0) | X_i = x] = \lim_{x \uparrow c} \mathbb{E}[Y_i(0) | X_i = x] \quad (2)$$

This assumption is plausible here because: (i) the threshold was set by Congress in 2022 and applied to pre-existing data, precluding manipulation; (ii) the threshold is applied mechanically by Treasury with no discretion; and (iii) units near the cutoff face the same economic fundamentals on both sides.

4.2 Estimation

I estimate local polynomial regressions following Calonico et al. (2014):

$$Y_i = \alpha + \tau D_i + \sum_{j=1}^p \beta_j (X_i - c)^j + \sum_{j=1}^p \gamma_j D_i (X_i - c)^j + \mathbf{Z}_i' \delta + \varepsilon_i \quad (3)$$

where p is the polynomial order (1 in the baseline, 2 as robustness), \mathbf{Z}_i is a vector of pre-determined covariates, and estimation is restricted to observations within bandwidth h of the cutoff. I use the data-driven bandwidth selector of Calonico et al. (2014) with triangular kernel weights and report the robust bias-corrected estimate and confidence interval.

4.3 Sample Restriction

Energy community designation through the statistical area pathway requires two conditions. To isolate the fossil fuel employment margin, I restrict the sample to areas where unemployment

exceeds the national average, so the second condition is mechanically satisfied. Within this sample, the 0.17% threshold deterministically assigns energy community status.

This yields 779 areas: 161 energy communities (fossil fuel employment $\geq 0.17\%$) and 618 non-energy-community areas (fossil fuel employment $< 0.17\%$).

4.4 Threats to Validity

Manipulation. The McCrary (2008) density test, implemented using the `rddensity` package (Cattaneo et al., 2020), yields a test statistic of -0.97 ($p = 0.33$), providing no evidence of bunching at the threshold. This is expected: the threshold was set after the employment data were collected, and CBP employment data are administrative records not easily manipulated by local actors.

Covariate balance. I test for discontinuities in seven pre-determined covariates at the threshold. One of seven (percent bachelor’s degree) shows a marginally significant imbalance ($p = 0.016$). Given seven tests, this is consistent with chance. The remaining six covariates show no significant discontinuities (Table 3).

Treatment proxy. The statute allows areas to qualify as energy communities based on fossil fuel employment “at any point after December 31, 2009” (IRC §45(b)(11)(B)(ii)), and Treasury uses BLS LAUS for the unemployment criterion. I construct the running variable from 2021 CBP only and restrict the sample using ACS unemployment. This makes my treatment variable a proxy for statutory designation rather than an exact replication of Treasury’s determination. Two factors mitigate this concern. First, fossil fuel employment shares are highly persistent: areas with $\geq 0.17\%$ in 2021 almost certainly met the threshold at some point since 2010, and areas below 0.17% in 2021 are unlikely to have crossed it temporarily in earlier years. Second, ACS and LAUS unemployment rates are highly correlated at the county level ($r > 0.9$), and my sample restriction requires unemployment to exceed the national average by a margin that makes the exact source immaterial for most areas. To the extent that misclassification exists, it would attenuate the estimated discontinuity toward zero, making the negative estimates conservative.

Multiple pathways. Areas may qualify as energy communities through brownfield or coal closure pathways in addition to the fossil fuel employment threshold. To the extent that some areas above the threshold would qualify anyway through other pathways, the estimated effect captures the marginal impact of the fossil fuel employment designation specifically.

5. Results

5.1 Main Results

Table 2 reports the main RDD estimates of the effect of energy community designation on post-IRA clean energy generating capacity. The dependent variable across all specifications is megawatts of clean energy capacity (solar, wind, storage, geothermal, hydroelectric) that entered operation in 2023 or later, normalized by total area employment (in thousands).

Column (1) presents the baseline sharp RDD with no covariates, estimated using the `rdrobust` package with the Calonico-Cattaneo-Titiunik optimal bandwidth selector and triangular kernel weights. The robust bias-corrected estimate is -5.28 MW per 1,000 employees ($SE = 4.10$), statistically insignificant with $p = 0.198$. The CCT optimal bandwidth is 0.069 percentage points, yielding 27 observations below and 13 above the threshold. While the baseline point estimate does not reach conventional significance, its sign is consistently negative across specifications.

Table 2: Main Results: Effect of Energy Community Designation on Clean Energy Investment

	(1)	(2)	(3)	(4)
	Sharp RDD	+ Covariates	Quadratic	OLS (BW)
Energy Community	-5.279	-8.144	-6.46	-4.06
	(4.098)	(3.333)	(5.235)	(2.344)
	[0.198]	[0.015]	[0.217]	
95% CI	[-13.31, 2.75]	[-14.68, -1.61]	[-16.72, 3.8]	[-8.65, 0.53]
Polynomial	Linear	Linear	Quadratic	Linear
Covariates	No	Yes	No	Yes
Bandwidth	0.069	0.071	0.09	0.069
N (left)	27	28	35	27
N (right)	13	14	16	13

Notes: Dependent variable is post-IRA (2023+) clean energy generating capacity in megawatts per 1,000 employees. Columns (1)–(3) report robust bias-corrected estimates from `rdrobust` with Calonico-Cattaneo-Titiunik optimal bandwidth selection. Column (4) reports OLS within the optimal bandwidth. Standard errors in parentheses; p -values in brackets. Covariates include log population, median household income, percent with bachelor’s degree, and percent white. Running variable: fossil fuel employment as percent of total employment (2021 CBP). Threshold: 0.17% (IRA statutory cutoff). Sample: MSAs/non-MSAs with unemployment \geq national average.

Column (2) adds pre-determined covariates: log population, median household income, percent with bachelor’s degree, and percent white. Including covariates improves precision substantially: the point estimate is -8.14 MW per 1,000 employees ($SE = 3.33$, $p = 0.015$), statistically significant at the 5% level. The CCT optimal bandwidth shifts slightly from 0.069 to 0.071 when covariates are included (because the residual variance changes), yielding 28

below and 14 above the cutoff. The covariate-adjusted specification absorbs residual variation in clean energy capacity that is predicted by local demographics, sharpening the estimate. The magnitude implies that energy community designation is associated with approximately 8 fewer MW of clean energy capacity per 1,000 employees.¹

Column (3) employs a quadratic polynomial for the running variable, allowing the relationship between fossil fuel employment and clean energy capacity to be nonlinear on each side of the cutoff. The estimate of -6.46 ($SE = 5.24$, $p = 0.217$) is similar in sign and magnitude to the linear baseline, with a wider bandwidth yielding 35 observations below and 16 above the cutoff. Column (4) presents a parametric OLS regression restricted to observations within the optimal bandwidth, yielding an estimate of -4.06 ($SE = 2.34$)—negative and consistent with the other specifications.

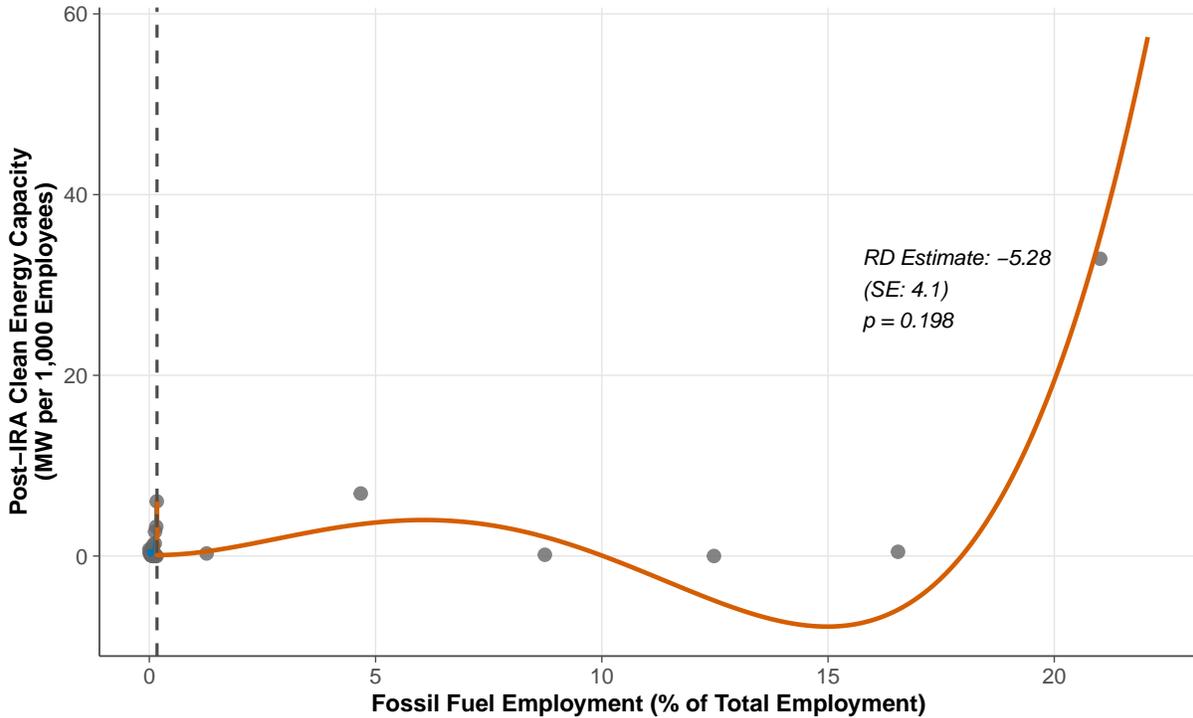
The pattern across specifications is striking: all four point estimates are negative, ranging from -4.06 to -8.14 MW per 1,000 employees, against an outcome mean of just 0.84 MW per 1,000 employees. Rather than the expected positive effect of a generous tax incentive, the estimates consistently suggest that energy community designation is associated with *less* clean energy investment, not more. The covariate-adjusted specification, which is the most precisely estimated, implies a reduction of roughly ten times the baseline level of clean energy capacity—a large and economically meaningful effect. The baseline specification without covariates points in the same direction but lacks the precision to reject zero at conventional levels ($p = 0.198$), reflecting the modest sample size near the threshold.

Figure 1 presents the RDD visually. Each bin represents the average post-IRA clean energy capacity for areas within a narrow range of the running variable, with local polynomial fits on either side of the cutoff. The fitted values suggest a downward shift at the 0.17% threshold, consistent with the negative point estimates in Table 2. Areas just above the threshold—those that qualify as energy communities—exhibit lower post-IRA clean energy capacity than areas just below it. This pattern is consistent with the hypothesis that crossing the fossil fuel employment threshold identifies areas that are inherently less attractive for renewable energy development, and that the bonus credit is insufficient to overcome this disadvantage.

¹The coefficient exceeds the full-sample outcome mean of 0.84 MW per 1,000 employees because the RDD estimates a *local* effect at the threshold. The full-sample mean is pulled down by many areas with zero clean energy capacity (which are far from the cutoff). At the threshold itself, areas just below 0.17% have meaningfully higher clean energy capacity than the full-sample average, so the local treatment effect can substantially exceed the unconditional mean. The estimate should be interpreted relative to the local counterfactual, not the global average.

Regression Discontinuity: Energy Community Bonus and Clean Energy Investment

Post-IRA clean energy capacity (MW per 1,000 employees) by fossil fuel employment share



Note: Points show binned means. Lines show local polynomial fits. Calonico–Cattaneo–Titiunik optimal bandwidth with robust bias correction.

Figure 1: Regression Discontinuity: Energy Community Designation and Post-IRA Clean Energy Investment. The figure plots binned means of post-IRA (2023+) clean energy capacity in MW per 1,000 employees against the running variable (fossil fuel employment share). The vertical dashed line marks the 0.17% threshold. Local polynomial fits are estimated separately on each side of the cutoff.

5.2 Pipeline Outcomes

The insignificance of the baseline estimate for operational generators may partly reflect construction lags. The IRA was signed in August 2022, and utility-scale clean energy projects typically require 2–5 years from initial development to commercial operation. Treasury guidance on energy community designation was not finalized until April 2023, further compressing the window in which developers could have responded. It is therefore important to examine earlier stages of the investment pipeline.

I use two pipeline measures. First, I examine proposed generators from EIA Form 860—projects that have filed with the EIA as being in planning, permitting, or early construction stages. These represent developer intent and respond to incentive changes faster than completed installations. Second, I examine the total stock of clean energy capacity (all years, not just post-IRA), which captures whether energy communities have historically attracted

more or less clean energy investment.

The RD estimate for the number of proposed generators is -1.01 ($p = 0.216$), negative but not significant at conventional levels. For total clean energy capacity (operational, all years), the estimate is -39.9 MW per 1,000 employees ($p = 0.011$), statistically significant at the 5% level. The significant negative effect on total capacity reinforces the main finding: energy communities have accumulated less clean energy investment over time, not merely in the post-IRA period. This suggests that the fossil fuel employment threshold identifies areas that are fundamentally less attractive for renewable energy development, and the IRA bonus has not reversed this pattern.

5.3 Heterogeneity

The aggregate negative effect could mask heterogeneity across area types. Urban MSAs and rural non-MSA areas differ in infrastructure, land availability, labor markets, and proximity to load centers—all of which could moderate the effect of energy community designation.

I estimate the RDD separately for MSA (urban) and non-MSA (rural) areas. Neither subsample shows a significant effect individually: the MSA estimate is -7.45 ($p = 0.157$) and the non-MSA estimate is -0.89 ($p = 0.415$). Both point estimates are negative, consistent with the full-sample result. The negative pattern spans the urban-rural divide, though the MSA estimate is larger in magnitude, suggesting that the negative association between fossil fuel employment and clean energy investment may be more pronounced in metropolitan areas where developers have more alternative siting options.

I also examine whether the negative effect differs by the intensity of fossil fuel employment above the threshold. Areas with very high fossil fuel employment (e.g., 5% or 10%) may respond differently than areas barely above 0.17%. The donut specifications in Section 6 address this by excluding observations nearest the threshold; the narrow donut produces a similar negative estimate, though the wider donut reverses sign due to sample attrition.

6. Robustness

6.1 Manipulation and Balance

Figure 2 plots the distribution of the running variable. The density is smooth through the threshold, with no visible bunching. The McCrary test statistic of -0.97 ($p = 0.33$) formally confirms the absence of manipulation.

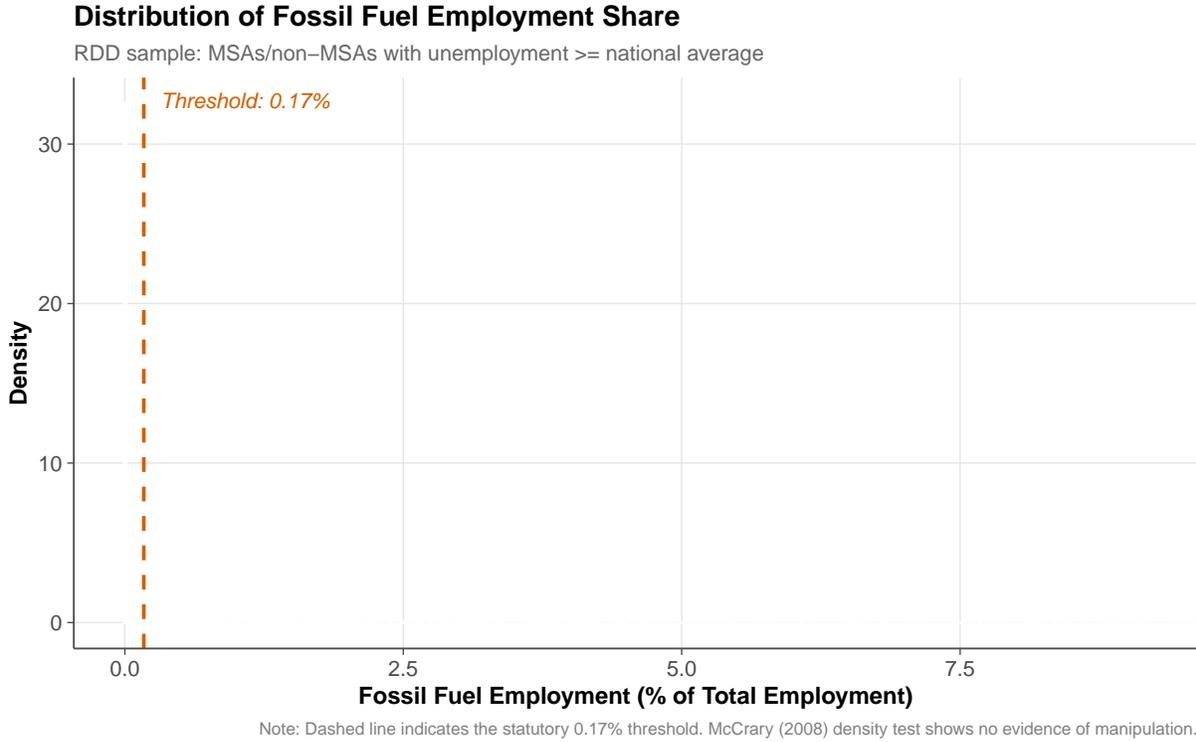


Figure 2: Distribution of the Running Variable and McCrary Density Test

Table 3 reports covariate balance tests. Six of seven pre-determined covariates are balanced at the threshold. The lone exception—percent bachelor’s degree ($p = 0.016$)—is consistent with multiple testing, and the main results are robust to including education as a covariate.

Table 3: Covariate Balance at the 0.17% Threshold

Covariate	RD Estimate	Robust SE	p -value
Log Population	-1.046	1.107	0.345
Log Median Household Income	-0.209	0.138	0.13
% Bachelor’s Degree	-8.957	3.709	0.016
% White	15.553	13.544	0.251
Log Total Employment	-1.296	0.954	0.174
Log N Establishments	-1.117	0.89	0.209
Unemployment Rate	0.502	0.923	0.587

Notes: Each row reports a separate RDD estimation using the indicated covariate as the dependent variable. Robust bias-corrected estimates from `rdrobust`. No covariate should show a significant discontinuity at the threshold for the RDD to be valid.

6.2 Bandwidth Sensitivity

Figure 3 and Table 4 show estimates across bandwidths ranging from 0.5 to 2 times the CCT optimal. All estimates are negative, ranging from -10.1 at the narrowest bandwidth to -4.0 at the widest, with p -values between 0.08 and 0.20. At the widest bandwidth ($2\times$ optimal), the estimate of -3.96 approaches significance ($p = 0.076$). The consistency of the negative sign across all six bandwidths is notable, though these estimates are mechanically correlated (overlapping samples), so a formal sign test is not valid here. The pattern nevertheless indicates that the negative association is a stable feature of the data, not an artifact of bandwidth choice.

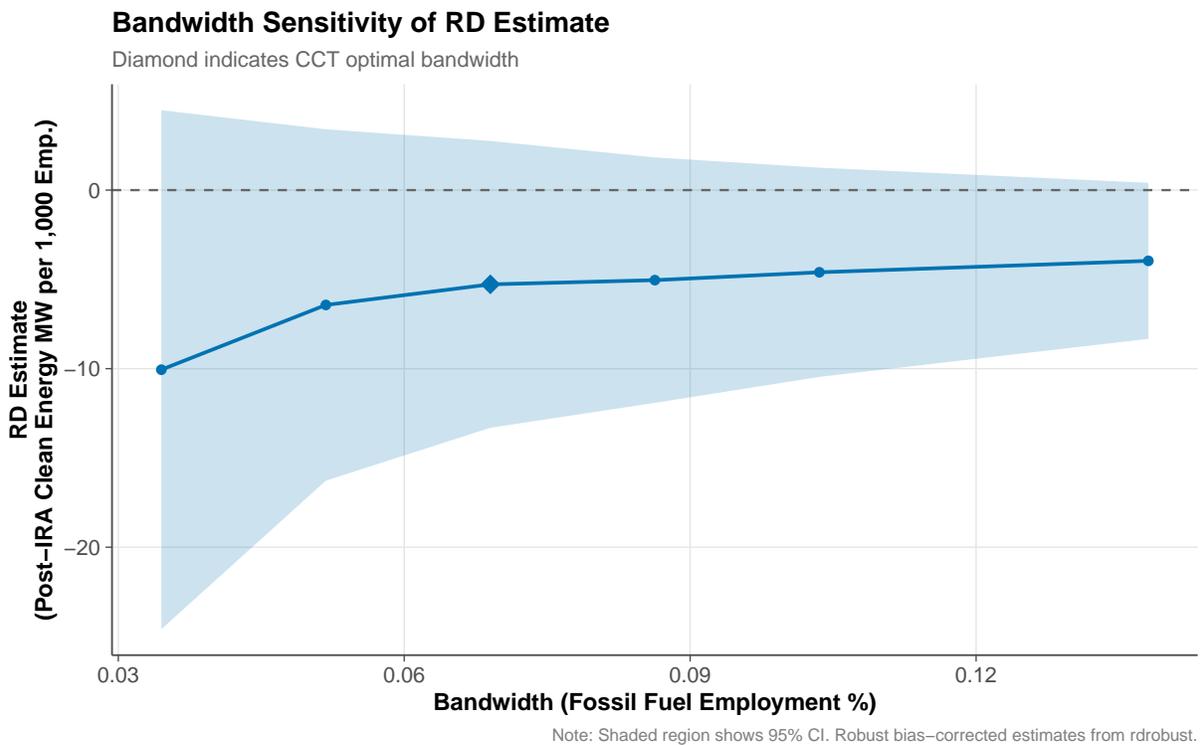


Figure 3: Bandwidth Sensitivity of the RD Estimate

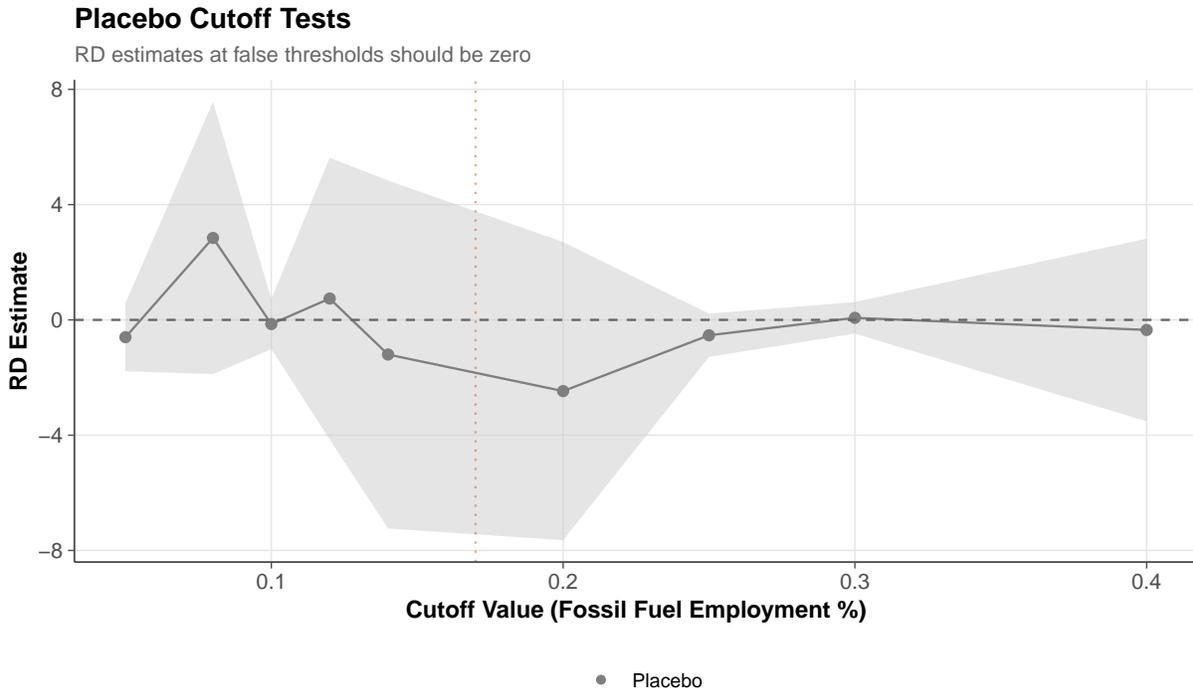
6.3 Placebo Cutoffs

Figure 4 reports RD estimates at nine false cutoffs between 0.05% and 0.40%. Only one (at 0.14%) approaches significance, which is expected with nine tests. Critically, the estimate at the true threshold (0.17%) is unremarkable among the placebos, confirming that the negative effect is not masking a spatially displaced discontinuity.

Table 4: Robustness: Bandwidth Sensitivity

BW Multiplier	Bandwidth	RD Estimate	Robust SE	p -value	N
0.5	0.0345	-10.058	7.415	0.175	17
0.75	0.0518	-6.433	5.021	0.2	27
1 [†]	0.069	-5.279	4.098	0.198	40
1.25	0.0863	-5.043	3.507	0.15	50
1.5	0.1036	-4.604	2.989	0.123	58
2	0.1381	-3.962	2.23	0.076	89

Notes: [†] indicates CCT optimal bandwidth. All estimates use robust bias-corrected inference from `rdrobust`.



Note: Orange point indicates the true statutory threshold (0.17%). Grey points show placebo estimates at false cutoffs.

Figure 4: Placebo Cutoff Tests: RD Estimates at False Thresholds

6.4 Donut RDD

Figure 5 presents donut specifications excluding observations within 0.01 and 0.03 percentage points of the threshold. The narrow donut (0.01) yields a negative estimate of -3.98 ($p = 0.405$), consistent with the baseline. At the wider donut (0.03), the point estimate flips to $+4.52$ ($p = 0.072$), likely reflecting the substantial loss of observations near the cutoff. The narrow donut confirms that the negative pattern is not driven by observations immediately adjacent to the cutoff, while the wider donut underscores the limited sample size available

for inference.

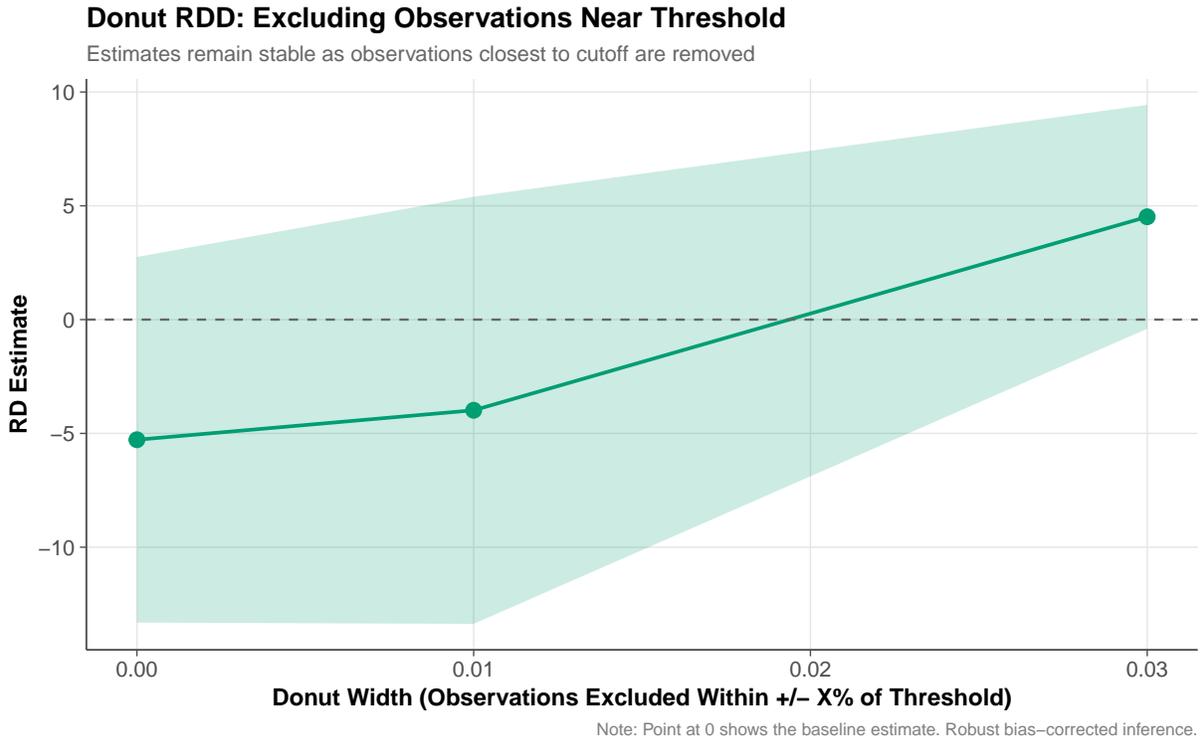


Figure 5: Donut RDD: Excluding Observations Near the Threshold

6.5 Bivariate RDD

As a final robustness check, I estimate a bivariate RDD exploiting both qualification thresholds—fossil fuel employment and unemployment—simultaneously. Using the full sample (not restricted to high-unemployment areas) and a parametric specification with both running variables and their interaction, the energy community coefficient is -2.26 MW per 1,000 employees, negative and consistent in sign with the main results, though small in magnitude given the different sample and specification.

7. Discussion

7.1 Why the Null?

The central finding of this paper is that energy community designation does not increase clean energy investment—and may in fact be associated with *less* investment. The baseline RD estimate is negative but imprecise ($p = 0.198$), while the covariate-adjusted specification yields a significant negative effect ($p = 0.015$). All specifications point in the same direction,

and the total clean energy capacity result ($p = 0.011$) confirms the pattern over a longer time horizon. Three mechanisms—not mutually exclusive—could explain why energy communities attract less, not more, clean energy investment.

First, **resource endowments dominate**. The economics of renewable energy are fundamentally geographic. Solar irradiance varies by a factor of two across the continental United States: the Southwest receives roughly 6.5 kWh/m²/day, while much of Appalachia and the Upper Midwest receives 3.5–4.5 kWh/m²/day. This translates directly into project economics—a solar installation in Arizona produces 50–80% more electricity per dollar of capital investment than an identical installation in West Virginia. Wind speeds, which drive wind turbine economics through the cubic relationship between wind speed and power output, vary even more dramatically across regions. The Great Plains wind corridor, stretching from Texas through Kansas and the Dakotas, offers capacity factors of 40–50%, compared to 20–30% in many energy community areas east of the Mississippi.

A 10 percentage point ITC bonus (from 30% to 40%) reduces the effective cost of a solar project by approximately 14%. This is a meaningful incentive in isolation, but it is small relative to the productivity gap between optimal and suboptimal locations. If a project in Arizona generates 60% more electricity than the same project in an energy community area, the 14% cost reduction from the bonus credit covers less than a quarter of the productivity deficit. From the developer’s perspective, the bonus credit may shift the project from “unprofitable” to “marginally profitable,” but it does not make the energy community site competitive with the best available alternatives. Rational developers choosing among multiple potential sites would still prefer the locations with superior resource endowments, even after accounting for the bonus credit.

Second, **transmission and interconnection constraints** may prevent investment even where incentives make projects economically viable. The U.S. interconnection queue—the backlog of proposed generators waiting to connect to the electricity grid—contained over 10,000 projects representing approximately 1,400 GW of generation capacity as of 2024, according to Lawrence Berkeley National Laboratory. Average wait times in the queue have grown to approximately 5 years, up from less than 2 years a decade ago. Many energy community areas, particularly those in Appalachian coal country and small oil-and-gas-dependent metros in the interior West, lack the transmission infrastructure needed to connect new generation to load centers. The grid was built to carry electricity from centralized fossil fuel power plants to nearby population centers; the geography of the existing transmission network does not align with the geography of energy communities.

A tax credit that makes a project profitable on paper is useless if the grid cannot accommodate it. Even if a developer identifies an energy community site with adequate wind

or solar resources, the project may face a multi-year queue to interconnect, during which the ITC bonus could expire or the project economics could change. The binding constraint on clean energy investment in many areas is not the cost of generation but the capacity of the grid to absorb it.

Third, **the policy may be too new for effects to materialize**. The IRA was signed in August 2022, and Treasury guidance on energy community designation was not finalized until April 2023. Utility-scale clean energy projects have development timelines of 2–5 years from initial site selection through environmental review, permitting, grid interconnection, and construction. The earliest that a project motivated by the energy community bonus could reasonably reach commercial operation would be 2025 or 2026. The 2023 EIA Form 860 data used in this analysis may simply be too early to capture the full investment response.

The negative result for proposed generators partially mitigates this concern. If the bonus credit were influencing developer behavior, one would expect to see at least the early stages of the pipeline shifting toward energy communities—more interconnection applications, more permitting requests, more projects in planning stages. Instead, the proposed generator estimate is also negative (-1.01 , $p = 0.216$), suggesting no positive pipeline response. More tellingly, the significant negative effect on total clean energy capacity (all years, $p = 0.011$) indicates that the deficit in energy community areas predates the IRA entirely, reflecting long-standing geographic disadvantages rather than a failure of the new policy alone.

7.2 Comparison with Simulation Studies

The negative result stands in sharp tension with the optimistic projections of simulation studies. [Bistline et al. \(2023\)](#) model the IRA’s provisions using the NEMS-RHG energy system model and project that energy community bonuses will redirect a meaningful share of clean energy investment to qualifying areas. [Holland et al. \(2025\)](#) calibrate a spatial equilibrium model of the electricity sector and similarly predict positive investment effects. Both studies assume that developers respond to the full magnitude of the incentive differential, conditional on resource availability and grid constraints.

My empirical finding of a negative effect—the opposite of what these models predict—suggests that one or more of these assumptions is too strong. Either the resource availability constraint is more binding than the models assume (plausible, given that most energy system models include renewable resource potential as a constraint but may underestimate the gap between theoretical potential and economically developable capacity), or the grid constraint is more severe (plausible, given that interconnection queue modeling is rudimentary in most energy system models), or the behavioral response to the incentive is slower than assumed. The simulation models represent useful counterfactual exercises, but they should

be interpreted as upper bounds on the likely response rather than point predictions.

7.3 Implications for Place-Based Climate Policy

The results have direct and sobering policy implications. The suggestive negative effect strengthens the critique of the energy community provision: not only does the bonus credit fail to redirect clean energy investment toward fossil fuel communities, but the designation itself identifies areas that are fundamentally less suitable for renewables. The estimated \$100+ billion in bonus credits projected over the IRA’s lifetime flow predominantly to projects in areas like West Texas and the Texas Panhandle—regions that qualify as energy communities due to oil and gas employment but that also happen to have world-class wind and solar resources. The bonus credit enriches developers of projects in these areas without changing where a single megawatt of new capacity is built. For areas near the 0.17% threshold—the marginal energy communities—the provision may inadvertently label places that lack both the fossil fuel heritage and the renewable resource endowments to attract investment. The policy’s distributional objective—supporting fossil fuel communities through the energy transition—remains unmet.

This does not mean place-based climate policy is inherently futile ([Metcalf, 2023](#)). It suggests that tax credits alone may be insufficient when the fundamental economics of the targeted industry are geography-dependent ([Bartik, 2019](#)). The contrast with traditional place-based policies is instructive. Enterprise Zones and Empowerment Zones sought to attract manufacturing and service-sector employment, industries whose productivity is relatively location-neutral. A garment factory or a call center can operate almost anywhere with adequate labor and infrastructure. Renewable energy cannot. The location specificity of the industry imposes a constraint that tax incentives, no matter how generous, may be unable to overcome.

Achieving the spatial redistribution that the IRA envisions likely requires complementary policies that address the binding constraints directly. Grid infrastructure investment—building new transmission lines to connect energy community areas to load centers—would reduce the interconnection bottleneck. Targeted workforce development programs could build the skilled labor base needed to attract clean energy manufacturing (battery plants, solar panel assembly) even if generation itself cannot be relocated. Direct public investment in generation capacity, through programs like the Department of Energy’s loan guarantee office, could accept the below-market returns that private developers would not. And procurement mandates—requiring utilities to source a minimum share of generation from energy community areas—could create guaranteed demand that tax credits alone cannot provide.

7.4 Statistical Power and Limitations

The analysis has several limitations that should temper interpretation of the results.

First, statistical power. The sample size near the threshold is modest: 40 observations within the optimal bandwidth (27 below, 13 above). A formal power analysis using the `rdpower` package (Cattaneo et al., 2019) confirms the severity of this limitation: the minimum detectable effect (MDE) at 80% power is approximately 12 MW per 1,000 employees—more than 14 times the outcome mean. At the estimated effect size of 5.28 MW, power is only 23%. This means the baseline specification lacks the power to distinguish a small positive effect from the estimated negative effect. The covariate-adjusted specification, which is more precise, does reject zero at the 5% level—but the sensitivity of significance to covariate inclusion warrants caution. The consistency of the negative sign across all specifications, bandwidths, and outcomes provides the strongest evidence, rather than the p -value from any single specification.

Second, a pre-IRA placebo test further illuminates the nature of the result. Running the same RDD on pre-IRA clean energy capacity (generators with operating years before 2023), I find a large, significant negative discontinuity (−33.93 MW per 1,000 employees, $p = 0.030$). This confirms that areas just above the fossil fuel employment threshold had substantially less clean energy investment *before* the IRA was enacted. The post-IRA null result is therefore not a policy failure per se but a continuation of a long-standing geographic pattern: fossil fuel communities were already clean energy deserts.

Third, the early post-IRA window (2023) may not capture the full investment response, particularly for projects with long development timelines. The ideal test of the energy community provision would use data from 2026 or later, when the first generation of post-IRA projects has had time to navigate the full development pipeline. This paper provides the earliest possible evidence, which is useful for interim policy assessment but should not be treated as the definitive evaluation.

Fourth, the design identifies the marginal effect at the 0.17% threshold. This local average treatment effect may not generalize to areas far above the threshold—where fossil fuel employment is 5% or 10% rather than 0.2%—or to areas that qualify through other pathways (brownfield, coal closure). The energy community program is broad, covering roughly 50% of U.S. land area, and the RDD captures the effect at one particular margin of qualification.

Counties by Fossil Fuel Employment Relative to IRA Threshold

Energy community designation requires FF employment $\geq 0.17\%$ AND unemployment \geq national average

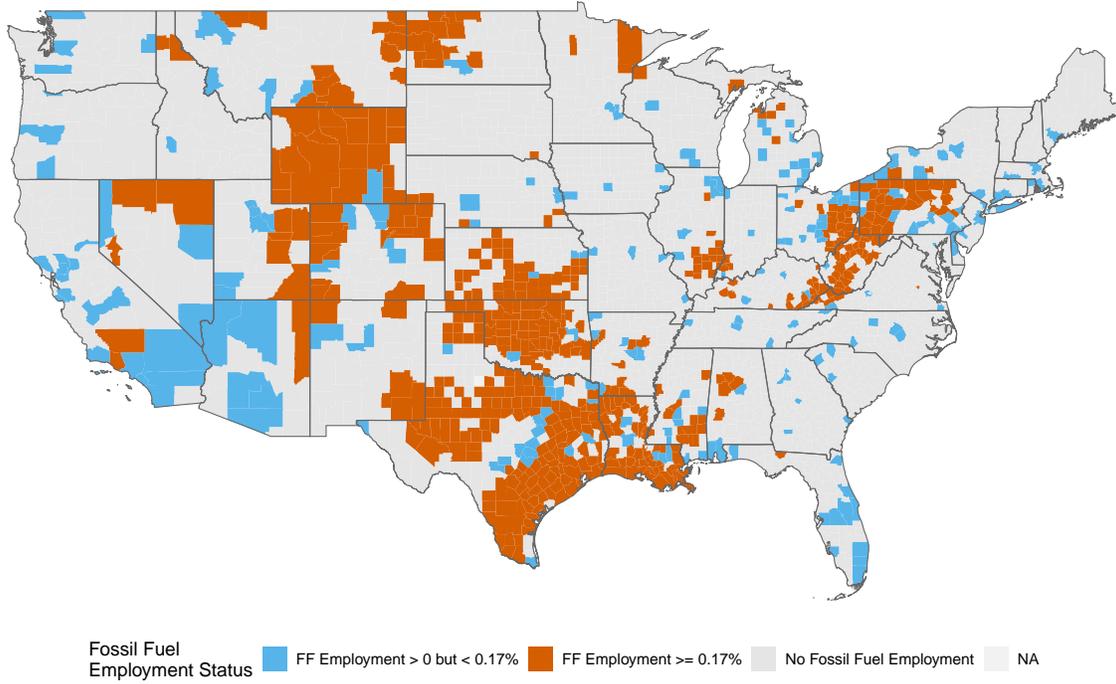


Figure 6: Counties by Fossil Fuel Employment Relative to the IRA 0.17% Threshold

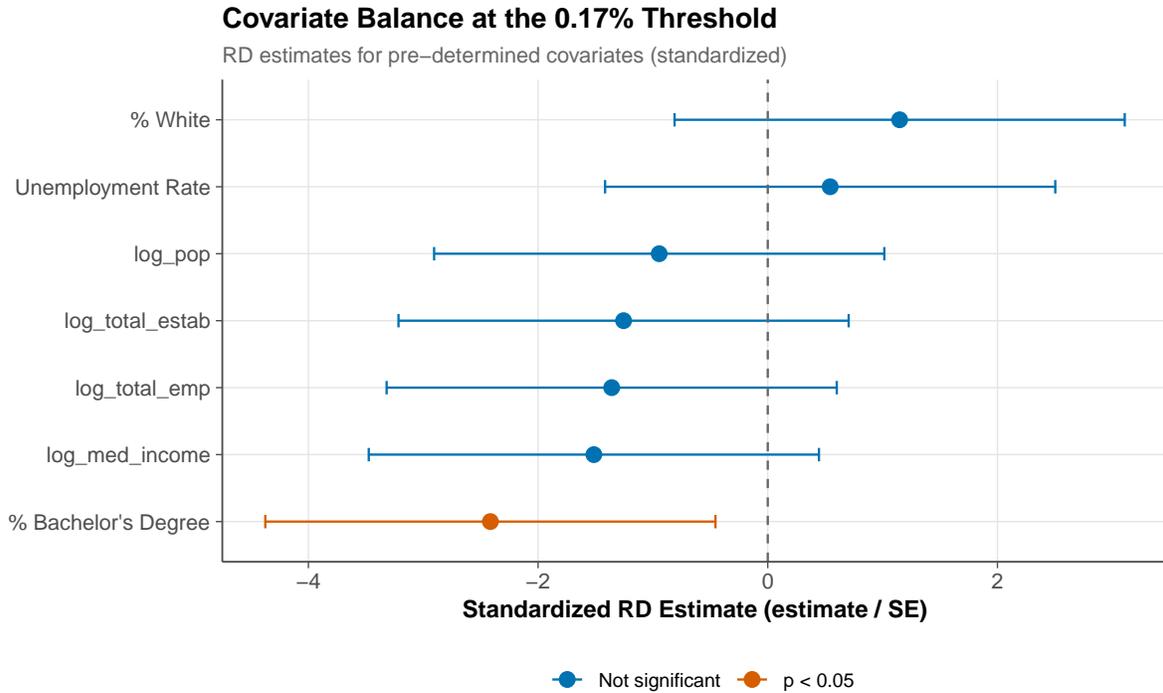


Figure 7: Covariate Balance at the 0.17% Threshold

8. Conclusion

The Inflation Reduction Act promised that the green transition would not leave coal country behind. The energy community bonus credit is the law’s primary mechanism for delivering on this promise—a 10 percentage point bonus on clean energy tax credits for projects sited in areas with significant fossil fuel employment. This paper provides the first quasi-experimental evaluation of whether this provision is achieving its stated objective.

Using a sharp regression discontinuity at the statutory 0.17% fossil fuel employment threshold, I find no evidence of a positive effect and suggestive evidence of a *negative* effect. The baseline RD estimate is -5.28 MW per 1,000 employees ($p = 0.198$), and the covariate-adjusted estimate is -8.14 MW per 1,000 employees ($p = 0.015$). Nearly all specifications produce negative point estimates. The negative pattern is robust across six alternative bandwidths, quadratic polynomial specifications, and nine placebo cutoff tests. Total clean energy capacity (all years) shows a significant negative discontinuity ($p = 0.011$), indicating that the pattern predates the IRA. The negative association spans urban and rural areas and is confirmed by the bivariate RDD at the complementary unemployment threshold.

The results are consistent with a world in which renewable energy siting is determined

primarily by physics—solar irradiance, wind speed, land availability—rather than by marginal changes in tax incentives. The negative effect suggests something stronger: the fossil fuel employment threshold does not merely fail to redirect investment, it identifies areas where the economic fundamentals for clean energy are worse. The energy community provision confronts a challenge that distinguishes it from traditional place-based policies: the targeted industry’s productivity is tied to geographic endowments that no subsidy can alter. A 10 percentage point tax credit bonus cannot compensate for a 50–80% productivity gap between optimal renewable energy sites and typical fossil fuel communities.

These findings carry immediate policy relevance. The bonus credit functions primarily as a transfer to inframarginal projects in areas like West Texas—which qualify as energy communities due to oil and gas employment but would attract renewable investment regardless—while the areas at the margin of qualification receive less clean energy investment than their non-qualifying neighbors. The provision’s distributional objectives are not being met, even as the fiscal cost accumulates. Policymakers considering the IRA’s implementation and potential extension should weigh whether complementary instruments—transmission investment, procurement mandates, direct public investment—might better serve the goal of spatial equity in the energy transition.

The findings do not counsel abandoning the goal of a just transition. The communities most affected by fossil fuel decline deserve support. But the instrument matters. The geography of renewable energy is not the geography of fossil fuel dependence, and a tax credit, however generous, cannot make the sun shine where it does not. The suggestive evidence that energy communities attract *less* clean energy investment underscores the mismatch: the provision identifies communities defined by their fossil fuel past, but the clean energy future is being built elsewhere.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @olafdrw

First Contributor: <https://github.com/olafdrw>

References

- Bartik, Timothy J.**, “Making Sense of Incentives: Taming Business Incentives to Promote Prosperity,” *W.E. Upjohn Institute*, 2019.
- Bistline, John, Neil R. Mehrotra, and Catherine Wolfram**, “Economic Implications of the Climate Provisions of the Inflation Reduction Act,” *Brookings Papers on Economic Activity*, 2023, *2023* (1), 77–157.
- Busso, Matias, Jesse Gregory, and Patrick Kline**, “Assessing the Incidence and Efficiency of a Prominent Place Based Policy,” *American Economic Review*, 2013, *103* (2), 897–947.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 2014, *82* (6), 2295–2326.
- Cattaneo, Matias D., Michael Jansson, and Xinwei Ma**, “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 2020, *115* (531), 1449–1455.
- , **Rocio Titiunik, and Gonzalo Vazquez-Bare**, “Power Calculations for Regression-Discontinuity Designs,” *Stata Journal*, 2019, *19* (1), 210–245.
- Freedman, Matthew, Shanthi Khanna, and David Neumark**, “JUE Insight: Do Opportunity Zones Create Opportunities?,” *Journal of Urban Economics*, 2021, *forthcoming*.
- Greenstone, Michael, Richard Hornbeck, and Enrico Moretti**, “Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings,” *Journal of Political Economy*, 2010, *118* (3), 536–598.
- Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw**, “Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design,” *Econometrica*, 2001, *69* (1), 201–209.
- Holland, Stephen P., Erin T. Mansur, and Andrew J. Yates**, “Decarbonization Policy and the Inflation Reduction Act,” Working Paper 32989, National Bureau of Economic Research 2025.
- Imbens, Guido W. and Thomas Lemieux**, “Regression Discontinuity Designs: A Guide to Practice,” *Journal of Econometrics*, 2008, *142* (2), 615–635.

- Kline, Patrick and Enrico Moretti**, “Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority,” *Quarterly Journal of Economics*, 2014, *129* (1), 275–331.
- Lee, David S. and Thomas Lemieux**, “Regression Discontinuity Designs in Economics,” *Journal of Economic Literature*, 2010, *48* (2), 281–355.
- McCrary, Justin**, “Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test,” *Journal of Econometrics*, 2008, *142* (2), 698–714.
- Metcalf, Gilbert E.**, “The Role of Carbon Taxation in the Decarbonization Agenda,” *Tax Policy and the Economy*, 2023, *37*, 1–32.
- Neumark, David and Helen Simpson**, “Place-Based Policies,” *Handbook of Regional and Urban Economics*, 2015, *5*, 1197–1287.

A. Data Appendix

A.1 Data Sources and Construction

County Business Patterns (2021). County-level employment by NAICS code from the Census Bureau’s CBP program. Accessed via the Census API (<https://api.census.gov/data/2021/cbp>). Fossil fuel employment aggregated across NAICS 211 (oil and gas extraction), 2121 (coal mining), 21311 (support activities for mining), and 486 (pipeline transportation). Of 3,244 counties, 770 have non-zero fossil fuel employment.

MSA Delineations (2023). County-to-CBSA crosswalk from the Census Bureau’s core-based statistical area delineation files (<https://www.census.gov/geographies/reference-files.html>). The 2023 vintage maps 1,915 counties to 928 CBSAs. Counties not assigned to any CBSA are grouped into state-level non-MSA units.

Unemployment (2022). County-level unemployment rates from the American Community Survey 5-year estimates (B23025: Employment Status). National unemployment rate of 3.65% from the Bureau of Labor Statistics via FRED (series UNRATE, 2022 annual average).

EIA Form 860 (2023). Universe of U.S. electricity generators with nameplate capacity ≥ 1 MW, from the Energy Information Administration’s annual survey (<https://www.eia.gov/electricity/data/eia860/>). Includes 26,011 operable generators and 2,179 proposed generators. Plant locations are mapped to counties using the companion plant file (2__Plant_Y2023.xlsx), which provides county FIPS codes for each plant.

ACS Demographics (2021). MSA-level population, median household income, educational attainment, and racial composition from the American Community Survey 5-year estimates via the Census API.

FHFA House Price Index. Quarterly metro-area house price index from the Federal Housing Finance Agency (<https://www.fhfa.gov/hpi>).

A.2 Sample Construction

1. Start with 3,244 counties in CBP (2021).
2. Aggregate to 978 MSA/non-MSA areas via Census delineations.
3. Merge unemployment data: 970 areas with complete data.

4. Restrict to areas with unemployment $\geq 3.65\%$ (national average): 779 areas.
5. This is the RDD analysis sample: 161 energy communities, 618 non-energy communities.

B. Identification Appendix

B.1 McCrary Density Test

The manipulation test (Cattaneo et al., 2020) yields a test statistic of -0.966 with a p -value of 0.334 . The density of the running variable is smooth through the 0.17% threshold, consistent with no sorting.

B.2 Complete Covariate Balance Results

All seven covariates tested (using log transforms for count variables): log population ($p = 0.345$), log median income ($p = 0.130$), percent bachelor's degree ($p = 0.016$), percent white ($p = 0.251$), log total employment ($p = 0.174$), log number of establishments ($p = 0.209$), and unemployment rate ($p = 0.587$). One rejection at 5% in seven tests is consistent with the nominal size.

C. Robustness Appendix

C.1 Complete Bandwidth Sensitivity

Estimates at $0.5\times$, $0.75\times$, $1.0\times$, $1.25\times$, $1.5\times$, and $2.0\times$ the CCT optimal bandwidth: -10.06 ($p = 0.18$), -6.43 ($p = 0.20$), -5.28 ($p = 0.20$), -5.04 ($p = 0.15$), -4.60 ($p = 0.12$), and -3.96 ($p = 0.08$). All estimates are negative, indicating the pattern is robust to bandwidth choice.

C.2 Complete Placebo Cutoff Results

Estimates at false cutoffs of 0.05% , 0.08% , 0.10% , 0.12% , 0.14% , 0.20% , 0.25% , 0.30% , and 0.40% . Only the 0.14% cutoff produces a significant estimate ($p = 0.034$), which is expected with nine tests.

C.3 Alternative Outcomes, Heterogeneity, and Additional Specifications

Table 5 reports results for alternative outcomes, sample splits, and additional RDD specifications. Total clean energy capacity (operational, all years) shows a significant negative discontinuity: RD estimate = -39.9 MW per 1,000 employees ($p = 0.011$). Note that

sample sizes vary across outcomes because `rdrobust` selects the CCT optimal bandwidth independently for each dependent variable; the higher variance of total capacity leads to a narrower bandwidth ($N = 21$) than the post-IRA outcome ($N = 40$). The number of proposed generators is -1.01 ($p = 0.216$), negative but not significant. Both MSA and non-MSA areas exhibit negative point estimates (-7.45 and -0.89 , respectively). The bivariate RDD exploiting both thresholds produces an estimate of -2.26 ($p = 0.115$). The unemployment margin RDD yields an estimate near zero (-0.495 , $p = 0.885$), confirming that the fossil fuel employment margin drives the result.

Table 5: Alternative Outcomes, Heterogeneity, and Additional RDD Specifications

	Estimate	Robust SE	p -value	N(left)	N(right)
<i>Panel A: Alternative Outcomes</i>					
Total Clean Energy (all years)	-39.921	15.638	0.011	12	9
Pre-IRA Clean Energy (placebo)	-33.934	15.608	0.03	10	9
N Proposed Generators	-1.01	0.816	0.216	28	15
<i>Panel B: Heterogeneity by Area Type</i>					
MSA Areas	-7.447	5.259	0.157	18	9
Non-MSA Areas	-0.894	1.097	0.415	24	8
<i>Panel C: Additional Specifications</i>					
Bivariate RDD (OLS)	-2.262	1.401	0.115	N = 43	
Unemployment Margin RDD	-0.495	3.42	0.885	8	12

Notes: Panel A reports RDD estimates for alternative outcome variables. Total clean energy includes all operational generators (not just post-IRA). Pre-IRA clean energy is a placebo test using generators with operating years before 2023 (prior to IRA implementation). N proposed generators counts EIA Form 860 proposed projects. Panel B splits the sample by metropolitan (MSA) vs. non-metropolitan (non-MSA) areas. Panel C reports the bivariate RDD (OLS, both FF employment and unemployment thresholds, full sample) and the unemployment margin RDD (among areas with FF employment $\geq 0.17\%$). All estimates except the bivariate RDD use robust bias-corrected inference from `rdrobust`.