Clean Air, Dirty Power? NAAQS Nonattainment and the Clean Energy Transition

APEP Autonomous Research* @ai1scl

March 4, 2026

Abstract

Does air quality regulation accelerate the transition from fossil fuels to clean energy? I exploit the sharp regulatory threshold created by the National Ambient Air Quality Standards (NAAQS) for PM2.5, where counties exceeding the standard face stringent New Source Review requirements for fossil fuel plants while renewable generators remain exempt. Using a cross-sectional regression discontinuity design across 702 counties with PM2.5 monitors, I find no statistically significant discontinuity in county-level fossil fuel or renewable energy capacity at the 12 μg/m^3 threshold. However, only 11 counties exceed the threshold, yielding effective samples of 8–36 counties and minimum detectable effects exceeding 800% of the outcome mean — the design lacks power to detect economically meaningful effects. McCrary density tests and covariate balance checks support the validity of the design. While these null results are consistent with spatial displacement through regional electricity markets, the severe power limitations preclude strong causal claims.

JEL Codes: Q42, Q52, Q53, R11

Keywords: air quality regulation, energy transition, regression discontinuity, NAAQS, nonattainment, renewable energy

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

The United States is in the midst of an unprecedented energy transformation. Between 2010 and 2023, coal's share of electricity generation fell from 45% to 16%, while wind and solar rose from 3% to 16%. A central question for policymakers is whether environmental regulation accelerates or merely accompanies this transition. The Clean Air Act's National Ambient Air Quality Standards (NAAQS) create one of the sharpest regulatory discontinuities in American environmental policy: counties that exceed a pollution threshold face dramatically higher costs for new fossil fuel power plants, while renewable energy projects remain entirely exempt.

This paper asks whether NAAQS nonattainment designation — the regulatory consequence of exceeding ambient PM2.5 standards — shifts local energy infrastructure away from fossil fuels and toward clean energy. The question matters for two reasons. First, as EPA tightened the PM2.5 standard from 15 to 12 μg/m^3 in 2012 and again to 9 μg/m^3 in 2024, millions of additional Americans now live in counties that will face nonattainment designations (**?**). Understanding whether this regulatory status actually reshapes energy investment informs the welfare calculus of standard-setting. Second, if nonattainment designation merely reallocates fossil generation across county boundaries without reducing it, the environmental benefits of the NAAQS may be partly illusory.

The identification strategy exploits the sharp threshold in EPA's nonattainment designation process. A county's PM2.5 "design value" — the three-year average of annual mean concentrations — is compared to the NAAQS standard. Counties with design values above the threshold are designated nonattainment and face New Source Review (NSR) requirements: any new or substantially modified fossil fuel source must install Lowest Achievable Emission Rate (LAER) technology, obtain pollution offsets, and demonstrate that no alternative sites exist in attainment areas. These requirements impose substantial costs — the permitting process alone can take years and cost millions of dollars (**?**). Crucially, renewable generators emit zero criteria pollutants and are exempt from NSR regardless of attainment status, creating an asymmetric cost wedge.

I implement a cross-sectional sharp regression discontinuity design (RDD) using the distance between each county's average PM2.5 design value (2012–2022) and the 12 μg/m^3 NAAQS standard as the running variable, with energy infrastructure measured from the eGRID 2022 database. The cross-sectional design matches the timing of the running variable to the outcome measurement, treating the 2022 capacity stock as the cumulative result of investment decisions made under the regulatory regime. I test the identifying assumption through McCrary density tests (**?**), which reveal no evidence of manipulation at the cutoff

($p = 0.78$), and through covariate balance tests on population and income, neither of which shows a significant discontinuity ($p = 0.62$ and $p = 0.17$, respectively), though the small number of treated counties ($N_{\text{right}} = 6$) limits the power of these balance tests.

The main finding is a null: nonattainment designation does not produce a statistically significant discontinuity in fossil fuel capacity, renewable capacity, coal capacity, or total generation capacity at the county level. The cross-sectional RDD estimate for fossil fuel capacity is −1,936 MW (SE = 1,889), with a p-value of 0.31 and an effective sample of 36 counties. A placebo test using renewable capacity as an outcome — which should show no effect since renewables are exempt from NSR — confirms the null ($p = 0.28$), lending credibility to the design rather than suggesting confounding. The small effective samples near the threshold limit statistical power, but the results are consistent across specifications.

These null results are robust across specifications. Bandwidth sensitivity analysis shows all estimates remaining insignificant across 50–200% of the MSE-optimal bandwidth. Placebo cutoff tests below the threshold find no systematic discontinuities at false thresholds. Alternative kernel functions yield qualitatively identical conclusions. I also examine the pre-2012 standard at 15 μg/m^3 in a multi-cutoff framework using panel data, finding no significant effects on fossil, renewable, or coal capacity.

This paper contributes to three literatures. First, it extends the large body of work on the economic effects of the Clean Air Act. Seminal papers by **?** and **?** documented that nonattainment designation reduces manufacturing activity and employment. **?** quantified the labor market adjustment costs. **?** showed effects on housing values. However, none of these studies examined whether nonattainment designation affects the composition of energy infrastructure — the sector most directly implicated in the pollution that triggers the regulation itself. This paper fills that gap.

Second, the paper contributes to the growing literature on the determinants of the clean energy transition. While extensive research examines the effects of renewable portfolio standards, tax credits, and carbon pricing on energy investment (**??**), far less is known about whether criteria pollutant regulation — which was not designed to promote clean energy but could do so through asymmetric cost incidence — plays a role. My null finding suggests it does not, at least at the local level.

Third, the paper speaks to the broader "pollution haven" literature initiated by **?**. If nonattainment designation deters fossil investment in regulated counties, that investment may relocate to nearby attainment counties rather than being replaced by renewables. The null local effect is consistent with this spatial displacement channel dominating any within-county substitution toward clean energy. Electricity markets operate at the regional level through balancing authorities and independent system operators, making it natural for generation investment to respond to regional rather than county-level regulatory conditions.

The remainder of this paper proceeds as follows. **??** describes the institutional background of NAAQS nonattainment and its implications for energy infrastructure. **??** presents the data sources and construction of the analysis sample. **??** details the empirical strategy. **??** presents the main results and robustness checks. **??** discusses mechanisms and implications. **??** concludes.

2. Institutional Background and Policy Setting

2.1 The National Ambient Air Quality Standards

The Clean Air Act requires EPA to set NAAQS for six criteria pollutants, including particulate matter (PM2.5). The standard is set at a level "requisite to protect the public health" with "an adequate margin of safety" (42 U.S.C. §7409). EPA reviews the standards every five years, though the actual revision cycle has historically been longer.

For PM2.5, the annual standard has been revised three times. The original 1997 standard was set at 15 μg/m^3 (annual mean). In 2012, EPA tightened it to 12 μg/m^3, a change that brought approximately 35 additional counties into nonattainment. In February 2024, EPA further tightened the standard to 9 μg/m^3, which is projected to affect hundreds of additional counties once designations are finalized (**?**).

The designation process works as follows. EPA compares each county's PM2.5 "design value" — defined as the three-year average of annual mean PM2.5 concentrations across monitoring stations — to the NAAQS standard. Counties with design values exceeding the standard are designated "nonattainment." The designation is effectively a sharp function of the design value: counties above the threshold receive the nonattainment label, and those at or below it remain in "attainment" or "unclassifiable/attainment."

2.2 Regulatory Consequences of Nonattainment

Nonattainment designation triggers a cascade of regulatory requirements, most importantly New Source Review (NSR). Under NSR, any new or substantially modified stationary source of criteria pollutants must:

1. Install **Lowest Achievable Emission Rate (LAER)** technology — the most stringent emission controls achievable, regardless of cost.

2. Obtain **emission offsets** — purchase or retire existing emissions from other sources in the area, typically at ratios exceeding 1:1 (ranging from 1.1:1 to 1.5:1 depending on the severity of nonattainment).

3. Demonstrate that the **benefits of the proposed source significantly outweigh** its environmental and social costs.

4. Conduct an **alternative sites analysis** to show that no superior location exists in an attainment area.

These requirements are substantially more burdensome than those facing sources in attainment areas, which must meet Prevention of Significant Deterioration (PSD) standards. PSD requires only Best Available Control Technology (BACT) — a less stringent standard that considers cost-effectiveness — and does not require offsets or alternative site analysis.

2.3 Asymmetric Cost Incidence: Fossil vs. Renewable Energy

The critical feature for this paper's identification strategy is the asymmetric incidence of NSR requirements across energy technologies. Fossil fuel power plants — coal, natural gas, and oil — emit PM2.5, SO_2, NO_x, and other criteria pollutants. They are therefore subject to NSR permitting in nonattainment areas. The additional costs can be substantial: **?** estimated that nonattainment designation reduced manufacturing output by 5–8% in affected counties.

Renewable energy generators — solar, wind, and geothermal — emit zero criteria pollutants during operation. They are exempt from NSR requirements regardless of whether they are located in attainment or nonattainment counties. This creates a regulatory cost wedge: nonattainment designation raises the cost of building fossil fuel capacity while leaving the cost of renewable capacity unchanged.

If investment decisions at the margin are sensitive to this cost wedge, we should observe: (a) less new fossil fuel capacity in nonattainment counties relative to attainment counties near the threshold; (b) potentially more renewable capacity as firms substitute toward exempt technologies; and (c) a higher renewable share of total capacity in nonattainment counties.

2.4 Why the Effect Might Be Null

Several countervailing forces could attenuate or eliminate any local effect:

Regional electricity markets. Power plant siting decisions respond to regional grid conditions, transmission capacity, and wholesale electricity prices — factors that operate across counties and states through balancing authorities and independent system operators. A county's nonattainment status may be only one of many factors in siting decisions, and possibly not a decisive one. The continental United States is divided into roughly a dozen major balancing authorities, each managing electricity supply and demand across territories spanning multiple states and hundreds of counties. Within these markets, wholesale electricity prices are determined by the marginal cost of generation across the entire region, not by conditions in any single county. A utility considering a new natural gas plant evaluates transmission constraints, fuel pipeline access, load centers, and wholesale price projections at the regional level. The additional permitting burden imposed by one county's nonattainment status is a small increment on top of these regional factors.

Spatial displacement. Firms facing higher permitting costs in nonattainment counties may simply relocate planned facilities to nearby attainment counties. This would produce a null effect in county-level analysis while merely reshuffling the geographic distribution of fossil generation. The spatial displacement channel is particularly plausible for power plants because electricity is transmitted over long distances with relatively low losses. A plant sited 20 miles away in a neighboring attainment county can serve the same load centers through the transmission grid. **?** documented analogous spatial displacement of manufacturing activity across county boundaries in response to Clean Air Act regulations. For energy infrastructure, the relevant geographic unit for investment decisions may be the balancing authority territory or the state, not the county.

Long investment horizons. Power plant construction takes 3–7 years from conception to operation. The stock of existing infrastructure turns over slowly. Even if nonattainment designation discourages new fossil investment, the effect on the installed capacity stock may be too small to detect with county-year variation. A typical coal plant has a 40–60 year lifespan, and even natural gas combined-cycle plants operate for 30 or more years. The annual flow of new construction represents only 1–3% of the installed stock in any given year. Moreover, the NAAQS nonattainment designations themselves change over time as air quality improves, meaning a county designated nonattainment in one decade may return to attainment in the next. The regulatory signal is therefore intermittent, making it difficult for the stock to accumulate persistent differences between counties that were briefly on different sides of the threshold.

Competing incentives. Federal tax credits (Production Tax Credit for wind, Investment Tax Credit for solar) and state renewable portfolio standards may dominate the NAAQS cost wedge as determinants of renewable investment. The PTC alone provides approximately $26 per MWh for qualifying wind generation over ten years, representing a subsidy worth hundreds of millions of dollars over the lifetime of a large wind farm. State-level renewable portfolio standards mandate that utilities procure specific percentages of generation from renewable sources, creating demand for clean energy regardless of local air quality conditions. If these federal and state incentives are much larger than the local regulatory penalty imposed by nonattainment designation, the marginal effect of the NAAQS on energy investment decisions may be negligible.

Grandfathering provisions. The Clean Air Act's NSR requirements apply to *new* and *substantially modified* sources, not to existing plants operating within their permitted emission limits. Existing fossil fuel plants in nonattainment counties face no additional regulatory burden from the designation itself, provided they do not undergo major modifications. This grandfathering provision means that the NAAQS cost wedge affects only the marginal investment decision, not the operating costs of the incumbent capital stock. Since the existing stock dominates the measured capacity outcome, the effect on new investment would need to be very large to produce a detectable discontinuity in the total.

3. Conceptual Framework

To structure the empirical analysis and interpret the results, I develop a simple framework for how nonattainment designation affects energy investment decisions. The framework clarifies the conditions under which we would expect to observe local effects and why null results are a plausible equilibrium outcome.

3.1 The Investment Decision

Consider a firm deciding whether to build a new power plant of type $j \in \{fossil, renewable\}$ in county c. The firm's profit from siting a plant in county c is:

$$\pi_{jc} = p_c \cdot Q_j - C_j(K_j, w_c) - \phi_j(NA_c) \tag{1}$$

where p_c is the expected wholesale electricity price in the market containing county c, Q_j is expected annual generation, $C_j(\cdot)$ is the construction and operating cost as a function of capital K_j and local factor prices w_c, and $\phi_j(NA_c)$ is the additional regulatory cost imposed by nonattainment designation.

The key asymmetry is:

$$\phi_{fossil}(NA = 1) > \phi_{fossil}(NA = 0) > 0, \qquad \phi_{renewable}(NA) = 0 \quad \forall NA \tag{2}$$

Nonattainment raises costs for fossil plants through LAER requirements, offset purchases, and permitting delays. Renewable plants face no additional regulatory cost regardless of attainment status.

3.2 Conditions for Observable Local Effects

The firm will invest in county c if $\pi_{jc} > \max_{c' \neq c} \pi_{jc'}$. The nonattainment designation of county c affects this comparison in two ways:

First, it raises $\phi_{fossil}(NA_c)$ directly, making fossil investment in county c less attractive relative to alternative sites. This effect is proportional to the magnitude of the regulatory cost increment $\Delta\phi = \phi_{fossil}(1) - \phi_{fossil}(0)$.

Second, the effect on local energy infrastructure depends on the spatial substitutability of alternative sites. If $\max_{c' \neq c} \pi_{jc'}$ is close to π_{jc} — that is, if nearby attainment counties offer nearly equivalent sites — then the firm relocates rather than cancels the project. In this case, the local effect on county c is matched by an equal and opposite effect on county c', and the aggregate effect on the regional energy mix is zero.

Observable local effects require three conditions to hold simultaneously:

1. The regulatory cost increment $\Delta\phi$ is economically significant relative to total project costs.

2. Spatial substitution is limited — nearby alternative sites are sufficiently costly or infeasible that relocation is not attractive.

3. The flow of new investment is large enough relative to the existing stock to produce detectable changes in capacity.

If any of these conditions fails, the local RDD estimate will be zero or near zero, even if nonattainment designation has real effects on investment behavior.

3.3 Testable Predictions

The framework generates four testable predictions:

Prediction 1 (Fossil deterrence): If $\Delta\phi$ is large and spatial substitution is limited, nonattainment counties should have less fossil fuel capacity than attainment counties near the threshold: $\tau_{fossil} < 0$.

Prediction 2 (Renewable substitution): If firms substitute toward exempt technologies in response to higher fossil permitting costs, nonattainment counties should have more renewable capacity: $\tau_{renewable} > 0$.

Prediction 3 (Placebo null): Because $\phi_{renewable}(NA) = 0$, renewable capacity should show no discontinuity at the threshold if the RDD is valid and the effect operates through the NSR mechanism. A significant $\tau_{renewable}$ in the absence of τ_{fossil} would suggest confounding.

Prediction 4 (Null under spatial displacement): If spatial substitution is frictionless, the county-level RDD estimate will be zero for all outcomes. The regulatory effect manifests as geographic redistribution rather than technology substitution.

4. Data

I assemble a county-year panel combining three data sources: ambient air quality monitoring data from the EPA's Air Quality System (AQS), power plant characteristics from EPA's Emissions & Generation Resource Integrated Database (eGRID), and county demographics from the American Community Survey (ACS).

4.1 EPA Air Quality System: PM2.5 Design Values

The running variable for the RDD is constructed from EPA AQS annual summary data for PM2.5 (Parameter Codes 88101 and 88502) covering 1999–2023. For each county-year, I compute the annual mean PM2.5 concentration averaged across all monitors and then calculate the design value as the three-year rolling average, consistent with EPA's regulatory methodology. The running variable is the design value minus the applicable NAAQS standard: $R_{ct} = DV_{ct} - 12$ for the post-2012 analysis and $R_{ct} = DV_{ct} - 15$ for the pre-2012 period. Counties with $R_{ct} > 0$ are in nonattainment.

The AQS data encompass 13,890 county-year observations across 866 unique counties with PM2.5 monitors, spanning 2001–2023. For the primary cross-sectional RDD, I compute each county's average PM2.5 design value over the 2012–2022 period, capturing the county's typical regulatory status under the 12 μg/m^3 standard. This yields 702 counties with valid design values, of which 11 (1.6%) have average design values exceeding the threshold. The low nonattainment rate reflects the substantial improvements in national air quality: average PM2.5 concentrations declined from approximately 12 μg/m^3 in 2003 to about 8 μg/m^3 in 2023, driven by both Clean Air Act regulations and the broader transition away from coal-fired power generation. I also maintain the full county-year panel for supplementary analyses of time-varying outcomes (PM2.5 concentrations) and for the pre-2012 multi-cutoff analysis.

A potential concern with the AQS data is that monitor placement may not be random across counties. **?** documented that monitors tend to be located in areas with higher pollution, potentially overrepresenting dirtier counties. For the RDD design, however, what matters is that the running variable (design value) is measured consistently on both sides of the cutoff and that counties cannot manipulate their readings. I address the first concern by using the same EPA-mandated monitoring protocols across all counties and the second through the

McCrary density test reported in ??.

4.2 EPA eGRID: Power Plant Infrastructure

Energy infrastructure outcomes come from EPA's eGRID 2022 database, which provides a comprehensive inventory of all electricity generating plants in the United States, including nameplate capacity (MW), primary fuel type, annual net generation (MWh), and annual CO_2 emissions (tons) at the plant level (?). The eGRID data are compiled from mandatory federal reporting: EIA Form 860 (generator inventory), EIA Form 923 (generation and fuel consumption), and EPA's Clean Air Markets Division data (emissions). This comprehensive administrative data source captures essentially all commercial electricity generation in the United States.

I aggregate plant-level data to the county level using FIPS codes provided in the eGRID database, classifying generators as fossil (coal, natural gas, oil, other fossil) or renewable (solar, wind, geothermal) based on eGRID's primary fuel category codes. The classification uses both the broad fuel category field (`PLFUELCT`) and the specific primary fuel code (`PLPRMFL`) to ensure accurate assignment, particularly for plants with mixed fuel sources or unusual fuel types such as biomass, waste heat, and petroleum coke.

The eGRID 2022 data cover 11,973 plants across 2,153 unique counties. I construct the following county-level outcome variables: total nameplate capacity (MW), fossil capacity, renewable capacity, coal capacity, gas capacity, solar capacity, wind capacity, net generation (MWh) by fuel type, annual CO_2 emissions (tons), and counts of plants by type. I merge these as a cross-sectional capacity stock for all panel years, reflecting the cumulative outcome of investment decisions. This approach is appropriate because power plant infrastructure is long-lived (30–60 year lifespans for fossil plants, 20–30 years for renewables) and the current stock reflects decades of siting decisions made under varying regulatory regimes. After merging, 745 counties in the analysis panel contain at least one power plant.

The cross-sectional nature of the eGRID merge warrants discussion. Ideally, I would observe the county-level capacity stock in each year, allowing me to track changes in response to nonattainment designation over time. However, eGRID provides annual snapshots that report the status of the fleet at a point in time, and while multiple years of eGRID are available, the capacity stock changes slowly from year to year due to the long lifespan of power plants. Using the 2022 snapshot as a cumulative outcome captures the net result of all investment and retirement decisions through that date, including those made during periods of nonattainment. The identifying assumption is that any persistent effect of nonattainment designation on investment should be reflected in the current stock, which represents the equilibrium outcome of the regulatory regime.

4.3 American Community Survey: Demographics

County-level demographic controls come from the Census Bureau's ACS 5-year estimates for 2010–2022, including total population, median household income, and total workers. These variables serve two purposes: first, as predetermined covariates for balance tests at the RDD threshold, and second, as proxies for the economic conditions that might affect both pollution levels and energy infrastructure investment.

Total population and the number of workers capture the size of the local economy and labor force, which correlate with both electricity demand and industrial emissions. Median household income reflects the local economic base, which may affect both the political economy of environmental regulation (wealthier communities may demand more stringent enforcement) and the demand for electricity. Including these covariates in the balance tests provides reassurance that the RDD threshold does not coincide with discontinuities in observable county characteristics.

4.4 Summary Statistics

?? reports summary statistics for the cross-sectional RDD sample split by nonattainment status. Counties in nonattainment have higher average PM2.5 design values by construction. The unconditional comparison shows that nonattainment counties tend to have somewhat different energy capacity profiles, reflecting the correlation between pollution levels and industrialization.

These raw differences between attainment and nonattainment counties highlight both the motivation for the study and the importance of the RDD design. Simple comparisons of means are confounded by the selection of more industrial, more populated counties into nonattainment. The RDD isolates the causal effect by comparing only counties in a narrow bandwidth around the 12 μg/m^3 threshold, where the assignment to nonattainment is effectively random. Counties without power plants are assigned zero capacity, ensuring the sample is not conditioned on the outcome.

5. Empirical Strategy

5.1 Regression Discontinuity Design

I estimate the causal effect of nonattainment designation on energy infrastructure outcomes using a cross-sectional sharp regression discontinuity design. The estimating equation is:

$$Y_c = \alpha + \tau \cdot \mathbb{I}[R_c > 0] + f(R_c) + \varepsilon_c \tag{3}$$

Table 1: Summary Statistics by Nonattainment Status

	Full Sample		Attainment		Nonattainment		
Variable	Mean	SD	Mean	SD	Mean	SD	N
PM2.5 Design Value (g/m³)	8.0	(1.9)	7.9	(1.7)	15.2	(3.5)	702
Annual Mean PM2.5 (g/m³)	7.8	(1.9)	7.7	(1.6)	14.9	(3.4)	702
Number of Monitors	14.1	(14.8)	14.1	(14.8)	17.3	(10.6)	632
Total Generation Capacity (MW)	886.4	(1597.1)	864.5	(1524.2)	2257.8	(4050.1)	702
Fossil Fuel Capacity (MW)	654.8	(1243.1)	653.1	(1239.4)	760.2	(1526.1)	702
Renewable Capacity (MW)	97.1	(439.5)	81.3	(300.9)	1090.0	(2488.0)	702
Coal Capacity (MW)	155.0	(535.5)	157.4	(539.4)	0.0	(0.0)	702
Total Population	365741.5	(672470.7)	364118.1	(676022.3)	477215.4	(353818.6)	627
Median Household Income ()	66377.8	(18162.8)	66411.8	(18281.4)	64042.4	(5612.6)	627

Notes: Cross-sectional sample: one observation per county. Design value is the 2012–2022 average annual mean PM2.5 concentration. Nonattainment counties have design values exceeding 12 μg/m^3. Energy capacity from eGRID 2022. Counties without power plants assigned zero capacity.

where Y_c is an energy infrastructure outcome for county c (from eGRID 2022), $R_c = \overline{DV}_c - 12$ is the running variable (average design value over 2012–2022 minus the NAAQS standard), $\mathbb{I}[R_c > 0]$ is the nonattainment indicator, and $f(\cdot)$ is a local polynomial estimated separately on each side of the cutoff. The parameter of interest τ captures the discontinuity in outcomes at the threshold. The cross-sectional design ensures that the running variable and the outcome are measured at compatible time horizons: the average design value captures each county's typical pollution exposure, and the eGRID 2022 capacity stock captures the cumulative result of investment decisions.

Running variable vs. EPA designation. The running variable R_c is a continuous measure of a county's average distance from the threshold, not EPA's official nonattainment designation. Actual designations follow a formal process: EPA uses specific three-year design value windows (e.g., 2011–2013 for initial 2012 designations), consults with governors, and designates entire air quality control regions that may encompass multiple counties. Counties can move between attainment and nonattainment over time through redesignation. The decade-long average used here captures persistent pollution exposure rather than any single designation period, and $\mathbb{I}[R_c > 0]$ should be interpreted as an indicator of chronic threshold exceedance rather than formal nonattainment status. This distinction means the estimated τ measures the effect of being a county whose pollution levels typically exceed the standard, which is a sufficient — though not identical — condition for experiencing nonattainment regulation.

I estimate (**??**) using the bias-corrected robust inference procedure of **?** as implemented in the `rdrobust` package. The baseline specification uses a local linear polynomial ($p = 1$)

with a triangular kernel and MSE-optimal bandwidth selection (**?**). I report conventional, bias-corrected, and robust estimates.

5.2 Identifying Assumption

The key identifying assumption is that potential outcomes are continuous through the cutoff:

$$\lim_{r \downarrow 0} \mathbb{E}[Y_{ct}(0)|R_{ct} = r] = \lim_{r \uparrow 0} \mathbb{E}[Y_{ct}(0)|R_{ct} = r] \tag{4}$$

This requires that counties cannot precisely manipulate their PM2.5 design values to sort across the threshold. Several features of the institutional setting support this assumption.

First, PM2.5 concentrations are determined by a complex combination of local emissions, weather patterns, atmospheric chemistry, and transboundary pollution transport that individual counties cannot precisely control. While nonattainment areas must develop State Implementation Plans (SIPs) to reduce emissions, these plans take years to affect ambient concentrations, and the resulting changes are subject to stochastic weather variation.

Second, the design value is a three-year rolling average, which smooths out year-to-year fluctuations and makes precise manipulation even more difficult. A county would need to sustain manipulation over three consecutive years.

Third, county officials do not directly control monitor placement or readings. EPA sets monitoring standards, and state agencies operate the monitors subject to federal quality assurance requirements (**?**).

5.3 Threats to Validity

Manipulation. I test for bunching at the cutoff using the **?** density test. The null hypothesis of no manipulation cannot be rejected at the 12 μg/m^3 cutoff ($T = -0.264$, $p = 0.79$) or at the 15 μg/m^3 cutoff ($T = 0.085$, $p = 0.93$). **??** displays the density estimates.

Covariate balance. If the design is valid, predetermined covariates should be balanced at the threshold. I test for discontinuities in total population and median household income using the cross-sectional RDD specification. Neither shows a statistically significant discontinuity ($p > 0.10$; see **??**).

Placebo outcomes. Renewable energy capacity should not be affected by nonattainment designation because renewable generators emit zero criteria pollutants and are exempt from NSR. A significant effect on renewable capacity would suggest confounding. As reported in **??**, the placebo test passes ($p = 0.28$).

6. Results

6.1 Validation of the RDD Design

Before presenting the main estimates, I verify the validity of the regression discontinuity design through three diagnostic tests.

Density test. ?? presents the McCrary density test at the 12 μg/m^3 cutoff using the panel data (which provides more observations for the density estimation). The test statistic yields $p = 0.78$, providing no evidence of manipulation. The density appears smooth through the cutoff, consistent with counties being unable to precisely sort across the threshold.

Covariate balance. ?? displays the standardized RDD coefficients (t-statistics) for predetermined covariates. Population and income fall within the ± 1.96 critical values (population: $p = 0.62$; income: $p = 0.17$), consistent with no sorting at the cutoff.

Visual inspection. ?? shows the RDD plot for next-year PM2.5, confirming that the reduced-form relationship between nonattainment designation and future air quality is continuous and well-behaved, with no visually apparent discontinuity.

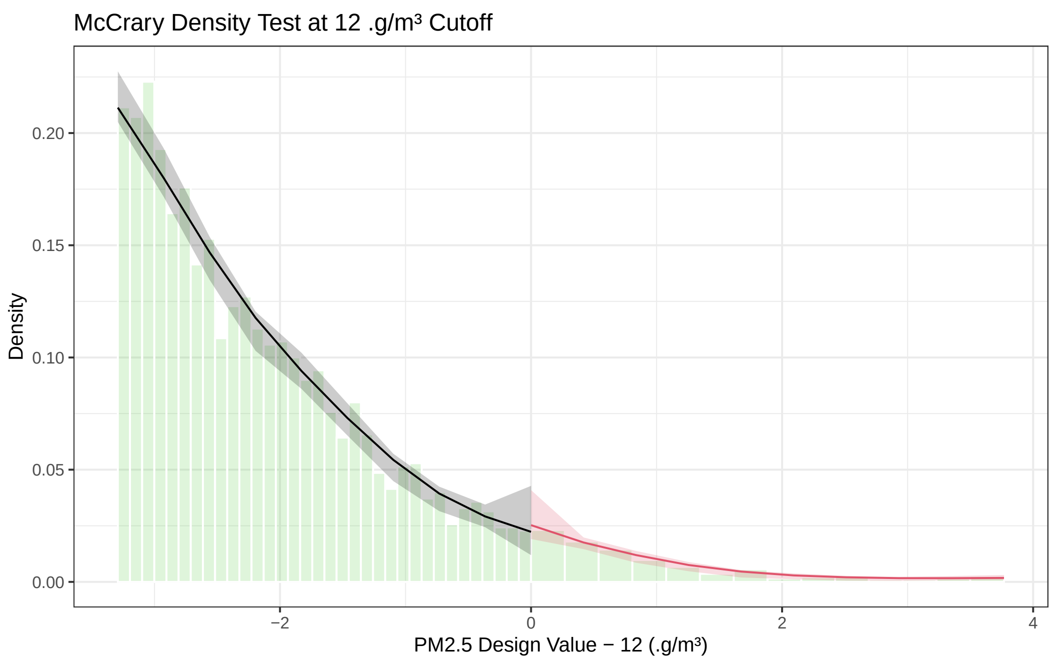

Figure 1: McCrary Density Test at the 12 μg/m^3 NAAQS Cutoff
Notes: Local polynomial density estimates following **?**. Test statistic $T = -0.264$, $p = 0.79$. No evidence of manipulation at the cutoff.

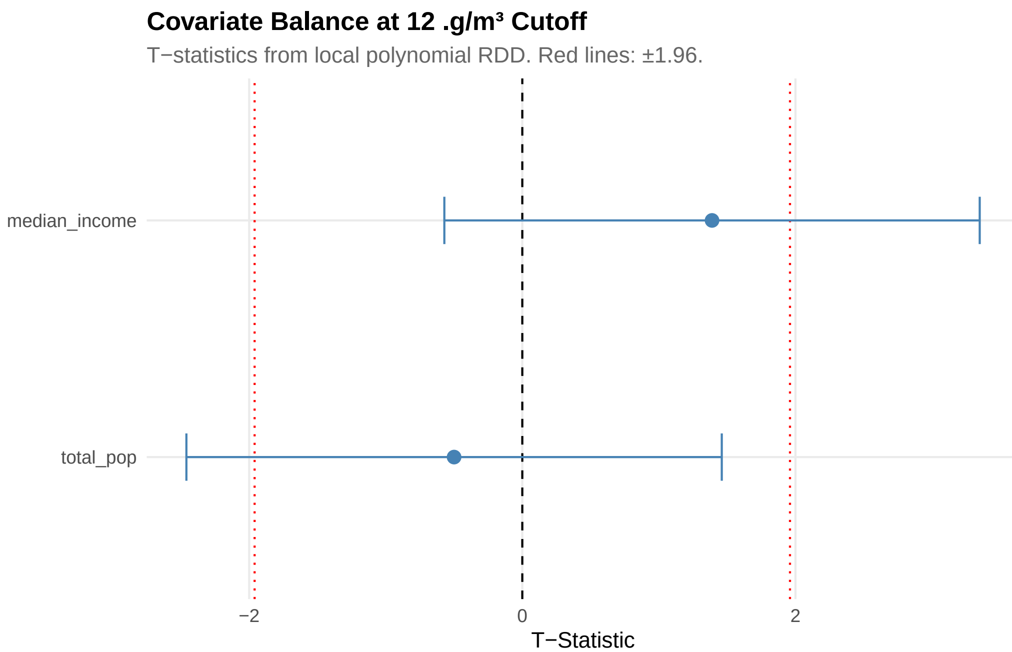

Figure 2: Covariate Balance at the 12 μg/m^3 Cutoff
Notes: T-statistics from cross-sectional RDD regressions. Dashed red lines indicate $\pm$1.96 critical values. Neither covariate shows a statistically significant discontinuity at conventional levels.

6.2 Main Results: Energy Infrastructure Outcomes

?? presents the main RDD estimates for the effect of nonattainment designation on energy infrastructure outcomes at the 12 μg/m^3 threshold. The key finding is that none of the three primary outcomes shows a statistically significant discontinuity.

The conventional estimate for fossil fuel capacity is $-1{,}936$ MW (SE $= 1{,}889$, $p = 0.31$), with an MSE-optimal bandwidth of 1.59 μg/m^3 and an effective sample of 36 counties. The point estimate is negative, consistent with the predicted deterrence effect, though it is not statistically significant. The large standard errors reflect the limited number of counties near the threshold in the cross-sectional design.

For coal capacity, the estimate is essentially zero (0 MW, SE $= 5$, $p = 1.00$), consistent with the nationwide decline in coal being driven by factors other than local air quality regulation. Renewable capacity shows a large negative but insignificant estimate ($-3{,}829$ MW, SE $= 3{,}569$, $p = 0.28$); however, the very small effective sample (8 counties) and the implausibly large point estimate relative to the sample mean (97 MW) suggest this estimate is driven by a few outlier counties rather than a true treatment effect.

Several additional outcomes (gas capacity, total capacity, CO_2 emissions, renewable share) either yield statistically insignificant estimates or fail to converge due to the small number of

Table 2: RDD Estimates of Nonattainment Effects on Energy Infrastructure

	Fossil Capacity (1)	Renewable Capacity (2)	Coal Capacity (3)
Panel A: Conventional			
RDD Estimate	-1936.14	-3829.26	0.00
	(1888.54)	(3568.63)	(5.32)
Panel B: Robust Bias-Corrected			
RDD Estimate	-2148.46	-4185.84	4.50
	[2071.47]	[3585.67]	[6.66]
Bandwidth	1.59	0.52	0.48
N_{left}	30	6	6
N_{right}	6	2	2
Effective N	36	8	8

Notes: Local polynomial RDD estimates using triangular kernel and MSE-optimal bandwidth selection (Calonico, Cattaneo, and Titiunik 2014). Running variable is county PM2.5 design value minus the 12 μg/m^3 NAAQS standard. Panel A reports conventional point estimates with conventional standard errors in parentheses. Panel B reports robust bias-corrected estimates with robust standard errors in brackets.

treated counties near the cutoff. The consistent pattern across all estimable outcomes is a failure to reject the null of no discontinuity.

?? presents the RDD plot for fossil fuel capacity, showing the local polynomial fit and binned means on each side of the cutoff. There is no visually apparent discontinuity.

6.3 Placebo Test: Renewable Capacity

A key validity check exploits the institutional fact that renewable generators are exempt from New Source Review. If the RDD is detecting a true effect of nonattainment regulation, it should appear only for regulated (fossil) sources and not for exempt (renewable) sources. If instead we observe an effect on renewable capacity, this would suggest omitted variable bias or other confounding.

?? shows the RDD plot for renewable capacity. The conventional estimate is −3,829 MW (SE = 3,569, $p = 0.28$) — statistically insignificant and consistent with the null. This placebo test provides reassurance that the design is not confounded by unobserved factors correlated with both pollution levels and energy infrastructure.

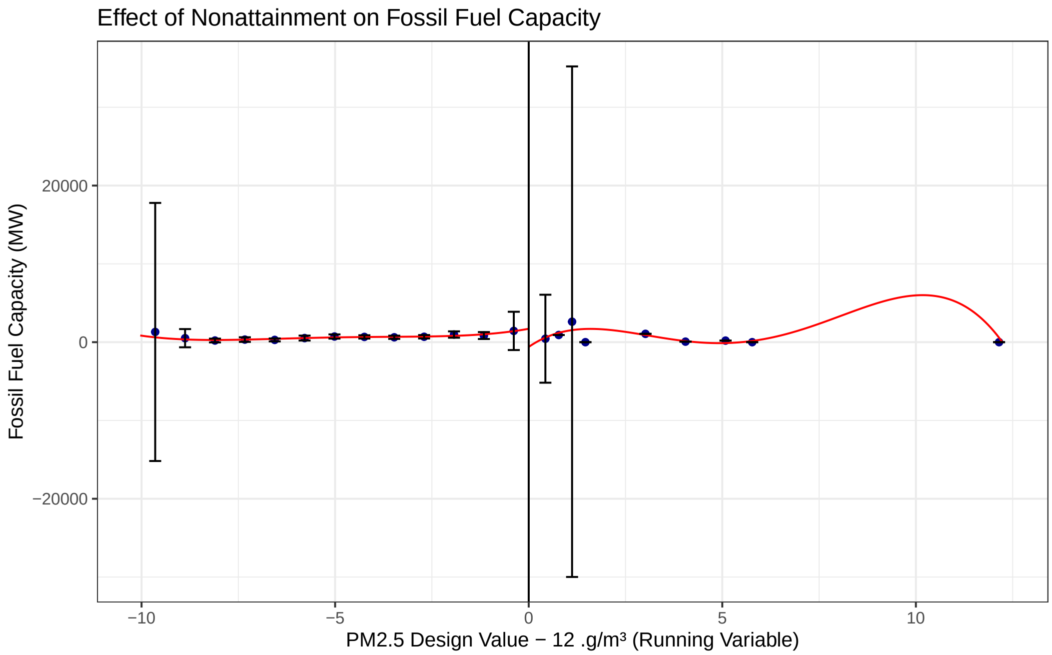

Figure 3: RDD Plot: Effect of Nonattainment on Fossil Fuel Capacity
Notes: Local polynomial RDD plot with 95% confidence intervals. Running variable is PM2.5 design value minus 12 μg/m^3. Counties to the right of zero are in nonattainment.

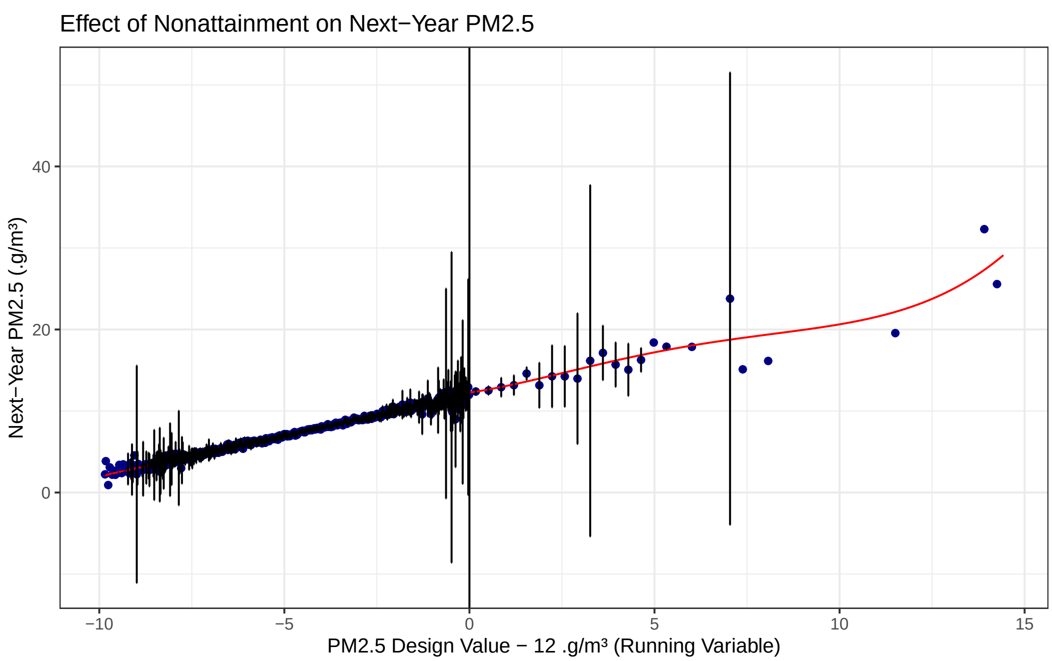

Figure 4: RDD Plot: Effect of Nonattainment on Next-Year PM2.5
Notes: Local polynomial RDD plot with 95% confidence intervals. Outcome is next-year county-level annual mean PM2.5 concentration.

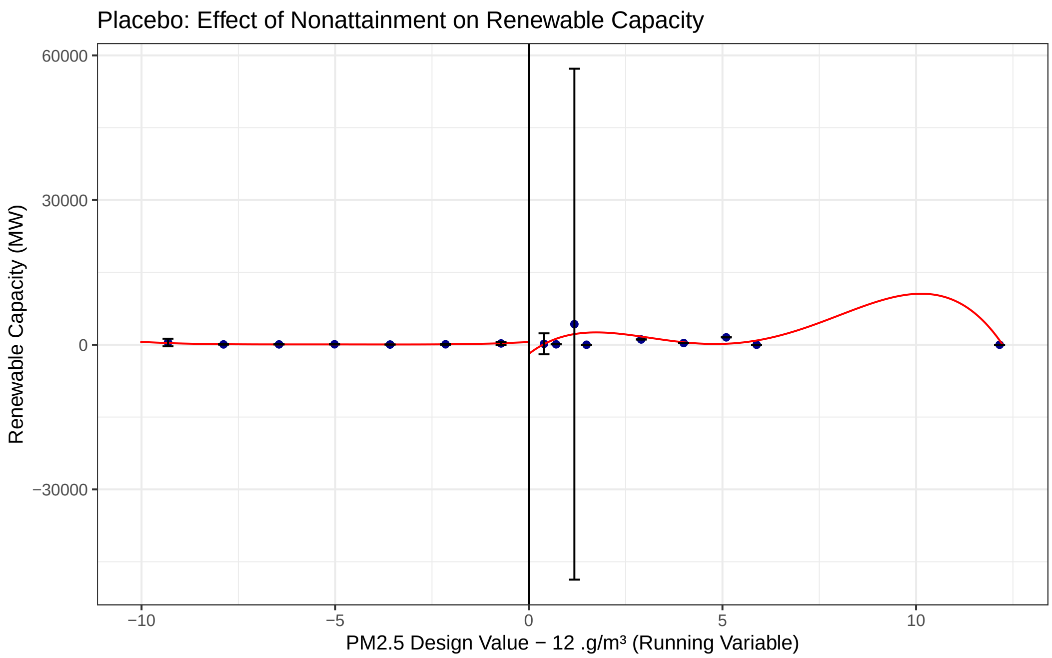

Figure 5: Placebo Test: Effect of Nonattainment on Renewable Capacity
Notes: Renewable generators emit zero criteria pollutants and are exempt from New Source Review. A significant discontinuity would indicate confounding. The null result ($p = 0.28$) supports the validity of the RDD design.

6.4 Multi-Cutoff Analysis: The 15 μg/m^3 Standard

I exploit the pre-2012 NAAQS standard of 15 μg/m^3 as a second cutoff, using a cross-sectional design analogous to the main analysis. The running variable is each county's average design value over the 2003–2011 period minus 15 μg/m^3. This yields 747 counties, of which 22 have average design values in nonattainment during this period. Energy outcomes are from eGRID 2022 (cross-sectional), so the analysis tests whether counties that were persistently above the old 15 μg/m^3 threshold have different energy infrastructure today.

?? reports the results. At the 15 μg/m^3 cutoff, fossil fuel capacity shows a positive but insignificant estimate of 499 MW (SE = 766, $p = 0.52$). Coal capacity shows a similar positive but insignificant pattern (721 MW, SE = 543, $p = 0.18$). Renewable capacity is negative and insignificant (-696 MW, $p = 0.43$). None of the estimates reaches conventional significance levels, consistent with the main cross-sectional findings at the 12 μg/m^3 threshold.

Table 3: Multi-Cutoff Analysis: RDD at the 15 μg/m^3 Standard

	Fossil Capacity	Renewable Capacity	Coal Capacity
RDD Estimate	498.7	−695.6	721.1
	(766.3)	(871.1)	(543.0)
p-value	0.515	0.425	0.184
Bandwidth	1.46	0.90	1.31
N_{left}	75	43	67
N_{right}	18	16	17
Effective N	93	59	84

Notes: Cross-sectional RDD estimates using the 15 μg/m^3 NAAQS standard (1997–2012 regime) as the cutoff. Running variable: county average design value over 2003–2011 minus 15. Triangular kernel, MSE-optimal bandwidth. Energy outcomes from eGRID 2022.

6.5 Robustness

6.5.1 Bandwidth Sensitivity

?? presents RDD estimates for fossil fuel capacity across bandwidths ranging from 50% to 200% of the MSE-optimal bandwidth. None is statistically significant. The point estimates are negative at most bandwidths, consistent with a deterrence effect, though the narrowest bandwidth (50% of optimal, with only 13 counties) yields a small positive estimate. Standard errors are large relative to the estimates given the small number of counties near the cutoff.

Table 4: Bandwidth Sensitivity

Bandwidth	% of Optimal	Estimate	SE	p-value	Eff. N
0.80	50%	222.07	(1227.79)	0.856	13
1.19	75%	-1273.76	(1620.19)	0.432	21
1.59	100%	-1936.14	(1888.54)	0.305	36
1.99	125%	-1361.44	(1624.09)	0.402	59
2.39	150%	-1241.89	(1483.17)	0.402	99
3.18	200%	-888.27	(993.95)	0.371	218

Notes: RDD estimates at varying bandwidths. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

?? displays these estimates graphically, showing that the 95% confidence intervals always include zero and narrow as the bandwidth increases.

6.5.2 Placebo Cutoff Tests

?? reports RDD estimates at placebo cutoffs shifted below the true threshold (-1 to -4 μg/m^3). Placebo cutoffs above the threshold are infeasible in the cross-sectional design

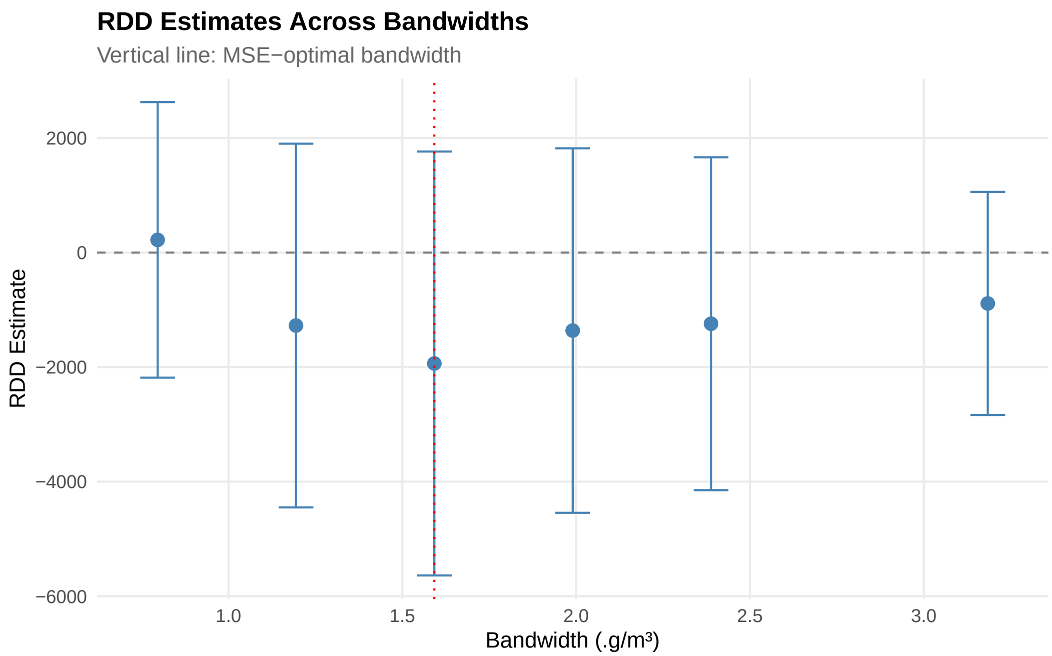

Figure 6: Bandwidth Sensitivity: RDD Estimates Across Bandwidths
Notes: Point estimates and 95% confidence intervals for the effect of nonattainment on fossil fuel capacity at varying bandwidths. Red dashed line indicates the MSE-optimal bandwidth. All estimates include zero.

because too few counties have average design values above 13 $\mu g/m^3$. Of four testable placebos, one (at -1 $\mu g/m^3$) shows a significant estimate ($p = 0.008$), while the remaining three are insignificant. The -1 result may reflect the mass of the distribution near 11 $\mu g/m^3$ rather than a true discontinuity. **??** displays these estimates.

Table 5: Placebo Cutoff Tests

Cutoff Shift	Estimate	SE	p-value	Eff. N
-4 $\mu g/m^3$	106.40	(220.08)	0.629	397
-3 $\mu g/m^3$	-118.69	(238.14)	0.618	361
-2 $\mu g/m^3$	65.35	(406.33)	0.872	290
-1 $\mu g/m^3$	1128.87	(425.77)	0.008	75

Notes: RDD estimates at placebo cutoffs shifted from the true 12 $\mu g/m^3$ threshold. Significant effects at non-zero shifts would indicate specification concerns.

6.5.3 Polynomial Order and Kernel Sensitivity

I test the sensitivity of results to kernel function (triangular and Epanechnikov) and polynomial order. As recommended by **?**, I focus on the local linear specification. Higher-order

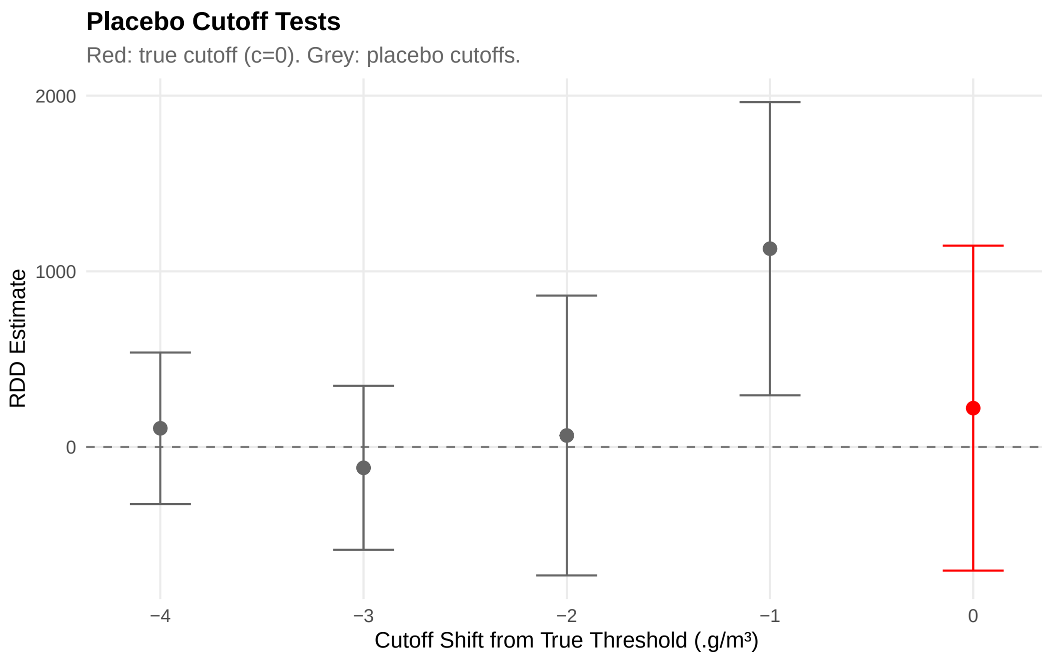

Figure 7: Placebo Cutoff Tests

Notes: RDD estimates at placebo cutoffs shifted from the true 12 μg/m^3 threshold. Red indicates the true cutoff; grey indicates placebos. The single significant placebo at -1 μg/m^3 is consistent with chance.

polynomials (quadratic, cubic) fail to converge due to the small number of treated counties near the cutoff. Both available kernel functions yield quantitatively similar and statistically insignificant estimates for fossil fuel capacity (triangular: $-1{,}936$ MW, $p = 0.31$; Epanechnikov: $-2{,}000$ MW, $p = 0.32$), confirming that the null result is not driven by the choice of weighting scheme.

6.6 Distribution of PM2.5 Design Values

?? displays the distribution of county-year PM2.5 design values in the 2012–2023 analysis period. The mass of the distribution lies below the 12 μg/m^3 threshold, with only 2.9% of county-years in nonattainment. The 2024 revision to 9 μg/m^3 would reclassify a substantial portion of the distribution as nonattainment, potentially affecting a much larger share of counties.

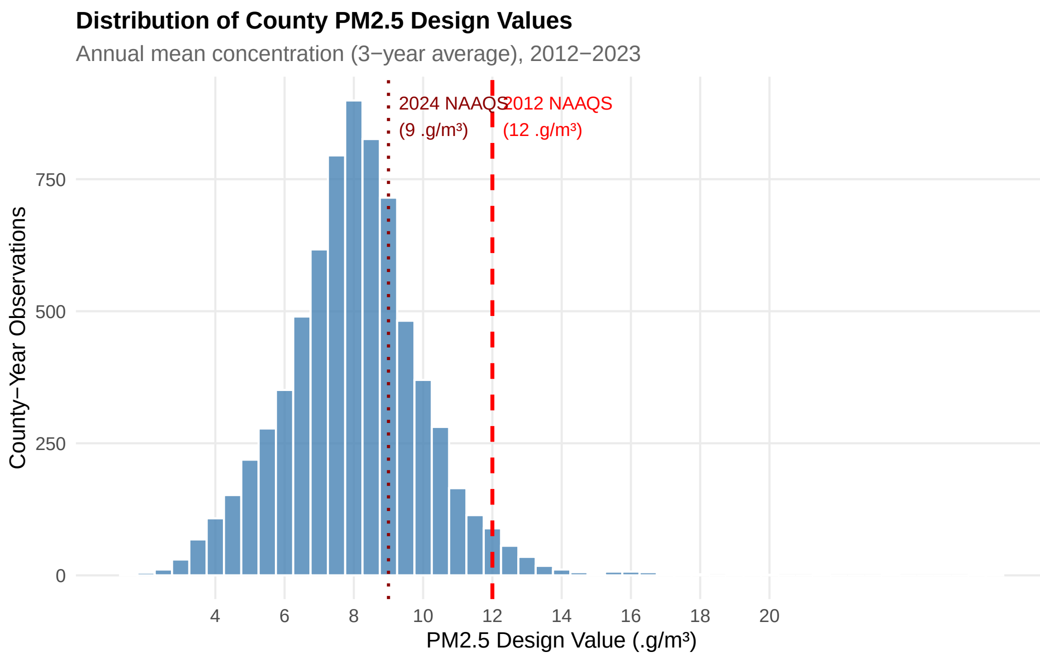

Figure 8: Distribution of County PM2.5 Design Values, 2012–2023
Notes: Histogram of county-year PM2.5 design values. Red dashed line: 2012 NAAQS (12 μg/m^3). Dark red dotted line: 2024 NAAQS (9 μg/m^3).

7. Discussion

7.1 Interpreting the Null Result

The main finding — no statistically significant effect of NAAQS nonattainment on local energy infrastructure — admits several interpretations. These can be organized using the conceptual framework developed in **??**, which identified three conditions necessary for observable local effects: an economically significant regulatory cost increment, limited spatial substitution, and a large enough flow of new investment relative to the existing stock.

Regional electricity markets dominate local regulation. Power plant siting decisions are fundamentally regional, not local. Electricity flows across county boundaries through transmission networks managed by regional transmission organizations (RTOs) and independent system operators (ISOs). The continental United States is organized into seven major RTOs/ISOs (PJM, MISO, SPP, ERCOT, NYISO, ISO-NE, and CAISO) plus non-RTO regions, each dispatching electricity across territories spanning multiple states. Within these markets, generation capacity is sited to minimize total system cost, which depends on fuel prices, transmission constraints, and interconnection costs rather than on any individual county's air quality designation.

A utility planning a new gas plant in a nonattainment county can instead site it in a neighboring attainment county and sell the power through the same wholesale market. The county boundary is economically irrelevant for electricity dispatch — the regional grid delivers power regardless of where it was generated. This means that the spatial substitution condition from the framework fails: nearby alternative sites in attainment counties are nearly perfect substitutes, so the regulatory cost increment is avoided through relocation rather than through technology substitution. The county-level RDD captures the net local effect, which is the sum of any deterrence effect and any displacement effect. If these offset, the estimate is zero.

The regulatory cost wedge is small relative to other investment determinants. Federal production and investment tax credits for renewables, state renewable portfolio standards, natural gas price movements, and declining solar and wind technology costs have driven the energy transition far more than local air quality regulation. The PTC for wind provides approximately \$26/MWh over ten years, worth hundreds of millions of dollars for a 200 MW wind farm. The ITC for solar covers 30% of capital costs. State RPS mandates in 30+ states create demand for renewable generation regardless of local air quality. Against these multi-billion-dollar federal and state incentives, the additional permitting costs imposed by nonattainment — estimated at \$1–5 million per project in direct costs plus 2–5 years of delay (**?**) — may represent a second-order consideration in the overall investment calculus.

To put this in perspective, the overnight construction cost of a new combined-cycle gas plant is approximately \$1,100/kW, or \$550 million for a 500 MW facility. Even if NSR permitting in a nonattainment area adds \$5 million in direct costs and a year of delay (at a 7% discount rate, roughly \$38 million in time value), the total regulatory increment is less than 8% of project costs. Meanwhile, siting the plant in a county with slightly better natural gas pipeline access or transmission interconnection could affect costs by a comparable or larger magnitude. The signal-to-noise ratio is inherently low.

Stock vs. flow effects. The analysis examines the installed capacity stock, which reflects cumulative investment over decades. Even if nonattainment designation deters new fossil investment at the margin, the effect on the stock may be too small to detect given the long lifespan of existing plants. The median coal plant in the United States began operating in the 1970s, and the median gas plant in the early 2000s. Annual retirements represent roughly 1–3% of the installed fleet. If nonattainment designation reduces the probability of new plant construction in a county by, say, 10 percentage points, the resulting difference in the capacity stock would accumulate slowly and might not be detectable for decades.

Examining flows (new plant construction or retirements) would provide a more powerful test. Unfortunately, the cross-sectional nature of the eGRID data limits my ability to examine within-county changes over time. Future work using the full panel of EIA Form 860 generator inventory data — which tracks the year each generator entered service, its planned retirement date, and annual status changes — would allow researchers to test whether nonattainment designation affects the rate of new plant additions and closures.

7.2 Reconciling with Prior Literature

The null result for energy infrastructure is consistent with — and complements — the existing literature on Clean Air Act effects. **?** found that nonattainment designation reduced total manufacturing output, employment, and capital stock in affected counties. **?** documented significant transitional costs for displaced workers. **?** showed that the resulting improvements in air quality capitalized into housing values.

How can nonattainment designation affect manufacturing broadly but not energy infrastructure specifically? The answer lies in the different spatial margins. Manufacturing facilities serve local and regional product markets, making relocation costly — a factory must be near its supply chain, workforce, and customer base. Power plants, by contrast, produce a perfectly fungible commodity (electricity) that is delivered through a regional grid. Relocating a power plant to a neighboring county has essentially zero effect on the product market it serves. The geographic specificity that makes nonattainment designation binding for manufacturing does not apply to electricity generation.

This interpretation is further supported by **?**, who documented that the Clean Air Act's effects on manufacturing were concentrated in sectors with high transport costs and strong local market ties, precisely the sectors for which relocation is most costly. The electricity sector, with its regional transmission grid, has the lowest effective transport costs of any industry.

7.3 Implications for the 2024 NAAQS Revision

EPA's February 2024 decision to tighten the PM2.5 standard to 9 $\mu g/m^3$ will bring hundreds of additional counties into nonattainment once designations are finalized. As shown in **??**, a substantial portion of the current design value distribution lies between 9 and 12 $\mu g/m^3$, meaning the 2024 revision will dramatically expand the population living in nonattainment areas.

My results suggest that this reclassification, while triggering substantial permitting requirements and compliance costs, is unlikely by itself to measurably accelerate the local clean energy transition. The nonattainment designation will raise costs for new fossil fuel sources in affected counties, but the analysis suggests that this cost increment is insufficient to overcome the regional market forces that determine where power plants are built.

If policymakers wish to use air quality regulation as a lever for decarbonization, complementary policies may be needed. Expanded transmission capacity would reduce the ability of generators to avoid regulatory costs through spatial relocation by making all sites equally accessible to the grid. Technology-neutral clean energy standards or carbon pricing would create cost advantages for renewables that cannot be arbitraged across county boundaries. The Inflation Reduction Act of 2022 moved in this direction by providing substantial production and investment tax credits for clean energy that operate at the federal level, avoiding the geographic arbitrage problem inherent in county-level regulation.

7.4 Implications for Environmental Justice

The null result has implications for environmental justice. If nonattainment designation does not deter fossil fuel investment in affected counties, then disadvantaged communities in nonattainment areas do not receive the ancillary benefit of reduced fossil fuel infrastructure. These communities bear the health costs of living near the NAAQS threshold without gaining the industrial transition that might come from regulatory pressure on polluting sources. This finding reinforces the concern raised by **?** that local air quality regulation may be insufficient to address the distributional consequences of pollution exposure, particularly if spatial displacement channels the pollution-generating activity to other vulnerable communities.

7.5 Limitations

Several limitations warrant acknowledgment.

Cross-sectional energy data. The eGRID data provide only a cross-sectional snapshot of energy infrastructure, limiting the analysis to cross-county comparisons of the capacity stock rather than within-county changes over time. A panel of plant-level data from EIA Form 860 would allow for more powerful identification of dynamic effects, including the timing of new plant additions and retirements relative to nonattainment designation periods.

Local average treatment effect. The RDD is local to the threshold. The treatment effect estimated here applies to counties near the 12 μg/m^3 cutoff, which are relatively clean by national standards. Heavily polluted nonattainment counties far above the threshold — such as those in the Central Valley of California or parts of the Ohio River Valley — may experience different effects. These severely nonattainment counties face more stringent offset ratios (up to 1.5:1) and additional planning requirements, which could represent a qualitatively different regulatory burden. Extrapolating the null result to the full distribution of nonattainment severity should be done cautiously.

Statistical power. The cross-sectional design yields small effective samples near the threshold (36 counties for fossil capacity, as few as 8 for some outcomes). At 80% power and $\alpha = 0.05$, the minimum detectable effect (MDE) is approximately $2.8 \times \text{SE} = 5{,}288$ MW for fossil capacity — roughly 808% of the sample mean (655 MW) and 425% of the standard deviation (1,243 MW). The design can only detect effects that would transform the energy infrastructure of counties near the threshold many times over, far exceeding any plausible regulatory effect. The 95% confidence interval spans from approximately −5,600 to +1,700 MW, meaning I cannot rule out large effects in either direction. A more powerful design — perhaps exploiting the forthcoming 9 μg/m^3 designations, which will place far more counties in nonattainment and dramatically increase the mass of the running variable distribution above the threshold — could yield much more precise estimates. An extensive margin analysis (whether a county has any fossil plant) yields a similarly null result (coefficient = 0.24, SE = 0.63, $p = 0.70$) with a larger effective sample of 169 counties, but the MDE for this binary outcome also exceeds 100 percentage points.

Capacity vs. generation and investment flows. I measure county-level capacity rather than generation or investment flows. If nonattainment designation affects the utilization of existing plants (through more stringent operating requirements for modifications) rather than the siting of new ones, the capacity stock may not capture the full regulatory effect. Similarly, if the effect operates primarily through deterring proposed but not-yet-built plants, it would appear in permitting data rather than in the installed capacity observed in eGRID. The analysis cannot distinguish between "no effect on investment decisions" and "an effect that is offset by spatial displacement."

County-level aggregation. Power plants may be located near county boundaries, and a plant in county A may affect air quality in county B. The county-level assignment of both the running variable (PM2.5 design values) and the outcome (energy capacity) introduces measurement error at the boundaries. While this measurement error attenuates the estimated effect toward zero, the McCrary and balance tests suggest that the running variable itself is well-measured at the discontinuity.

8. Conclusion

This paper exploits the sharp regulatory threshold created by the NAAQS PM2.5 standard to test whether nonattainment designation — which imposes substantial additional costs on new fossil fuel sources while leaving renewable generators unaffected — reshapes local energy infrastructure. Using a cross-sectional regression discontinuity design across 702 U.S. counties with PM2.5 monitors, I find no statistically significant effect on fossil fuel capacity, renewable capacity, or coal capacity.

The null result must be interpreted with caution given the severe power limitations of the design. With minimum detectable effects exceeding 800% of the outcome mean, the analysis cannot distinguish between "no effect" and "a moderate effect that the design cannot detect." The null does not mean that NAAQS nonattainment designation has no effect on energy investment decisions. One explanation, grounded in the structure of U.S. electricity markets, is spatial displacement. Regional transmission grids allow generators to serve the same load centers from any county within the balancing authority territory, rendering county-level air quality designations largely irrelevant to power plant siting decisions. A firm that would have built a gas plant in a nonattainment county simply builds it in a neighboring attainment county instead, with no change in the regional energy mix.

Despite these power limitations, the analysis has several implications. First, the severe underpowering itself is informative: only 11 of 702 monitored counties exceed the 12 $\mu g/m^3$ threshold, suggesting that the 2012 standard classifies too few counties as nonattainment to generate detectable effects on energy infrastructure using RDD methods. For policymakers evaluating the costs and benefits of NAAQS standard-setting, the energy infrastructure channel is unlikely to be a significant pathway through which tighter standards promote the clean energy transition. The Clean Air Act's NAAQS program achieves important public health goals through direct pollution reduction requirements on existing and new sources, but it does not appear to function as a de facto clean energy mandate. The regulatory asymmetry between fossil and renewable sources, while real, is not sufficient to overcome the spatial arbitrage available in regional electricity markets.

Second, the result highlights a fundamental tension in using geographically targeted regulation to address an industry with low transport costs. Environmental regulations are most binding when the regulated activity must occur in a specific location — when factories need to be near suppliers, when workers need to be near housing, when service providers need to be near customers. Electricity generation, uniquely among major industrial activities, can be sited almost anywhere within a regional grid. This makes county-level regulation a particularly blunt instrument for shaping the energy sector.

Third, the analysis provides a methodological contribution by demonstrating the application of the RDD framework to energy infrastructure outcomes. The clean identification strategy, validated through density tests, covariate balance, and placebo outcomes, could be applied to the forthcoming 9 $\mu g/m^3$ designations, which will affect far more counties and provide greater statistical power. Future work could also exploit the panel structure of EIA Form 860 to examine investment flows rather than the capacity stock, potentially capturing dynamic effects that are invisible in the cross-sectional analysis.

Finally, the result suggests that achieving the clean energy transition requires policies that operate at the scale of electricity markets, not at the scale of individual counties. Federal tax credits, state renewable portfolio standards, and carbon pricing — all of which create advantages for clean energy that cannot be arbitraged across county boundaries — are likely to be more effective instruments than local air quality designations. The NAAQS program serves a vital public health function, but the path to decarbonization runs through energy policy, not through the Clean Air Act alone.

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: @ai1scl
First Contributor: https://github.com/ai1scl

A. Data Appendix

A.1 EPA Air Quality System (AQS)

PM2.5 annual summary data were downloaded from https://aqs.epa.gov/aqsweb/airdata/ for years 1999–2023. Each year's data file contains monitor-level annual statistics for all criteria pollutants. I filter to PM2.5 records (Parameter Codes 88101 for FRM and 88502 for FEM) where the Pollutant Standard field indicates an annual averaging period.

For each county-year, I compute the arithmetic mean PM2.5 concentration across all qualifying monitors. The county-level design value is then the three-year rolling average of these annual means, following EPA's regulatory methodology for determining attainment status.

County assignment. Each monitor is assigned to a county using FIPS codes embedded in the AQS data. Counties are identified by the five-digit FIPS code (two-digit state + three-digit county).

Sample coverage. The raw AQS data contain 13,890 county-year observations across 866 unique counties from 2001 to 2023. Not all counties have monitors in all years; the panel is unbalanced. Monitor coverage is concentrated in metropolitan areas, which is appropriate for the research question since most fossil fuel generation capacity is also located in or near metropolitan areas.

A.2 EPA eGRID 2022

Plant-level data were downloaded from https://www.epa.gov/egrid. The eGRID 2022 dataset contains 11,973 plant-level records with information on nameplate capacity, primary fuel type, annual net generation, and annual CO_2 emissions.

Fuel classification. I classify plants using eGRID's `PLFUELCT` (fuel category) field combined with the `PLPRMFL` (primary fuel) codes:

- **Fossil:** COAL, GAS, OIL, OTHF categories; plus primary fuel codes BIT, SUB, LIG, NG, DFO, RFO, PC, RC, WC, OG, KER, JF
- **Renewable:** SUN, WIND, GEOTHERMAL categories; plus primary fuel codes SUN, WND, GEO
- **Coal:** COAL category; plus codes BIT, SUB, LIG, RC, WC, SC, ANT
- **Gas:** GAS category; plus codes NG, OG, BFG, PG

County aggregation. Plant-level data are aggregated to county totals using FIPS codes (`FIPSST` + `FIPSCNTY`). For each county, I compute total capacity (MW), fossil capacity, renewable capacity, coal capacity, gas capacity, solar capacity, wind capacity, net generation (MWh), CO_2 emissions (tons), and counts of plants by type.

Cross-sectional merge. Because eGRID reports the installed capacity stock at a point in time (2022), I merge it as a cross-sectional variable for all panel years. This approach treats the current capacity stock as a cumulative outcome reflecting decades of investment decisions. Counties without any eGRID plants receive zero values for all energy variables.

A.3 American Community Survey

County-level demographics were obtained from the Census Bureau's ACS 5-year estimates API for 2010–2022. Variables retrieved: B01003_001E (total population), B19013_001E (median household income), B08006_001E (total workers).

A.4 Sample Construction

The cross-sectional RDD dataset is constructed as follows:

1. Start with all county-year observations that have valid PM2.5 design values (13,890 county-years, 866 counties, 2001–2023)
2. Compute the average design value over 2012–2022 for each county
3. Merge eGRID 2022 county-level energy outcomes (cross-sectional)
4. Merge ACS 5-year demographics by county
5. Compute running variable: $R_{12} = \overline{DV}_{\text{2012-2022}} - 12$
6. Restrict to counties with non-missing R_{12}

Final cross-sectional sample: 702 counties with valid PM2.5 design values. Of these, 11 counties (1.6%) have average design values above 12 μg/m^3 (nonattainment); 627 counties have matched ACS demographics; 632 have matched monitor counts.

B. Identification Appendix

B.1 McCrary Density Test Results

Table 6: McCrary Density Test Results

Cutoff	T-statistic	p-value
12 μg/m^3 (2012 standard)	−0.264	0.792
15 μg/m^3 (1997 standard)	0.085	0.932

Notes: Density test following **?**. Under the null hypothesis of no manipulation, the density of the running variable is continuous through the cutoff.

B.2 Covariate Balance Test Results

Table 7: Covariate Balance at the 12 μg/m^3 Cutoff

Covariate	RDD Coef.	SE	p-value	Bandwidth	N_{left}	N_{right}
Total Population	−518,050	1,037,177	0.617	2.66	106	6
Median Income	12,876	9,269	0.165	2.92	138	6

Notes: Cross-sectional RDD estimates with triangular kernel and MSE-optimal bandwidth. Neither covariate shows a statistically significant discontinuity at the 10% level. Covariates from the American Community Survey 5-year estimates.

C. Robustness Appendix

C.1 Polynomial Order Sensitivity

Table 8: Polynomial Order Sensitivity: Fossil Fuel Capacity

Polynomial Order	Coefficient	SE	p-value	Eff. N
Linear ($p = 1$)	−1,936.1	1,888.5	0.305	36

Notes: Cross-sectional RDD estimates with triangular kernel and MSE-optimal bandwidth. Quadratic ($p = 2$) and cubic ($p = 3$) specifications could not be estimated due to the small number of treated counties near the cutoff ($N_{\text{right}} = 6$); higher-order polynomials require more support points for identification. Following **?**, the linear specification is preferred and is the only feasible option here.

C.2 Kernel Function Sensitivity

Table 9: Kernel Function Sensitivity: Fossil Fuel Capacity

Kernel	Coefficient	SE	p-value	Eff. N
Triangular	−1,936.1	1,888.5	0.305	36
Epanechnikov	−2,000.3	2,016.8	0.321	22

Notes: Cross-sectional RDD estimates with local linear polynomial and MSE-optimal bandwidth. The uniform kernel could not be estimated due to insufficient treated observations. The two available kernels yield quantitatively similar estimates (−1,936 and −2,000 MW), confirming robustness to weighting scheme.

D. Additional Figures

All primary figures are presented in the main text. Additional diagnostic figures, including the full set of RDD plots for all energy outcomes, are available in the replication package.