Burnout to Buildout: Coal Exit and the Geography of Transition Costs

Hao Zhao*

September 9, 2026

Abstract

Does exiting polluting assets necessarily deliver a green transition? Using coal-targeted federal rules as a nationwide shock to U.S. utility firms, I trace firm-level generation portfolios and show that coal exit is followed mainly by path-dependent substitution within fossil generation. Renewables instead expand where firms face binding renewable mandates and targeted financial incentives for specific technologies. I then build a stylized model to clarify the tradeoff between the fossil and renewable paths, and quantify two geography-related costs that arise in the policy-driven renewable transition. First, policy-induced renewable buildout can be geographically misaligned: firms earn \$290–\$300 less in annual electric operating revenue per kilowatt from solar projects in low-irradiance regions. Moreover, the renewable transition creates a systematic and overlooked social cost via increased workplace remoteness. Counterfactual analysis shows that reallocating renewable mandates toward higher-resource areas increases the solar generation share with fewer solar projects and reduces workplace remoteness without increasing emissions.

Keywords: energy transition; polluting assets; geographic misalignment; E vs. S dilemma

JEL classifications: D22, G31, Q42, Q53, Q58, L94

*Finance Department, Durham University Business School (email: hao.zhao@durham.ac.uk). I am indebted to Xian Gu, Felix Irresberger, and Louis Nguyen for their continuous guidance and support. I am grateful to Xingchen Zhu for the detailed and helpful discussion, and to Mayank Kumar, Chenyan Lyu, and Thomas Schmid for their field insights and suggestions. I also thank Kumushoy Abduraimova, Rajdeep Chakraborti, Haifeng Guo, Danny McGowan, Dimitris Petmezas, Zacharias Sautner, Enrique Schroth, Peter Tufano, Chendi Zhang, Shu Zhang, Junli Zhao, Yang Zhao, and Jiaqi Zheng, as well as organizers and participants at the Durham Finance Internal Seminar, the 2026 Exeter Sustainable Finance Conference, and the 2026 China International Conference in Finance (CICF) for helpful comments and suggestions and for the opportunities to present this paper.

1 Introduction

Coal power plants supplied more than half of U.S. electricity in 2002 but only 16.2 percent in 2023. During this period, the renewable share rose from 8.9 percent to 21.4 percent.[1] This transition reflects both tighter regulatory constraints on emissions-intensive generation and the increasing commercial viability of clean generation technologies such as solar photovoltaics, a policy–technology combination that Acemoglu et al. (2016) model as capable of overcoming fossil-fuel inertia in the energy sector.

In the U.S., policy targeting coal generation is national, but the adjustment problem is local. As shown in Figure 1, when facing a coal-targeted policy shock, major electric utilities are still expected to meet local demand and reliability by rotating their generation portfolios rather than shrinking them. Prior work mainly focuses on the coal-exit margin at the generation-unit level (e.g., Fell and Kaffine, 2018; Linn and McCormack, 2019; Gowrisankaran et al., 2025). I provide a new angle by studying this adjustment through firms' generation asset portfolios, treating coal as one component alongside other technologies. Importantly, this design allows the firm-level transition to be further tracked through changes in the portfolio's relative composition.

A natural question follows: if exiting coal is a step away from "brown" assets, does the transition necessarily make utility firms "green"? I dissect this process in three stages: the phaseout of coal assets, the subsequent (green) transition in firms' generation portfolios, and the outcomes of these decisions. I address three questions: What regulatory and economic forces drive coal exit and shape firms' exit paths? How do firms redirect generation toward cleaner technologies under regulatory and geographic constraints? Given the tradeoffs created by these constraints, what are the financial and social costs of the resulting renewable transition?

To empirically answer these questions, I combine utility-plant-generator-level operational data from the U.S. Energy Information Administration (EIA) and utility-level financial data from the Federal Energy Regulatory Commission (FERC), with additional measures mainly including regulatory exposure to emissions and renewable requirements, financial incentives, and geospatial resource data from 2007 to 2023. The analysis proceeds in three parts.

The first part of the analysis examines how firms' generation portfolios adjust after two major federal air pollution regulations targeting coal-linked pollutants: CSAPR (proposed in 2010 for eastern states) and MATS (proposed nationwide in 2011). The interval from proposal to implementation is about five years in both cases. While CSAPR allows allowance trading, MATS imposes non-tradable unit-level standards that further raise compliance costs. Using a stacked dynamic event study framework, I find that coal-reliant firms significantly reduce the coal portfolio share after the two regulations take effect, with a corresponding increase in the gas-fired share (fossil fuel path dependence) and no significant change in the renewable share.

On the mechanisms of coal exit, I show that both divestiture and closure rates rise at the firm level after the two regulations. To interpret firms' choice between these margins, I build on Darmouni and Zhang (2024) and extend the model to allow local gas–coal fuel cost advantages to affect the coal plant asset market. Where gas has a larger cost advantage, buyers' willingness

[1] Source: U.S. Energy Information Administration (EIA), *Total Energy: Annual Energy Review*, https://www.eia.gov/totalenergy/data/annual/.

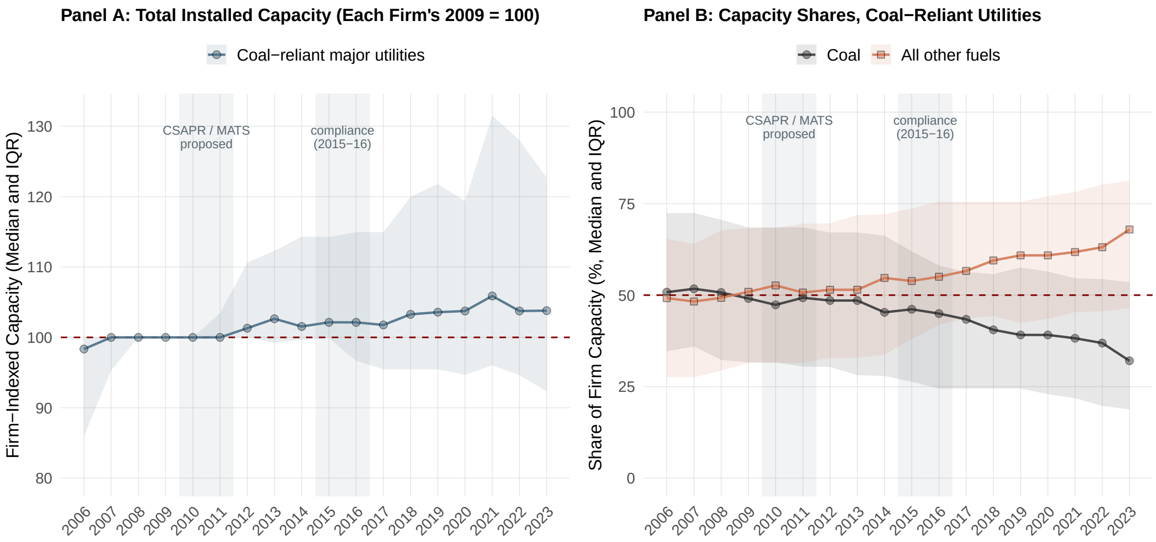

Figure 1. Continued Service and Coal Exit of Major Utilities. This figure documents that utilities exposed to the coal-targeted air pollution regulations kept serving their markets at their pre-policy scale while rotating generation capacity away from coal. The sample comprises the coal-reliant major utilities (FERC Form 1 filers) with complete EIA capacity records over 2006–2023; "coal reliance" is defined before the proposal of coal-targeted regulations (CSAPR and MATS). Capacity is ownership-weighted nameplate capacity aggregated to the firm level. Panel A indexes each firm's total capacity to its own 2009 level (converted to 100) and plots the cross-firm median with the interquartile range. Panel B plots the cross-firm median and interquartile range of the firm-level shares of coal versus all other fuels in total capacity; the two series are complements and sum to 100% for each firm. In both panels, shaded vertical bands mark the proposal window of CSAPR and MATS (2010–2011) and the compliance window (2015–2016, including one-year MATS extensions); dashed horizontal lines mark the pre-policy index level (Panel A) and the equal-share benchmark of 50% (Panel B).

to pay for coal plants falls, making more potential transactions on the equilibrium pricing line shift into the retirement area. Consistent with this implication, I isolate a substitution effect: a one-standard-deviation increase in the coal–gas fuel cost gap raises the probability of closure by about five percentage points, and has no significant effect on divestiture. Correspondingly, coal capacity declines by an amount comparable to the increase in gas-fired capacity.

In the second part of the analysis, I focus on the redirection dimension of the transition: how firms redirect capital toward renewable generation as they exit coal. To isolate the investment response induced by renewable policy, I estimate a fuzzy regression discontinuity design at the annual targets in state Renewable Portfolio Standards (RPS). I find that firms positioned just below the RPS targets increase construction-related capital expenditures, and this additional spending translates into higher renewable generation portfolio shares in subsequent years. A ten percent increase in instrumented construction CapEx raises the share of renewables in the following year by approximately half a percentage point, with similar cumulative effects observed over a three-year horizon. Falsification tests using gas and other non-renewable clean energy generation as outcomes show no comparable effects, and standard diagnostics support the validity of the local identification around the RPS thresholds.

I next examine firms' technology choices within the broader renewable transition. A key distinction in policy design is that regulatory mandates such as RPS often target an overall increase in renewable generation, whereas financial incentives tend to be structured around specific technologies. Given this design difference, I isolate how technology-specific incentives shape investment across renewable options. Using an event-study specification, I find that broad renewable incentives raise the overall share of renewable capacity, while solar-specific and wind-specific programs expand their respective technologies, without affecting untargeted generation categories such as gas or other clean capacity.

The third part of the analysis begins from a key implication of the earlier results: if policies can reshape generation portfolios by retiring coal and promoting renewables, then policies that do not fully account for geographic variation in renewable potential may push new renewable capacity into sites with less favorable conditions. Because renewable output depends critically on local natural conditions, siting plants in low-quality resource areas can reduce their operational efficiency. I find such geographic misalignment is present: though the adoption of solar projects is most concentrated in areas with very high solar irradiance, a substantial share of investment occurs in regions with only moderate resource quality. This misalignment has economically meaningful consequences: firms that build solar plants in areas with irradiance below National Renewable Energy Laboratory Class 3 standards earn around \$290–\$300 less in electric operating revenue per kilowatt of installed capacity than firms that invest in higher-quality locations.

Although state-level policies account for much of the geographic misalignment in renewable energy siting, they do not fully determine firms' local siting decisions. Conditional on similar external policy conditions, firms are expected to choose sites based on engineering and economic criteria such as resource quality, grid access, and land costs. Within these tradeoffs, I find that firms tend to overlook a social consideration that matters for workers but not directly for plant operations: the remoteness of the site from population centers. I show that utility-scale solar projects are systematically sited in less populated areas compared to the same firms' earlier fossil fuel plants. This pattern is robust across alternative observation radii around plants, within state-year variation, and under different sample matching strategies. Because the specification directly compares each firm's renewable projects to its own historical fossil plants, this evidence specifically implies a friction for workers retrained from fossil to renewable positions.

Finally, I use structural estimation to examine how reallocating RPS targets toward higher-irradiance areas affects electricity generation and solar project adoption, and assess the implications for the financial cost of geographic misalignment and the social cost of workplace remoteness. I model firms' investment choices between two representative technologies, solar and gas, using a static discrete choice framework in which payoffs are a parsimonious function of geographic resource quality, policy and market conditions, intermittency costs, capital and operating costs, and firm-specific path dependence and project scale. The estimates show how geography, policy, and path dependence jointly shape the transition. Higher irradiance and stronger policy incentives favor solar, while existing fossil portfolios and large project requirements favor gas. Policy effects are strongest earlier in the transition period, when solar faced a larger capital-cost disadvantage. Two counterfactual scenarios illustrate the policy tradeoff. The environmentally focused scenario raises solar adoption and the solar generation share, and lowers emissions at the cost of slightly greater workplace remoteness. The socially focused scenario leaves solar adoption approximately unchanged, and reduces workplace remoteness without increasing emissions.

This study first contributes to the literature on how air pollution regulations affect firms' fossil generation and reallocation decisions. Prior research shows that emission regulations such as the Clean Air Act, NOx trading programs, and the Regional Greenhouse Gas Initiative change plant-level pollution controls and impose compliance costs (e.g., Becker and Henderson, 2000; Greenstone, 2002; Fowlie et al., 2012; Shapiro and Walker, 2018; Kumar and Purnanandam, 2023). I extend this literature by decomposing the firm-level divestiture and closure margins, and isolating a fuel-cost substitution mechanism: firms reduce coal-fired generation and shift toward natural gas more in markets where the coal–gas fuel cost gap is larger under CSAPR and MATS. These findings relate most closely to Fell and Kaffine (2018) and Linn and McCormack (2019) on the central role of gas prices in driving the coal exit. The analysis of coal exit and technology transition choice also relates to work on policy uncertainty around MATS (Gowrisankaran et al., 2025), production flexibility under demand uncertainty (Lin et al., 2026), and the role of private equity ownership in retirement responses to gas pipeline expansions (Kumar, 2025).

Second, I quantify a firm-level financial cost from geographic misalignment in renewable investment, contributing to the literature on operational inefficiencies associated with key features of renewable technologies, such as intermittency (Gowrisankaran et al., 2016), location-specific grid value (D. S. Callaway et al., 2018), and the start-up costs of thermal plants when solar output declines in the evening (Jha and Leslie, 2025), etc. I show that siting utility-scale solar plants in lower-irradiance regions is associated with materially lower electric operating revenue per kilowatt of capacity, and this revenue gap is robust across alternative irradiance thresholds and firm groups. This evidence translates geographic misalignment into a direct firm-level revenue loss.

Third, I document an "environmental versus social" dilemma within the renewable energy transition, complementing studies that mostly focus on the social costs of fossil fuels where pollution generates negative externalities such as declines in local housing values and labor earnings (e.g., Davis, 2011; Isen et al., 2017), and how social factors can partly offset coal plant closures under environmental pressure (e.g., Darmouni and Zhang, 2024). I find that, even under similar policy and market conditions, firms tend to locate utility-scale solar projects in more remote and less populated areas than their prior fossil fuel plants. From the perspective of firms' ESG strategies, prioritizing the "E" dimension via renewable transition comes at the cost of the "S", since the local optimization of renewable siting tends to overlook workplace remoteness. This finding also relates to the discussion on how renewable energy transition affects workforce reallocation and earnings (Walker, 2013), local employment opportunities (Fabra et al., 2024), regional development under geographically targeted subsidies (Zhang, 2025), biodiversity impacts of renewable siting (Gong et al., 2025), and amenity costs from siting near residential communities (Jarvis, 2025).

Finally, this study contributes to policy design for a more efficient renewable energy transition. The counterfactual analysis shows that better aligning renewable mandates with resource quality can mitigate the environmental-social tradeoff. Policy pressure can be reallocated to reduce either emissions or workplace remoteness without materially worsening the other outcome. These results contribute to the literature proposing policies for a more efficient renewable energy transition, such as carbon pricing paired with incentives that direct innovation toward clean technologies and help overcome path dependence (e.g., Acemoglu et al., 2012; Acemoglu et al., 2016; Aghion et al., 2016).

The paper proceeds as follows. Section 2 reviews the institutional details. Section 3 describes the data and sample. Section 4 analyzes the coal exit and renewable transition channels. Section 5 examines geographic misalignment in renewable investment and its financial costs. Section 6 quantifies the social costs of renewable adoption. Section 7 estimates the static choice model and presents counterfactuals. Section 8 concludes.

2 Institutional Details

This section introduces the major federal air pollution regulations that target pollutants from coal power generation, the state-level Renewable Portfolio Standards (RPS) that require a minimum share of electricity from eligible renewables, and the technical constraints on energy transition that underlie my identification strategy. The focus is on institutional features that directly shape exit and transition choices. For additional background on the U.S. power sector, electricity markets and the energy transition, see Rauh and Garcia (2026).

2.1 Policy Driver for the Exit of Pollutive Operation

The Environmental Protection Agency (EPA) proposed and finalized two major rules that set binding limits on pollutants emitted by fossil fuel–fired power plants:

Cross-State Air Pollution Rule (CSAPR). The EPA proposed CSAPR on July 6, 2010 (Federal Register, 75 FR 45210) and finalized the rule in 2011, with the final rule published on August 8, 2011 (Federal Register, 76 FR 48208). CSAPR targets emissions of sulfur dioxide (SO_2) and nitrogen oxides (NO_x) that contribute to ozone and particulate matter ($PM_{2.5}$) pollution in downwind states. It sets state-specific caps enforced through allowance trading, which allows utilities to meet obligations either by abating emissions or by purchasing allowances. The proposal initially designated 32 eastern states as subject to the requirements, while the final rule applied to 27 states. Within the annual SO_2 program, states were divided into two groups with different stringency, where group 1 faced stricter two-phase reduction requirements and group 2 faced more moderate requirements. Although implementation was delayed until 2015 due to legal challenges, the 2010 proposal created immediate transition risk for coal-reliant utilities by signaling forthcoming binding constraints. Therefore, I define 2010 as the introduction year (event time $\tau = 0$), corresponding to the proposal for targeted states, and 2011 as the adoption year (event time $\tau = 1$), when the EPA finalized the rule and established binding compliance expectations.

Mercury and Air Toxics Standards (MATS). The EPA proposed MATS on May 3, 2011 (Federal Register, 76 FR 24976) and finalized the rule on February 16, 2012 (Federal Register, 77 FR 9304). MATS introduced the first nationwide emission standards for hazardous air pollutants, including mercury, acid gases such as hydrogen chloride, and nonmercury metals, from coal and oil fired electric generating units. Compliance required substantial upgrades such as activated carbon injection and advanced filtration. New units had to comply upon startup or by the effective date in 2012; existing units were given up to three years to comply (to April 2015), with limited one-year extensions to 2016, and in select cases to 2017 for reliability. Given the cost of controls and the age and scale of many coal units, retirements often dominated retrofits. Similarly, the 2011 proposal triggered immediate transition risk by signaling strict new limits on hazardous air pollutants from coal-fired power plants. Thus, I define 2011 as the introduction year (event time $\tau = 0$), and 2012 as the adoption year (event time $\tau = 1$) when the EPA finalized the rule.

Overlapped Timelines. The timelines of CSAPR and MATS largely overlap. CSAPR was proposed for eastern states in 2010, creating immediate compliance expectations for coal-reliant utilities; one year later, MATS was proposed nationwide. Given that both policies shock coal generation, I define treated status as firms that own coal assets and are exposed to the proposal of either CSAPR or MATS. Specifically, I implement a stacked event study design with two cohorts: (i) firms exposed to CSAPR (and subsequently to MATS) form one cohort, and (ii) firms exposed only to MATS form a second cohort. I then stack the cohort-specific windows by event time relative to the proposal year, which keeps the pre-treatment period clean (before proposal, no compliance is expected). Because binding compliance for both rules arrives roughly five years after proposal, I further estimate dynamic treatment effects and assess whether effects concentrate around the implementation years.

Other Earlier Programs. In this study, I focus on CSAPR and MATS as the primary shocks to coal power generation. Among earlier relevant federal emission rules, the major programs were the NO_x Budget Trading Program (NBP) and the Clean Air Interstate Rule (CAIR). NBP is a cap-and-trade program launched in 2003 and applied only in the ozone season (May–September) for eastern states, so it lacked the non-tradable and year-round constraints that characterize MATS's unit-level hazardous-air-pollutant standards and CSAPR's binding state emission budgets. CAIR was finalized by the EPA in 2005 to address interstate transport of SO_2 and NO_x via regional cap-and-trade, and legally challenged in 2008 before the start of its NO_x and SO_2 programs. Eventually, the EPA finalized CSAPR in 2011 to replace CAIR and address the court's objections.[2]

2.2 Policy Pressure on Renewable Power Transition

Renewable Portfolio Standards (RPS) are state-level policies that require utilities to supply a minimum share of their electricity from eligible renewable sources such as wind, solar, geothermal, and biomass. As of 2023, 28 states and the District of Columbia have binding RPS mandates, and 7 additional states have voluntary renewable portfolio goals.[3] Compliance is enforced through Renewable Energy Certificates (RECs), which allow utilities to meet

[2]For program details of NBP, see U.S. Environmental Protection Agency (EPA), *NO_x Budget Trading Program Historical Reports*, https://www.epa.gov/power-sector/nox-budget-trading-program-historical-reports. For CAIR, see U.S. Environmental Protection Agency (EPA), *Clean Air Interstate Rule Historical Reports*, https://www.epa.gov/power-sector/clean-air-interstate-rule-historical-reports.

[3]Source: U.S. Energy Information Administration (EIA), *Renewable Energy Explained: Portfolio Standards*, https://www.eia.gov/energyexplained/renewable-sources/portfolio-standards.php.

requirements either by building renewable capacity or by purchasing credits. Noncompliance results in financial penalties or alternative compliance payments. By design, most RPS programs become more stringent over time, creating durable obligations that shape utilities' investment strategies and generation portfolios. Unlike federal subsidies that reduce financing costs, RPS mandates impose direct operating requirements on utilities and systematically shift their generation portfolios toward renewables.

2.3 Technical Constraints on Transition Decisions

I introduce three technical constraints that form the basis of my empirical identification strategy. These constraints limit where new generation can connect to the grid and how far firms can reallocate production across locations.

First, compliance across states is constrained by the transmission system. Moving large volumes of power over very long distances requires dedicated high-voltage lines and available capacity, and such corridors remained scarce during my sample period.[4] In practice, this limits the ability of a utility facing a strict mandate in one state to meet that obligation by selling fossil generation into distant states with weaker or no mandates, or by siting large volumes of renewable capacity in remote high-resource regions that are poorly connected to its load. For identification, I rely on the idea that most compliance decisions must be made within the footprint of the relevant regulated market or balancing authority, which narrows firms' strategic options and makes local regulatory changes informative about their behavior.

Second, investment can be misallocated when high interconnection costs make the best resource sites uneconomic. Some locations have excellent solar or wind conditions but are weakly connected to existing transmission infrastructure. In these cases, investment decisions depend not only on natural resource quality but also on the cost of delivering electricity to the network. A firm that appears well positioned geographically may still find its highest-quality sites unattractive once grid access costs are included. In practice, this weakens the simple link between resource potential and actual investment, because promising projects can be delayed, scaled down, or dropped from the queue when interconnection costs become too large.[5] [6]

Third, grid reliability requirements limit how far firms can replace conventional generation with renewables within the period I study. Solar and wind output is variable, and in current U.S. practice these resources are not yet operated as the primary providers of voltage and

[4]High-voltage direct current (HVDC) technology provides an efficient option for long-distance transmission with lower line losses than conventional alternating-current lines, but it was sparsely deployed in the United States over my sample period. As of 2024, only five HVDC transmission lines were in operation. See National Renewable Energy Laboratory (NREL), *On the Road to Increased Transmission: High Voltage Direct Current* (2024), at https://www.nrel.gov/news/program/2024/on-the-road-to-increased-transmission-high-voltage-direct-current.html.

[5]See Energy Technologies Area, Berkeley Lab, *Queued Up: Characteristics of Power Plants Seeking Transmission Interconnection*, 2024, available at https://emp.lbl.gov/queues.

[6]Project-level interconnection costs are not modeled in the later analysis of policy-induced geographic misalignment, where the main variation is across states rather than across individual projects. In the social cost analysis, I compare each firm's renewable siting to its own previous siting, so interconnection costs affect both locations similarly. For related work on transmission constraints and the efficiency outcomes of limited transmission capacity, see, e.g., Fell et al. (2021), Davis et al. (2023), Hausman (2025).

frequency support at system scale.[7] To meet system-wide reliability standards set by the North American Electric Reliability Corporation (NERC) and regional entities, planners must retain a portfolio of dispatchable resources, including gas and coal units, that can stabilize the grid in real time and respond to contingencies. Under these conditions, it is not realistic over my sample period for the firms in a region to collectively eliminate conventional generation and rely only on renewables with limited storage, even if such a portfolio appeared attractive on a pure energy cost basis.

3 Data and Sample

I begin with utility-level data from the U.S. Energy Information Administration (EIA), which report detailed ownership and capacity of electricity generation assets. I link these data with the Federal Energy Regulatory Commission (FERC) to obtain financial information and construct the main sample covering 2007–2023.[8] [9]

Portfolio of Generation Capacity. Generator-level data from the EIA Form 860 Annual Electric Generator Report provide information on capacity, fuel type, ownership shares, and operational status. I use these data to construct ownership-weighted aggregate generation capacity by firm and year. These portfolio measures allow me to identify coal-reliant firms that are more exposed to air pollution regulation and to analyze how utilities adjust ownership across generation types. Plant closures are identified using generator operational status codes.

Renewable Requirements. Data on Renewable Portfolio Standards (RPS) come from Lawrence Berkeley National Laboratory. At the state-year level, the dataset provides the required share of renewable generation in total electricity supply. I use the "Total RPS" target for the main analysis, and incorporate carve-outs, which specify technology-specific minimums, as an additional source of variation in renewable obligations.

Financial Incentives. The Database of State Incentives for Renewables & Efficiency (DSIRE) provides a comprehensive collection of renewable energy program information. I focus on financial incentives implemented at the state level, and classify programs by technology category and detailed program type, such as loans, bonds, feed-in tariffs, and performance-based incentives.

Geospatial Resources. I measure renewable resource quality using ERA5 reanalysis data from the European Centre for Medium-Range Weather Forecasts (ECMWF), which provide globally consistent climate variables at an 11 km resolution. For each plant, I construct solar potential from surface solar radiation downwards, equivalent to global horizontal irradiance (GHI), and

[7] Battery storage and other balancing technologies existed during the sample period but were costly and deployed at relatively small scale. By 2023, large scale battery capacity was about 16 GW, compared with roughly 388 GW of U.S. utility scale renewable generation capacity. See U.S. Energy Information Administration (EIA), *Battery Storage in the United States: An Update on Market Trends*, available at https://www.eia.gov/analysis/studies/electricity/batterystorage/.

[8] Consistent state-level demand data from FERC Form 714 are available beginning in 2006; the use of lagged controls restricts estimation to 2007 and later.

[9] The linkage and data access for EIA and FERC are provided through PUDL, a data pipeline developed by Catalyst Cooperative that cleans, integrates, and standardizes widely used U.S. energy datasets. See: https://catalystcoop-pudl.readthedocs.io/en/latest/.

wind potential from the east–west and north–south components of near-surface wind, combined into wind speed.[10] These measures capture geographic variation in renewable resources that condition firms' investment choices.

Site and Workforce Accessibility. To examine the geographic accessibility of new generation projects, I measure both their proximity to population centers and their proximity to workforce communities associated with retiring coal plants. First, I use Oak Ridge National Laboratory's LandScan population data, which provide global settlement patterns at a 30 arc-second resolution (around 1 km), and calculate population density within specified ranges around each plant.[11] [12] Second, I use the U.S. Census Bureau's Longitudinal Employer-Household Dynamics Origin-Destination Employment Statistics (LODES), which provide block-level workplace employment and job flows between workplace and residence locations. I use these data to characterize the residential geography associated with utility workplaces near retiring coal plants, as well as utility-sector employment around newly built projects.

Firm-Level Financial Information. To study investment and performance, I construct utility-level financial variables from FERC Form 1, including construction capital expenditure, total assets, depreciation, O&M costs, leverage, ROA, and liquidity.[13] I supplement this with electric operating revenue from the EIA Form 861 Annual Electric Power Industry Report, fuel costs from the EIA Form 923 Power Plant Operations Report, and abatement costs from the EIA Form 860 Annual Electric Generator Report.

Other Firm-Level and State-Level Information. I use FERC Form 714 for state-level electricity demand, EIA Electricity Retail Sales data for electricity prices, and emissions data on sulfur dioxide (SO_2), nitrogen oxides (NO_x), and carbon dioxide (CO_2) from EPA's Continuous Emissions Monitoring System (CEMS) to track compliance at the firm level.

Analysis Sample and Summary Statistics. I construct two main samples for the period 2007–2023. The first is a firm-year sample that merges EIA and FERC data, which provides standardized financial information and is used to analyze how utilities adjust generation portfolios in response to regulation and financial incentives. The second is a broader EIA-based sample, available at both the firm-year and plant-year level, which is used to identify firm investment behaviors around the RPS requirement discontinuity and to study project-level outcomes of post-coal-exit transition. In terms of coverage, the EIA sample follows Form EIA-860, which includes all U.S. power plants with at least 1 MW of combined

[10] Global horizontal irradiance (GHI) is the standard benchmark for solar resource assessment in power system planning. See National Renewable Energy Laboratory (NREL), *Measurement, Modeling, and Database of Solar Irradiance*, https://docs.nrel.gov/docs/fy20osti/75573.pdf. Wind speed is derived from ERA5 by combining the squared east–west and north–south 10-meter wind components. While these standardized measures differ from site-specific conditions such as panel orientation or turbine hub height, they provide consistent proxies for renewable resource quality across plants and time.

[11] LandScan integrates census data, satellite imagery, and geographic information through spatial modeling to produce population distribution estimates. See Lebakula et al. (2025) for methodology.

[12] Geospatial processing is conducted in Google Earth Engine (GEE) using plant coordinates from EIA.

[13] FERC Form 1 is a regulatory filing required of major electric utilities, collecting standardized financial and operational data. These reports are public and non-confidential.

Table 1: Summary Statistics

This table presents summary statistics for the variables used in the analysis from 2007 to 2023. Panel A reports firm-year observations from the merged EIA–FERC sample of major utilities. Panel B reports firm-year observations from the broader EIA sample used in the local fuzzy regression discontinuity design (FRDD) around RPS thresholds. Panel C reports plant-year-level observations from the EIA Form 860 sample, which includes all plants with at least 1 MW of capacity, restricted to fossil fuel, solar, and wind plants. All variables are defined in the Appendix.

Panel A. Firm-Year Level EIA–FERC Matched Sample

	N	Mean	Median	*SD*	P25	P75
$\text{PortfolioShare}^{(k)}_{i,t}$ (%)						
$k = Coal$	2,312	43.24	39.12	33.98	15.98	66.73
$k = Gas$	2,312	35.30	34.61	30.77	0.00	58.44
$k = Renewables$	2,312	9.31	0.00	19.22	0.00	9.74
$k = Other\ Clean$	2,312	4.81	0.00	13.48	0.00	0.10
$k = Solar$	2,312	0.33	0.00	1.74	0.00	0.00
$k = Wind$	2,312	1.69	0.00	4.94	0.00	0.00
CoalReliance_i (0–1)	2,312	0.46	0.42	0.34	0.18	0.71
InRPS_i ($\mathbb{1}$)	2,312	0.62	1.00	0.48	0.00	1.00
$\text{Closure}_{i,t}$ ($\mathbb{1} \times 100$)	2,312	6.36	0.00	24.41	0.00	0.00
$\text{Divestiture}_{i,t}$ ($\mathbb{1} \times 100$)	2,312	3.42	0.00	18.17	0.00	0.00
$\text{Abatement}^{\text{log}}_{i,t}$ (log m$)	2,312	1.78	0.00	2.75	0.00	3.86
$\text{Capacity}_{i,t-1}$ (log MW)	2,312	6.45	6.78	2.04	4.91	8.16
$\text{FuelCost}_{i,t-1}$ ($/MWh)	2,312	20.31	20.01	14.74	10.62	28.68
$\text{StateDemand}_{s(i),t-1}$ (log TWh)	2,312	4.28	4.29	0.83	3.92	4.80
$\text{ElectricPrice}_{s(i),t-1}$ (¢/kWh)	2,312	11.95	11.54	2.58	10.37	12.73
$\text{CarveOut}_{s(i),t-1}$ (ratio)	2,312	0.20	0.00	0.62	0.00	0.00
$\text{FirmSize}_{i,t-1}$ (log m$)	2,312	8.75	8.99	1.28	8.18	9.66
$\text{Leverage}_{i,t-1}$ (ratio)	2,312	0.32	0.31	0.09	0.28	0.34
$\text{ROA}_{i,t-1}$ (ratio)	2,312	0.03	0.03	0.02	0.02	0.04
$\text{Liquidity}_{i,t-1}$ (ratio)	2,312	0.003	0.00	0.05	-0.02	0.02
$\text{Depreciation}_{i,t-1}$ (ratio)	2,312	0.06	0.05	0.05	0.05	0.06
$\text{Revenue}^{\text{perMW}}_{i,t}$ (m$/MW)	2,078	1.22	0.52	2.36	0.39	0.82

Panel B. Local FRDD Firm-Year Level EIA Sample

	N	Mean	Median	*SD*	P25	P75
$\text{PortfolioShare}^{(k)}_{i,\text{forward}}$ (%)						
$k = Renewables; [t+1]$	1,277	3.68	0.00	8.94	0.00	2.40
$k = Renewables; [t+1, t+3]$	1,277	4.04	0.00	9.68	0.00	2.97
$k = Gas; [t+1, t+3]$	1,277	35.66	35.41	32.75	0.00	59.99
$k = Other\ Clean; [t+1, t+3]$	1,277	5.88	0.00	18.10	0.00	0.00
$\text{ConstructCapEx}_{i,t}$ (log m$)	1,277	5.49	5.72	1.52	4.72	6.67
$\text{BelowRPS}_{i,t}$ ($\mathbb{1}$)	1,277	0.47	0.00	0.50	0.00	1.00

Panel C. EIA Plant-Year Level Sample

	N	Mean	Median	*SD*	P25	P75
$\text{Solar}_{p,t}$ (kWh/m^2/day)	215,541	4.56	4.36	0.70	4.03	5.05
$\text{Wind}_{p,t}$ (m/s)	215,541	1.15	1.02	0.52	0.83	1.33
$\text{Population}^{(b)}_{p,t}$ (count/km^2)						
$b = 0\text{–}10\ km$	215,541	247.57	42.09	689.94	9.09	180.12
$b = 10\text{–}25\ km$	215,541	186.21	37.70	450.29	9.75	154.04
$b = 25\text{–}60\ km$	215,541	126.06	43.19	214.94	13.91	138.19
$b = 60\text{–}100\ km$	215,541	96.90	46.03	130.15	15.91	118.94

nameplate capacity. By contrast, the merged EIA–FERC sample is restricted to entities that file FERC Form 1, classified as "major utilities".[14]

Table 1 reports summary statistics for the main variables. Panel A summarizes the firm-year-level EIA-FERC matched sample for 189 major electric utilities. Coal dominates the generation portfolios of major utilities, accounting for 43% of capacity on average, compared to 35% for natural gas and only 9% for renewables. Solar and wind remain negligible, together contributing less than 2%, and most firms report no renewable capacity at all. Before the air pollution regulations CSAPR and MATS, coal made up an average of 46% of generation capacity in the sample portfolio, with about 62% of firms operating plants in states subject to RPS requirements. This variation provides scope to compare firms with high, low, and no coal reliance in their compliance response, and to examine how pre-policy RPS mandates shape the direction of their subsequent transition.

Panel B reports summary statistics for the broader EIA sample of 295 utilities used in the fuzzy regression discontinuity design around RPS thresholds. Forward-looking portfolio shares show that renewable capacity averages about 4% within three years, below the 9% renewable share observed in the EIA–FERC sample. About 47% of firm-years fall just below the relevant RPS cutoff, indicating a balanced distribution around the threshold.

Panel C reports plant-year-level statistics from the EIA Form 860 sample for 11,841 plants owned by 8,485 utilities, restricted to fossil fuel, solar, and wind plants to enable cleaner later empirical comparisons across transition types. Solar irradiance averages 4.56 kWh/m^2/day and wind speed 1.15 m/s. Population density around plants is highly skewed, for instance, with a median of 42 people/km^2 within 10 km but a mean close to 250, and in the 25–60 km band, a median of 43 compared to a mean of 126.

4 What Drives Exit and How Firms Redirect

This section examines two response dimensions of coal-dependent utilities following the federal air pollution regulations CSAPR and MATS. The first is compliance. Firms operating coal-fired generation may reduce coal exposure by selling assets, accelerating retirements, or installing abatement equipment to continue operations. Beyond direct compliance costs, relative fuel prices after the regulations are also expected to shape these decisions: higher coal costs increase the burden of maintaining coal units, while the relative affordability of gas may make a shift

[14] Form EIA-860 collects generator-level data on existing and planned units and associated environmental equipment at plants with at least 1 MW of nameplate capacity, providing broad coverage of U.S. power plants. FERC Form 1 is filed annually by "major" electric utilities, defined as those that, in each of the three previous calendar years, exceeded one of the following thresholds: (1) one million megawatt hours of total annual sales, (2) 100 megawatt hours of annual sales for resale, (3) 500 megawatt hours of annual power exchanges delivered, or (4) 500 megawatt hours of annual wheeling for others (deliveries plus losses). This narrower coverage provides richer firm-level financial detail. The broader renewable buildout also involves substantial investment by private equity, institutional investors, and foreign corporations (Andonov and Rauh, 2025), with project finance and long-term power purchase agreements supporting investment in renewable generation (Darmouni et al., 2026). The merged sample follows major electric utilities' ownership of generation assets across technologies, rather than electricity purchased from third-party generators.

Table 2: Transition Matrices of Firm-level Generation Portfolios

This table reports three-year transition probabilities across firm-level generation portfolio categories before and after the adoption of the Cross-State Air Pollution Rule (CSAPR) and the Mercury and Air Toxics Standards (MATS). Columns (1)–(4) show transition probabilities in the pre-policy period; columns (5)–(8) report transition probabilities in the post-policy period. Portfolio categories include: (1) coal only; (2) gas or oil, no renewables; (3) non-renewable, non-fossil (nuclear or waste); and (4) any renewable capacity. Each row reports destination distributions for firms starting in a given category.

Portfolio category at $t+3$:	*1*	*2*	*3*	*4*	*1*	*2*	*3*	*4*
Subsample period:	*Pre-policy period*				*Post-policy period*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Portfolio at t: *1* (Coal only)	100.00%	0.00%	0.00%	0.00%	87.05%	8.48%	1.34%	3.12%
2 (Gas/oil, no renewables)	1.72%	82.76%	3.45%	12.07%	1.56%	87.30%	1.56%	9.57%
3 (Nuclear/waste)	0.00%	2.27%	84.09%	13.64%	0.82%	10.66%	77.05%	11.48%
4 (Any renewables)	0.00%	1.57%	0.52%	97.91%	0.25%	1.51%	0.38%	97.86%

toward gas capacity more attractive. To assess policy effectiveness, I further test whether these compliance strategies translate into measurable reductions in targeted air pollutants.

The second dimension is redirection of generation. Path dependence is expected: because of technological and operational continuity, coal-dependent firms may favor shifting toward less pollutive fossil fuel generation such as natural gas. In contrast, state RPS mandates can act as external drivers that break this path dependence and push firms toward renewables. The evolving annual RPS targets create a regulatory discontinuity, which I use to examine whether firms just below the threshold increase construction-related capital expenditures and expand their renewable portfolios in subsequent years. Finally, I dissect the renewable transition by technology type, assessing the role of financial incentives in steering investment toward specific renewable assets.

4.1 Responses to Regulation: Path Dependence and Renewable Redirection

To provide an overview of generation portfolio movements around the CSAPR and MATS, Table 2 reports three-year transition probabilities across firm-level portfolio categories before and after the regulation. The categories are: (1) coal-only firms; (2) firms with gas or oil but no renewables; (3) firms with non-renewable, non-fossil capacity (nuclear or waste); and (4) firms with any renewable generation including solar, wind, or hydro. The left panel shows transition probabilities during the pre-policy period (event time $\tau < 0$), while the right panel shows those during the post-policy period (event time $\tau \geq 0$).

In the pre-policy period, coal-only firms showed complete persistence: 100% remained coal-only after three years. By contrast, firms starting with gas/oil or with nuclear/waste capacity already displayed some diversification into renewables, with adoption rates of 12.1% and 13.6%, respectively. In the post-policy period, coal-only firms begin to diversify: 8.5% transition toward gas/oil and 3.1% add renewables. Gas/oil and nuclear/waste firms also continued to adopt renewables, though at slightly lower rates (9.6% and 11.5%) than in the pre-policy period. Across both periods, entry into coal is rare, with rates below 2%. Overall, the comparison indicates that the air pollution regulations primarily loosen the persistence of coal-only portfolios and accelerate the diversification.

To formally assess how firms adjust their electricity generation portfolios in response to the air pollution regulations, I implement a difference-in-differences (DiD) framework that leverages variation in regulatory timing and firms' pre-policy exposure. The design exploits the fact that CSAPR was proposed and adopted one year earlier than MATS, with CSAPR binding differentially across states while MATS later applied nationwide. I combine this variation with firm-level heterogeneity in pre-policy coal dependence and RPS exposure. This framework distinguishes path-dependent adjustments reflected in substitution from coal to other less pollutive fossil fuels, from redirection toward renewables under pre-existing state renewable requirements. I start by estimating the following specification:

$$\begin{aligned}\text{PortfolioShare}^{(k)}_{i,t} = \alpha_0 + \beta\left(\text{Exposed}_i \times \text{PostPolicy}^{\text{Air}}_{i,t}\right) \\ + X'_{i,t-1}\gamma + Z'_{s(i),t-1}\delta + \theta_i + \lambda_\tau + \varepsilon_{i,t}.\end{aligned} \tag{1}$$

The dependent variable $PortfolioShare^{(k)}_{i,t}$ denotes the share of generation type k, including coal, gas, renewables, and other clean fuels, in firm i's total generation capacity in year t.[15] It is calculated as the ratio of owned capacity in type k to total owned capacity, as:

$$\text{PortfolioShare}^{(k)}_{i,t} = \frac{\sum_{j\in k} \text{GeneratorSize}_{j,t} \times \text{Ownership}_{i,j,t}}{\sum_{j} \text{GeneratorSize}_{j,t} \times \text{Ownership}_{i,j,t}} \times 100, \tag{2}$$

where j indexes generation units, and generator j's nameplate capacity ($GeneratorSize_{i,t}$) is weighted by the firm's ownership share. The time indicator $PostPolicy^{Air}_{i,t}$ equals one if firm i's state has adopted the air pollution regulations (from 2011 for CSAPR, or from 2012 for MATS) by year t. The group indicator $Exposed_i$ is proxied by two separate conditions: (i) the firm's pre-policy reliance on coal generation ($CoalReliance_i$), given coal's central role in the emissions targeted by the policy; (ii) whether the firm operated in a state with a Renewable Portfolio Standard (RPS) adopted prior to the regulations ($InRPS_i$). $CoalReliance_i$ is calculated as the ownership-weighted share of coal generation capacity for firm i at event time $\tau = -1$, constructed from equation (2) without multiplying by 100. $InRPS_i$ is an indicator equals one if the firm is located in an RPS state at event time $\tau = -1$.

The vector $X_{i,t-1}$ includes lagged firm-level controls: *FirmSize* (log of total assets plus one), *Leverage* (total debt over total assets), *ROA* (net profit over total assets), and *Liquidity* (working capital ratio, current assets minus current liabilities over total assets). To capture natural asset retirement, I also include *Depreciation* (accumulated depreciation over gross plant assets). Generation-related variables include *Capacity* (log of total installed capacity in MW plus one) and *FuelCost* (USD/MWh, capturing variation in input price pressure). The vector $Z_{s(i),t-1}$ contains lagged state-level characteristics, averaged across the states in which firm i operates. These include *StateDemand* (log of market demand in TWh plus one, proxying for

[15] For the analysis, I focus on two major fossil fuel types: coal, which is the most pollutive, and gas, which is comparatively less pollutive. Renewables include solar, wind, and hydro. Other clean generation refers to nuclear and waste-based sources.

Table 3: Firms' Responses to Regulatory Pressures

This table reports the effect of federal air pollution regulations on the composition of firms' electricity generation portfolios. The dependent variable $PortfolioShare^{(k)}_{i,t}$ denotes the share of generation type k (coal, gas, renewables, or other clean) in firm i's total owned capacity in year t, computed by weighting each generator's capacity by the firm's ownership share. $PostPolicy^{Air}_{i,t}$ equals one if firm i's state has adopted the air pollution regulations (CSAPR or MATS) by year t. Panel A uses $CoalReliance_i$, the pre-policy share of coal capacity, to measure exposure intensity to the regulations. Panel B uses $InRPS_i$, an indicator equal to one if the firm operated in a state with a Renewable Portfolio Standard (RPS) adopted before regulation. All regressions include firm fixed effects (θ_i) and event-time fixed effects (λ_τ), where $\tau = t - g_i$ denotes the number of years since firm i first became subject to the regulations, with cohort g defined as the year of initial exposure. Firm-level controls include lagged values of $FirmSize$, $Leverage$, ROA, $Liquidity$, $Depreciation$, $Capacity$, $FuelCost$, and $PortfolioShare^{(k)}_{i,t-1}$ (for $k \neq$ coal, Panel A only); State-level controls are averaged at the firm level, including lagged $StateDemand$, $ElectricPrice$, and $CarveOut$. All variable definitions are in the Appendix. Standard errors are clustered at the firm level. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Panel A. Exiting Coal and Path Dependence

Dep. Var.: $\text{PortfolioShare}^{(k)}_{i,t}$				
$k =$	*Coal*	*Gas*	*Renewables*	*Other Clean*
	(1)	(2)	(3)	(4)
$\text{CoalReliance}_i \times \text{PostPolicy}^{\text{Air}}_{i,t}$	-9.503***	3.029***	1.121	0.434
	(3.622)	(1.073)	(1.059)	(0.470)
Firm FE	Yes	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes	Yes
Controls (Firm & State)	Yes	Yes	Yes	Yes
Sample (pre-policy defined)	All firms	All firms	All firms	All firms
Observations	2,312	2,312	2,312	2,312
Adjusted within R^2	0.038	0.618	0.689	0.629

Panel B. Redirecting Toward Renewables

Dep. Var.: $\text{PortfolioShare}^{(k)}_{i,t}$				
$k =$	*Renewables*	*Gas*	*Renewables*	*Gas*
	(1)	(2)	(3)	(4)
$\text{InRPS}_i \times \text{PostPolicy}^{\text{Air}}_{i,t}$	3.241**	-2.969	2.418**	-3.887*
	(1.446)	(2.052)	(1.138)	(2.249)
Firm FE	Yes	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes	Yes
Controls (Firm & State)	Yes	Yes	Yes	Yes
Sample (pre-policy defined)	All firms	All firms	Coal-based	Coal-based
Observations	2,312	2,312	1,897	1,897
Adjusted within R^2	0.111	0.127	0.040	0.067

local electricity market size), $ElectricPrice$ (cents/kWh, capturing variation in local electricity pricing conditions), and $CarveOut$ (ratio, measuring binding renewable mandates that target specific technologies such as rooftop solar). Firm fixed effects (θ_i) control for unobserved time-invariant differences across firms. Event-time fixed effects (λ_τ) account for relative-year trends around initial regulatory exposure, where $\tau = t - g_i$ denotes time since firm i first became subject to the air pollution regulations, with cohort g as the year of initial regulatory exposure.

Table 3 Panel A presents the results of the first test, which examines whether the CSAPR and MATS lead coal-reliant firms to exit pollutive capacity and reallocate elsewhere. Firms without coal-based generation provide a natural benchmark, because they are not exposed to the primary emission sources (e.g., sulfur dioxide and mercury) that the CSAPR and MATS are designed to reduce. Coal-reliant firms face direct compliance pressure, as coal is the principal target of the regulations. Using a continuous measure allows the specification to reflect differences in exposure intensity. Because generation capacity shares must sum to one, I also include the lagged value of $PortfolioShare^{(k)}_{i,t-1}$ (for $k \neq$ coal) to purge the deterministic negative correlation between $CoalReliance_i$ and pre-policy k-share.

Column (1) shows that, relative to firms without any coal exposure, a firm entirely reliant on coal prior to the regulation reduces its coal share by 9.5 percentage points following the adoption of CSAPR and MATS. This decline aligns with the regulatory mechanism: both policies impose binding emission limits on pollutants predominantly emitted by coal-fired units. Column (2) shows that the same fully coal-reliant firm increases its gas share by 3.0 percentage points over the same period, suggesting that natural gas served as the primary substitute. Columns (3) and (4) report no statistically significant changes in the shares of renewables or other clean sources. These patterns indicate that the emission regulations prompted a sizable exit from coal among the most exposed firms, but that short-run capacity redirection was concentrated in other less pollutive fossil sources rather than toward cleaner technologies. The absence of a significant renewable response is consistent with Blonz et al. (2026), who find that MATS reduces solar investment by tightening firms' financing constraints.

While Panel A highlights the role of emission regulations in pushing firms to exit coal-based generation, Panel B examines whether pre-existing renewable portfolio standards (RPS) influence where firms redirect capacity once that exit becomes necessary. $InRPS_i$ is defined as one if the firm was located in a state with an RPS adopted before the CSAPR and MATS took effect.[16] Although RPS mandates typically impose long-term renewable targets, they may shape short-run reallocation by making renewables a more salient or institutionally supported option. Emission rules restrict the status quo by raising the cost of pollutive capacity, whereas RPS policies act as forward-oriented signals that guide investment toward a specific substitute. The key question is whether firms exposed to both regulatory pressures respond differently than those subject only to the air pollution regulatory push.

Panel B column (1) shows that, relative to firms in non-RPS states, those exposed to a renewable mandate increase their renewable generation share by 3.2 percentage points following the adoption of CSAPR and MATS. Column (2) shows no significant difference in the change in gas portfolio share between the two groups, implying that gas remains the average reallocation path, as observed in Panel A. Columns (3) and (4) restrict the sample to firms with coal-based generation prior to the regulation: those for whom the transition is most directly binding. Among these firms, those exposed to RPS policies also increase their renewable share by 2.4 percentage points. Column (4) shows a 3.9 percentage point reduction in gas share for RPS group, which is marginally significant. This pattern reinforces the interpretation that RPS mandates affect where firms reallocate once a response is required.

[16] In the year prior to the CSAPR and MATS, 56% of firms operated in only one state, and 94% operated in three or fewer states. For firms with multi-state operations, $InRPS_i$ is set to one if any of those states had adopted an RPS one year prior to the emission regulations.

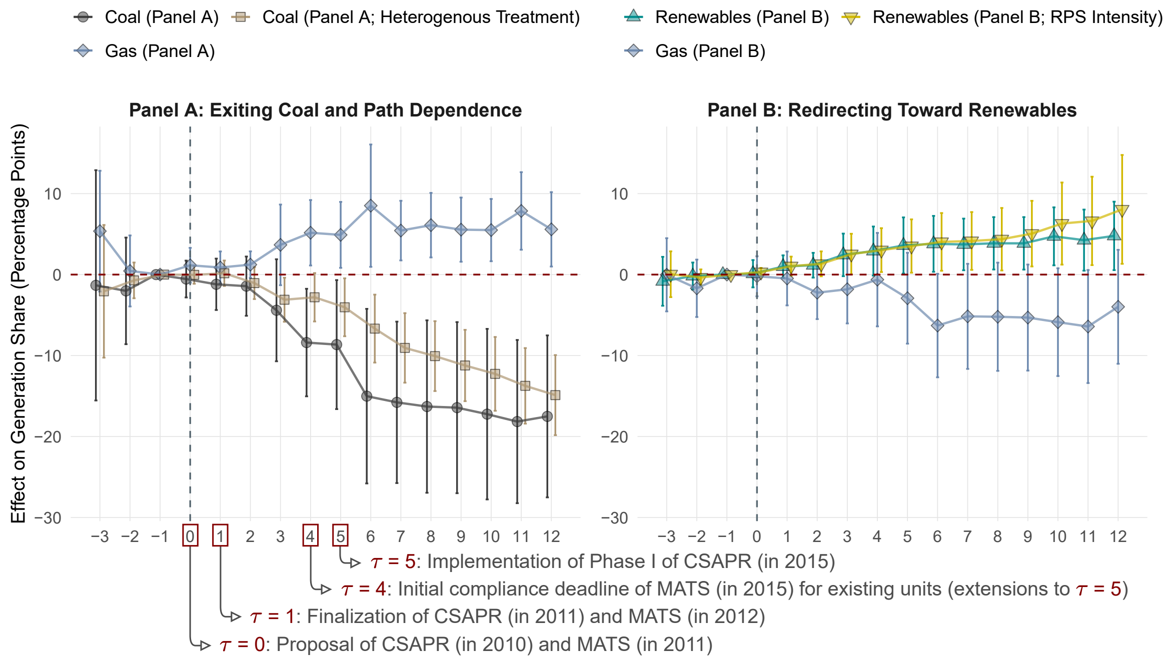

Figure 2. Regulatory Pressure and Dynamic Shifts in Generation Portfolios. This figure presents coefficient estimates from the dynamic difference-in-differences specification that traces changes in electricity generation portfolios around the introduction of CSAPR and MATS. The dependent variable is $PortfolioShare^{(k)}_{i,t}$, defined as the share of generation type k (coal, gas, or renewables) in firm i's total owned capacity in year t, weighted by ownership share. Panel A examines the treatment effects on coal and gas generation shares for firms with different reliance on coal generation, measured by the pre-policy share of coal capacity ($CoalReliance_i$). For robustness, a heterogeneous treatment effect estimator is also plotted for the effect on coal generation shares. Panel B presents two specifications. The first compares firms operating in states with versus without a pre-policy RPS, captured by the indicator $InRPS_i$. The second uses a time-variant RPS exposure measure based on the maximum state-level RPS target a firm is exposed to in each year. Both the heterogeneous treatment effect and the time-variant RPS intensity specifications are estimated using the method of de Chaisemartin and D'Haultfoeuille (2024). Event time ranges from $\tau = -3$ to $\tau = 12$, with $\tau = -1$ as the reference year. In the main specification, it refers to the year prior to the air pollution regulations. In the RPS intensity specification, it marks the last year before a firm is exposed to both a non-zero RPS target and the air pollution regulation adoption. Vertical bars represent 95% confidence intervals based on standard errors clustered at the firm level.

Overall, the results highlight a two-part structure of transition. Emission regulations create the need to exit pollutive capacity, and renewable mandates help define the reinvestment margin by shaping where firms go next. To examine the timing of the transition, I estimate:

$$\begin{aligned}\text{PortfolioShare}^{(k)}_{i,t} = \alpha_0 + &\sum_{\tau=-3,\,\tau\neq-1}^{12} \phi_\tau \left(\text{Exposed}_i \times \mathbb{1}\{\text{Period}_{i,t} = \tau\}\right) \\ &+ X'_{i,t-1}\gamma + Z'_{s(i),t-1}\delta + \theta_i + \lambda_\tau + \varepsilon_{i,t}.\end{aligned} \tag{3}$$

The specification mirrors the setup in Table 3, but replaces the single post-policy interaction with a full set of event-time interactions. Each ϕ_τ captures the differential change in portfolio share between treated and control firms in year τ, relative to the year prior to regulation. I estimate effects from $\tau = -3$ to $\tau = 12$, omitting $\tau = -1$ as the reference period. The lower bound of $\tau = -3$ reflects data constraints: consistent electronic demand data from FERC Form 714 begins in 2006, and lagged controls restrict estimation to 2007 and later. Because

RPS targets evolve annually across states, I complement the current approach with a separate specification that uses the maximum state-level RPS target a firm is exposed to in each year as a continuous measure of exposure. I estimate its dynamic treatment effect using the method of de Chaisemartin and D'Haultfoeuille (2024), which captures variation in policy intensity over time and assesses whether higher targets lead to a sharper shift toward renewables.

Figure 2 plots the estimated ϕ_τ coefficients, with event time (τ) 0 marks the introduction of CSAPR and MATS, and vertical bars denote 95% confidence intervals. Panel A compares the differential evolution of coal and gas portfolio shares between coal-reliant and non-coal firms. The decline in coal share becomes statistically significant in the fourth year after the regulations, coinciding with a significant increase in gas generation in the same year. The timing aligns with the adoption-to-compliance windows of CSAPR and MATS.

Panel B compares firms in RPS versus non-RPS states. The increase in renewable share becomes statistically significant in the fifth year after the regulations, while no significant difference is observed in gas generation throughout the period. Estimates from the RPS intensity specification show a consistent pattern, with stronger shifts toward renewables under higher policy targets after the air pollution regulations. The absence of differential pre-trends in both panels lends support to the identification strategy. In summary, the dynamic patterns reinforce the two-part structure of the transition: coal-reliant firms begin reducing pollutive capacity after the air pollution regulations, while those under renewable mandates increasingly redirect capacity toward renewables.

4.2 Exit Channel: Divestiture vs. Closure

Given that the federal air pollution regulations induced the exit of coal power generation, I move on to first examine whether the affected firms respond through asset divestiture or operational shutdown, as these choices represent distinct exit strategies. Divestiture reflects a financial adjustment that reallocates ownership while keeping the physical capacity in operation, whereas closure marks a more fundamental retreat from coal-based production, signaling a structural step in the broader transition process.

To clarify the tradeoff a coal asset owner faces among (i) continued operation, (ii) divestiture, and (iii) retirement, Internet Appendix Figure B.3 illustrates the equilibrium allocation of coal assets in the model of Darmouni and Zhang (2024). The horizontal axis is a coal unit's operating profitability π^c, and the vertical axis is the equilibrium coal asset price P^c. The owner chooses among "*Keep*", "*Sell*", and "*Retire*", with payoffs $\pi^c - \epsilon$, $P^c - \theta\epsilon$, and 0, respectively. Here, ϵ is the environmental externality cost faced by owners, and mandate stringency $\theta \in [0, 1]$ determines how much of this cost is reflected when selling. The figure contrasts two buyer environments that determine the equilibrium pricing line. In Panel A, buyers face no financial frictions, thus $P^c = \pi^c$. In Panel B, buyers are financially constrained, and the equilibrium pricing line is $P^c = (\xi/r)\pi^c$ with $\xi/r < 1$.[17]

[17] Following Darmouni and Zhang (2024), the buyer has no cash on hand and must finance the purchase externally. Only a fraction $\xi \leq 1$ of future operating profits π^c is pledgeable to financiers, who require an expected gross return $r \geq 1$. The purchase is feasible only if $\xi\pi^c \geq rP^c$, which implies $P^c \leq (\xi/r)\pi^c$.

Given the portfolio path dependence documented in Section 4.1, I extend the baseline environment in Figure B.3 by allowing buyers to allocate capital to gas-fired generation as an alternative to acquiring coal assets. This alternative is represented by $V^g \geq 0$, where V^g is a function of Δ that captures coal's marginal fuel-cost disadvantage relative to gas. This channel lowers the owner's divestiture payoff to $P^c - \theta\epsilon - V^g(\Delta)$. In other words, in locations where gas-fired generation is a better alternative for potential buyers (higher $V^g(\Delta)$), buyers require a larger discount to acquire coal assets, reducing the value of divestiture at a given price P^c. Accordingly, the sell-versus-retire cutoff shifts from $P^c = \theta\epsilon$ to $P^c = \theta\epsilon + V^g(\Delta)$.

The model yields two empirical implications. First, divestiture and continued operation are always feasible choices, and retirement becomes an equilibrium option only when regulatory pressure is present (e.g., for a given $\theta\epsilon$ from CSAPR and MATS). Second, a larger coal fuel-cost disadvantage relative to gas (higher $V^g(\Delta)$) lowers the sell payoff, reducing divestiture and increasing retirement in equilibrium. Particularly in the financially constrained scenario, such reduction in divestiture is larger: units that no longer divest shift into both retirement and continued operation in equilibrium. Overall, in both buyer environments, a larger coal fuel-cost disadvantage increases the propensity for the owner to retire rather than sell coal units.

Panel A of Table 4 provides an overview of how often each type of exit strategies occurs following the adoption of federal air pollution regulations. The panel estimates a multinomial logit model where the outcome variable $ExitType_{i,t}^{(m)}$ classifies firm i's action in year t as no exit ($m = 0$), divestiture without closure ($m = 1$), or closure of coal-fired generators ($m = 2$). Divestiture is defined as a reduction in ownership share of coal assets without a corresponding EIA retirement flag, while closure is marked by full generator shutdown. In cases where both occur in the same year, the observation is coded as closure. The estimated marginal effects indicate that regulation raises the likelihood of divestiture by 3.34 percentage points and closure by 3.85 percentage points, relative to continued operation.

Panel B examines whether firms' exit strategies respond to post-regulation differences in fuel costs. The analysis tests two things: (i) a substitution effect, i.e., whether a relative increase in coal prices makes gas generation a more attractive alternative for these path-dependent firms and thereby accelerates coal exit, and (ii) whether such effect is more likely to take the form of plant closures rather than divestitures, given that generation assets are inherently immobile and the fuel price differences apply to all owners and buyers operating within the same jurisdiction. I estimate a DiD specification that interacts regulatory exposure with $FuelCostGap_i^{Coal\text{-}Gas}$, a firm-level measure capturing the standardized cost disadvantage of coal relative to gas. The original gap is measured in USD per megawatt-hour (MWh), and is computed by subtracting the average gas costs across all firms with gas generation in the same year from each firm's coal costs, measured in the year following the air pollution regulations. Because the sample includes only coal-reliant firms, each firm's coal costs are directly observed. The interaction of $PostPolicy_{i,t}^{Air}$ and $FuelCostGap_i^{Coal\text{-}Gas}$ identifies how this relative price signal shapes firms' response to regulation. Dependent variables include an indicator equal to one if the firm retires at least one coal-fired generator in a given year, the log of one plus abatement costs ($m), and

The ratio ξ/r therefore represents how loose financial frictions are: a lower ξ/r implies tighter constraints and an equilibrium price below the unconstrained benchmark $P^c = \pi^c$.

Table 4: Firm Exit Channels Following Air Pollution Regulation

This table examines how coal-reliant electric firms respond to the federal air pollution regulations through divestiture, closure, pollution abatement, and portfolio adjustment. In Panel A, the dependent variable $ExitType_{i,t}^{(m)}$ is a multinomial indicator where $m = 1$ denotes divestiture without closure and $m = 2$ denotes closure (with or without divestiture). $PostPolicy_{i,t}^{Air}$ equals one if firm i's state has adopted the air pollution regulations (CSAPR or MATS) by year t. In Panel B, the dependent variables are closure and divestiture indicators for coal-fired generators (each coded as 100 if the event occurs) and the log of one plus abatement cost. Panel C reports generation portfolio outcomes: the portfolio shares and levels of gas and coal capacity. $FuelCostGap_i^{Coal-Gas}$ denotes the standardized difference between coal and gas fuel costs measured in the year following regulation proposal. All regressions include firm fixed effects (θ_i) and event-time fixed effects (λ_τ), where $\tau = t - g_i$ denotes the number of years since firm i first became subject to the regulations, with cohort g defined as the year of initial exposure. Firm-level controls include lagged values of *FirmSize*, *Leverage*, *ROA*, *Liquidity*, *Depreciation*, *Capacity*, and *FuelCost*. State-level controls are averaged at the firm level and include lagged *StateDemand*, *ElectricPrice*, and *CarveOut*. All variable definitions are in the Appendix. Standard errors are clustered at the firm level. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Panel A. Firm Choice Between Divestiture and Closure

Dep. Var.: $\text{ExitType}_{i,t}^{(m)}$				
$m =$	*1.Divestiture Only*	*2.Closure*		
	(1)	(2)		
$\text{PostPolicy}_{i,t}^{\text{Air}}$	1.333***	1.364***	Benchmark (m)	*0.No Exit*
	(0.428)	(0.388)	Observations	1,897
Marginal Effect	3.34%***	3.85%***	Pseudo-R^2	0.141
	(0.011)	(0.012)		
Controls (Firm & State)	Yes	Yes		
Sample (pre-policy defined)	Coal-based	Coal-based		

Panel B. Fuel Relative Prices and Exit Outcomes

Dep. Var.:	$\text{Closure}_{i,t}$	$\text{Divestiture}_{i,t}$	$\text{Abatement}_{i,t}^{log}$
	(1)	(2)	(3)
$\text{PostPolicy}_{i,t}^{\text{Air}} \times \text{FuelCostGap}_i^{\text{Coal-Gas}}$	5.468***	1.368	1.922***
	(2.038)	(1.740)	(0.619)
Firm FE	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes
Controls (Firm & State)	Yes	Yes	Yes
Sample (pre-policy defined)	Coal-based	Coal-based	Coal-based
Observations	1,820	1,820	1,820
Adjusted within R^2	0.011	0.003	0.190

Panel C. Substitution Effect: Fuel Relative Prices and Generation Portfolio

Dep. Var.:	$\text{PortfolioShare}_{i,t}^{(k)}$ (%)		$\text{Capacity}_{i,t}^{(k)}$ (GW)	
	$k = Gas$	$k = Coal$	$k = Gas$	$k = Coal$
	(1)	(2)	(3)	(4)
$\text{PostPolicy}_{i,t}^{\text{Air}} \times \text{FuelCostGap}_i^{\text{Coal-Gas}}$	5.199***	-4.873***	0.314**	-0.248**
	(1.525)	(1.366)	(0.130)	(0.099)
Firm FE	Yes	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes	Yes
Controls (Firm & State)	Yes	Yes	Yes	Yes
Sample (pre-policy defined)	Coal-based	Coal-based	Coal-based	Coal-based
Observations	1,820	1,820	1,820	1,820
Adjusted within R^2	0.157	0.048	0.099	0.136

an indicator equal to one if it divests at least one generator. The indicators are multiplied by 100 so that the coefficients can be interpreted as percentage point effects.

Firms facing a larger post-policy coal cost disadvantage are significantly more likely to close coal-fired generators: those with a one-standard-deviation larger fuel cost gap (around \$21.13 per MWh) are 5.5 percentage points more likely to close (column 1). In contrast, the effect on divestiture (column 2) is positive but not statistically significant, indicating that higher coal costs relative to gas primarily lead firms to close plants rather than liquidate ownership. This pattern is consistent with the model's prediction that the relative price pressures reduce the scope for asset sales by limiting the pool of potential buyers facing the same cost conditions. Column (3) shows a similarly strong response for abatement spending, consistent with firms investing in the coal capacity they continue to operate when sale is not a practical alternative.

Panel C shows the corresponding portfolio adjustment. Columns (1) and (2) indicate that firms with a one-standard-deviation larger gap shift 5.2 percentage points more of their capacity into gas and 4.9 percentage points more out of coal, a near one-for-one reallocation. Columns (3) and (4) confirm that this is a reallocation of capacity rather than a compositional artifact: the firms' gas capacity is 0.314 GW higher and their coal capacity 0.248 GW lower.

In the current specification, measuring the fuel cost gap in the year following the policy proposal captures the cost environment in which a firm plans its transition response. This timing also leaves an ambiguity, since the regulations may move the cost gap at the same time. Internet Appendix Section B.1 re-estimates both panels with the cost gap measured from the year before the proposal through three years after ($\tau \in \{-1, 0, +1, +2, +3\}$). The estimates remain statistically significant at every event time in the window, including at $\tau = -1$, where the gap is measured before the regulations were proposed. This suggests that the substitution effect mainly captures the fuel cost heterogeneity that already existed across firms.

Finally, to assess whether the coal exits also deliver measurable emission reductions, Internet Appendix Section B.2 presents two results. First, coal-reliant firms significantly reduce emissions of SO_2 and NO_x, the policy-targeted pollutants, and CO_2 falls alongside them as an indirect consequence of the post-policy transition. Second, the coal-to-gas substitution captured by the fuel cost gap accounts disproportionately for the targeted pollutants: the aggregate post-policy reductions in SO_2 and NO_x correspond to 1.1 and 1.7 standard deviations of the fuel cost gap, whereas the aggregate CO_2 reduction corresponds to roughly 3.0 standard deviations. This implies that even though gas is a less pollutive fossil fuel, the coal-to-gas substitution channel alone cannot explain the entire post-policy carbon reduction, with the remaining carbon reduction potentially from additional abatement or a broader shift toward non-fossil generation.

4.3 Transition Channel: RPS-Induced Construction and Renewable Capacity

Having examined firms' coal exit strategies and compliance outcomes, I next turn to another dimension of the transition process: how firms redirect their generation portfolios after moving away from coal generation. Earlier evidence from Table 3 Panel B shows that firms' exposure to state Renewable Portfolio Standards (RPS) may signal a green transition through policy-driven incentives for renewable investment. To identify this investment channel, I implement a fuzzy regression discontinuity design (FRDD) exploiting state-level RPS annual targets. RPS

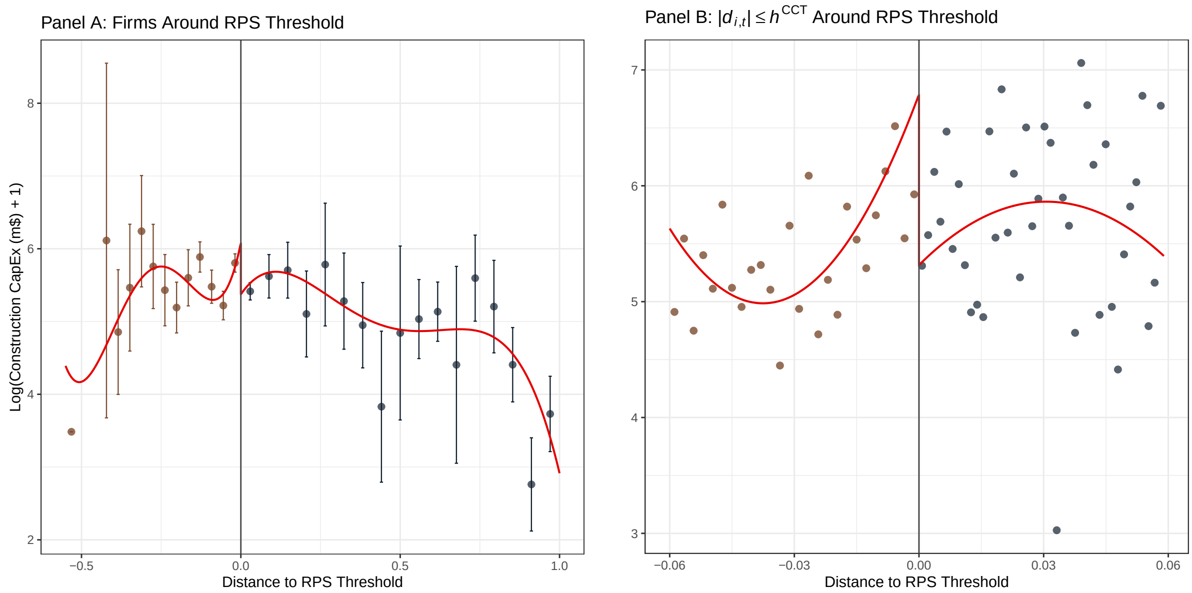

Figure 3. Discontinuity in Construction CapEx. This figure presents regression discontinuity plots of firm-level construction capital expenditures ($ConstructCapEx$) against the distance to the state-level Renewable Portfolio Standard (RPS) threshold, defined as the difference between a firm's renewable generation capacity share and the annual RPS target. The sample is restricted to post–federal air pollution regulation years to ensure that firms face clear energy transition incentives. Panel A uses the entire observed range of the running variable and estimates a fourth-order polynomial fit on either side of the threshold. Panel B restricts the sample to firms within the optimal Calonico–Cattaneo–Titiunik (CCT) bandwidth, defined as $|d_{i,t}| \leq h^{\text{CCT}}$ with $h^{\text{CCT}} = 0.06$, selected based on forward three-year renewable portfolio share estimates. Within this narrow window, a second-order polynomial is estimated on each side using a triangular kernel. Dots indicate binned means of the outcome variable; 95% pointwise confidence intervals in Panel A reflect within-bin variation.

requirements set annual minimum targets for renewable generation, with thresholds increasing over time. This institutional feature creates discontinuous incentives for firms to initiate or accelerate renewable investment projects when their renewable share falls just short of the threshold.

Two considerations motivate the identification strategy. First, the fuzzy nature of the group assignment. Firms that do not meet the RPS generation target are not necessarily those that fail to comply. For instance, compliance can also occur through the purchase of Renewable Energy Credits (RECs) from Independent Power Producers (IPPs), allowing firms to meet RPS obligations without increasing their own renewable generation. In addition, for firms operating in multiple states, I measure exposure using the highest RPS target among the states in which they operate. Because proximity to the threshold (RPS target) affects the likelihood but not the certainty of treatment, the design is characterized as fuzzy rather than sharp.

Second, the choice of the assignment variable. I use installed renewable capacity share rather than realized generation to define a firm's position relative to the RPS threshold. RPS targets change annually, thus past generation is not mechanically linked to the current compliance status, and cannot produce a sharp discontinuity either. Alternatively, using same-year generation against the RPS threshold may bring potential look-ahead bias, since renewable capital expenditure can be made before production outcomes are realized and before the firm's compliance position is known. In comparison, the installed capacity is determined prior to the actual generation, and reflects a firm's potential to meet RPS requirements in that year, capturing how the compliance pressure affects investment decisions before the actual production and compliance status are observed. Together, these features create a quasi-experimental environment around the RPS annual targets. Firms just below the threshold have more incentives to make renewable investment, allowing identification of the causal effect of RPS-induced renewable transition.

Figure 3 visualizes the relationship between construction-related capital expenditures and the distance to the RPS threshold, defined as the difference between a firm's renewable generation capacity share and the binding annual RPS target. Panel A includes the full post-policy sample, and estimates separate fourth-order polynomials on either side of the threshold. Panel B narrows the window to firms with threshold distances $|d_{i,t}| \leq 0.06$, selected using the optimal Calonico–Cattaneo–Titiunik (CCT) bandwidth (Calonico et al., 2014) for estimating three-year forward renewable generation portfolio shares, and applies a second-order polynomial fit on each side. In both panels, the plotted dots represent binned averages of *ConstructCapEx*, the natural logarithm of one plus construction-related capital expenditures. Panel A shows binned means also with 95% pointwise confidence intervals based on within-bin variation. The discontinuous increase in expenditures just below the threshold provides visual evidence of firms' investment responses to compliance incentives.

The estimation follows a FRDD framework implemented via two-stage least squares (2SLS), using the symmetric window of ±6% around the RPS threshold based on the optimal CCT bandwidth. The sample covers firm-year observations after the adoption of the air pollution regulations, focusing the renewable investment behavior during the redirection period following the coal generation exit. In the first stage, I instrument construction-related capital expenditures using an indicator for whether a firm's renewable generation capacity share falls below its state's binding RPS annual target:

$$\text{ConstructCapEx}_{i,t} = \alpha_0 + \alpha_1 \text{BelowRPS}_{i,t} + f(d_{i,t}) + \theta_{f(i)} + \lambda_t + \varepsilon_{i,t}. \tag{4}$$

Here, $ConstructCapEx_{i,t}$ denotes the natural logarithm of one plus construction-related capital expenditures for firm i in year t. The indicator $BelowRPS_{i,t}$ equals one if the firm's renewable share falls below the RPS target in year t. The running variable $d_{i,t}$ measures the distance between a firm's renewable share and the RPS threshold. To control for smooth trends around the threshold, I include a polynomial in $d_{i,t}$, defined as $f(d_{i,t}) = \sum_{k=1}^{K} \gamma_k d_{i,t}^k$, with the baseline specification using $K = 2$ as recommended by Gelman and Imbens (2019). Both stages of the 2SLS also include firm fuel-type fixed effects ($\theta_{f(i)}$) and year fixed effects (λ_t). The fuel type corresponds to the generation technology that accounts for the largest share of a firm's installed capacity, i.e., technology k such that $PortfolioShare_{i,t}^k$ is the highest. This fixed effect absorbs persistent operational differences across firms dominated by different generation technologies, ensuring that the estimated effects reflect variation in firms' responses to RPS incentives rather than structural differences inherent to each technology type.

Table 5: RPS-Induced Construction and Renewable Transition

This table reports the FRRD estimates of the effect of renewable policy-induced construction capital expenditures (*ConstructCapEx*) on firms' generation portfolio composition. The endogenous variable is instrumented using an indicator for whether the firm falls below its state's binding annual RPS target (*BelowRPS*), along with second-order polynomial controls in the firm's distance to the threshold (degree $K = 2$, including d and d^2). The sample covers the period after the air pollution regulations. Column (1) reports estimates for all firms, using renewable generation portfolio share in year t+1 as the outcome. Column (2) applies the same specification but replaces the outcome with the average renewable portfolio share over years t+1 to t+3. Column (3) repeats the specification in column (2) using only coal-based firms (defined before CSARP and MATS), which are the primary targets of the air pollution regulations. Columns (4) and (5) present falsification tests using all firms, with the outcomes as the average gas and other clean generation portfolio shares over years t+1 to t+3. All regressions include firm fuel-type fixed effects and year fixed effects. All variable definitions are in the Appendix. Standard errors are clustered at the firm level. The Kleibergen–Paap Wald F-statistic assesses instrument strength under heteroskedasticity, with values above 10 commonly used as a benchmark for single-instrument relevance. For reference, the Cragg–Donald Wald F-statistic and Stock–Yogo critical values are also reported (at the 10% and 15% maximal IV size thresholds), though these rely on the assumption of homoskedastic and i.i.d. errors. The Anderson–Rubin Wald test evaluates whether the instrumented regressor has any effect under valid instruments and remains robust to weak identification.

Dep. Var.: $\text{PortfolioShare}^{(k)}_{i,forward}$					
$k =$	*Renewables*			*Gas*	*Other Clean*
$forward =$	$t+1$	$[t+1, t+3]$	$[t+1, t+3]$	$[t+1, t+3]$	$[t+1, t+3]$
	(1)	(2)	(3)	(4)	(5)
Second Stage					
$\widehat{\text{ConstructCapEx}}_{i,t}$	5.270***	5.766***	5.407**	-7.368	1.202
	(2.001)	(2.186)	(2.208)	(4.495)	(0.927)
First Stage (Instrumenting $ConstructCapEx_{i,t}$)					
$\text{BelowRPS}_{i,t}$	1.015***	1.015***	0.995***	1.015***	1.015***
	(0.288)	(0.288)	(0.288)	(0.288)	(0.288)
Specification Category	Main	Main	Main	Placebo	Placebo
Firm Fuel-Type FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Polynomial Control in d	$K = 2$	$K = 2$	$K = 2$	$K = 2$	$K = 2$
Sample (pre-policy defined)	All firms	All firms	Coal-based	All firms	All firms
Observations	1,277	1,277	1,064	1,277	1,277
Weak Instrument Test					
Kleibergen-Paap Wald F-stat	12.42	12.42	11.91	12.42	12.42
Cragg-Donald Wald F-stat	70.99	70.99	71.73	70.99	70.99
Stock-Yogo critical value (10%)	16.38	16.38	16.38	16.38	16.38
Stock-Yogo critical value (15%)	8.96	8.96	8.96	8.96	8.96
Test of Endogeneity					
Anderson-Rubin Wald Chi-sq	16.19	16.44	12.27	3.08	1.67
p-value	0.00	0.00	0.00	0.08	0.20

In the second stage, I estimate the effect of instrumented investment on renewable generation capacity shares over the subsequent years:

$$\text{PortfolioShare}^{\text{Renewables}}_{i,forward} = \beta_0 + \beta_1 \widehat{\text{ConstructCapEx}}_{i,t} + f(d_{i,t}) + \theta_{f(i)} + \lambda_t + \varepsilon_{i,t}. \tag{5}$$

Here, $PortfolioShare^{Renewables}_{i,forward}$ represents the average share of renewable generation in firm i's portfolio in the following years. I consider two windows in the analysis: the immediate next year (t+1) and a medium-range window averaging years t+1 through t+3.

Table 5 column (1) indicates that for firms close to the RPS target, a 10% increase in instrumented construction-related capital expenditure raises the generation portfolio share of renewables in the following year by roughly 0.53 percentage points (5.270×0.10). Similarly, the same 10 % spending rise lifts the average renewable share over the horizon $t+1$ to $t+3$ by roughly 0.58 and 0.54 percentage points for all firms (column 2) and among pre-policy coal-reliant firms (column 3), respectively. Placebo regressions for gas and other clean technologies (columns 4 and 5) show no comparable effect, reinforcing the channel's specificity on renewable transition.

The first-stage estimates confirm the validity of the instrument. The indicator for firms below the RPS threshold strongly predicts construction-related capital expenditures, with coefficients positively significant at the 1% level in every specification. Kleibergen–Paap Wald F-statistics exceed the conventional benchmark of 10. In addition, the Anderson-Rubin Wald tests reject the null of no effect of the instrumented regressor on renewables share in columns (1)–(3), supporting robustness of the causal interpretation. In contrast, placebo tests using gas and other clean technologies (columns 4 and 5) yield non-significant Anderson–Rubin statistics at the 5% level, suggesting that the regulatory discontinuity primarily affects renewable energy investments.

Overall, these findings indicate that RPS-induced increases in construction-related capital expenditures accelerate the transition of generation portfolios toward renewable energy. The estimates provide direct evidence that firms respond to renewable regulations through capital investment, and highlight how targeted generation requirements shape the direction of that investment under broader transition pressures from the air pollution regulations.

4.4 Transition Targeting: Financial Incentives and Technology Choice

Given that firms expand renewable investment to meet RPS targets, a natural follow-up question is which specific generation type they choose for the transition. By design, regulatory mandates such as RPS often set broad requirements because the main goal is to direct firms away from non-renewable toward renewable generation, while leaving the choice of technology largely flexible.[18] Financial incentives, in contrast, operate more specifically by lowering the cost of particular technologies that related to power generation. In this sense, regulation establishes the boundary of acceptable energy sources, while incentives are designed to support specific generation types within that boundary.

Therefore, to investigate how the transition is targeted across technologies, I categorize state-level financial incentive programs into three categories: (i) those promoting renewables of any type, (ii) solar-specific, and (iii) wind-specific technologies, to test whether firms in these states shift their portfolios toward the corresponding targets. The model specification builds on equation (1), where the dependent variable $PortfolioShare_{i,t}^{(k)}$ denotes the share of generation type k in firm i's owned capacity in year t. The treated group indicator $Incentive_{i}^{(r)}$ equals one if, before CSAPR or MATS, the firm operated in a state offering financial incentives for

[18]For instance, only a few RPS programs contain technology-specific carveouts like solar set-asides, which require a minimum share from designated technologies. See, e.g., Lawrence Berkeley National Laboratory, *U.S. State Renewables Portfolio & Clean Electricity Standards: 2024 Status Update*, https://eta-publications.lbl.gov/sites/default/files/lbnl_rps_ces_status_report_2024_edition.pdf.

Table 6: Renewable Financial Incentives and the Direction of Transition

This table examines how targeted financial incentives affect firms' generation portfolios following the adoption of federal air pollution regulations. The dependent variable $PortfolioShare_{i,t}^{(k)}$ denotes the share of generation type k in firm i's owned capacity in year t. The key variable $Incentive_i^{(r)}$ is a firm-level indicator equal to one if before the regulations the firm operated in a state offering financial incentives targeting technology r, including (i) any renewable, (ii) solar-specific, and (iii) wind-specific technologies. $PostPolicy_{i,t}^{Air}$ equals one if firm i's state has adopted the air pollution regulations (CSAPR or MATS) by year t. The dependent variables vary by column and include portfolio shares of renewables, solar, wind, gas, and other clean sources. Panel A reports estimates for firms that relied on coal prior to the air pollution regulations (labeled as "coal-based" firms). Panel B restricts to coal-reliant firms whose generation portfolio shares in the targeted technology r were stable or zero three years prior to the air pollution regulations (labeled as "coal-static" firms). Columns (4) and (5) in each panel serve as falsification tests for gas and other clean generation. All regressions include firm fixed effects (θ_i) and event-time fixed effects (λ_τ), where $\tau = t - g_i$ denotes the number of years since firm i first became subject to the air pollution regulations, with cohort g defined as the year of initial exposure. Firm-level controls include lagged values of *FirmSize*, *Leverage*, *ROA*, *Liquidity*, *Depreciation*, *Capacity*, and *FuelCost*. State-level controls are averaged at the firm level and include lagged *StateDemand*, *ElectricPrice*, and *CarveOut*. All variable definitions are in the Appendix. Standard errors are clustered at the firm level. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Panel A. Targeted Incentives and Transition Response (Coal-based Firms)

Dep. Var.: $\text{PortfolioShare}_{i,t}^{(k)}$					
$k =$	*Renewables*	*Solar*	*Wind*	*Gas*	*Other Clean*
Incentive technology (r=)	*Renewables*	*Solar*	*Wind*	*Renewables*	*Renewables*
	(1)	(2)	(3)	(4)	(5)
$\text{Incentive}_i^{(r)} \times \text{PostPolicy}_{i,t}^{\text{Air}}$	2.488***	0.239**	1.848***	-5.371	0.257
	(0.827)	(0.106)	(0.424)	(3.379)	(0.397)
Specification Category	Main	Main	Main	Placebo	Placebo
Firm FE	Yes	Yes	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes	Yes	Yes
Controls (Firm & State)	Yes	Yes	Yes	Yes	Yes
Sample (pre-policy defined)	Coal-based	Coal-based	Coal-based	Coal-based	Coal-based
Observations	1,897	1,897	1,897	1,897	1,897
Adjusted within R^2	0.061	0.010	0.050	0.125	0.044

Panel B. Targeted Incentives and Transition Response (Coal-static Firms)

Dep. Var.: $\text{PortfolioShare}_{i,t}^{(k)}$					
$k =$	*Renewables*	*Solar*	*Wind*	*Gas*	*Other Clean*
Incentive technology (r=)	*Renewables*	*Solar*	*Wind*	*Renewables*	*Renewables*
	(1)	(2)	(3)	(4)	(5)
$\text{Incentive}_i^{(r)} \times \text{PostPolicy}_{i,t}^{\text{Air}}$	2.062*	0.223**	0.936***	-6.187*	0.366
	(1.046)	(0.104)	(0.341)	(3.430)	(0.501)
Specification Category	Main	Main	Main	Placebo	Placebo
Firm FE	Yes	Yes	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes	Yes	Yes
Controls (Firm & State)	Yes	Yes	Yes	Yes	Yes
Sample (pre-policy defined)	Coal-static	Coal-static	Coal-static	Coal-static	Coal-static
Observations	1,508	1,832	1,634	1,508	1,508
Adjusted within R^2	0.036	0.009	0.014	0.140	0.051

technology r, including the three categories (any renewable, solar-specific, and wind-specific). The post-policy indicator $PostPolicy_{i,t}^{Air}$ equals one if the firm is subject to either CSAPR or

MATS by year t. Financial incentives are from the DSIRE database and restricted to state-level programs classified as "financial incentives" and related to solar, wind, or hydro.[19]

Table 6 Panel A reports the results for the sample of firms with coal-fired generation before the air pollution regulations. Column (1) shows that exposure to renewables-related financial incentives is associated with a 2.5 percentage point increase in the renewable capacity share. Columns (2) and (3) decompose this effect: firms with solar-specific incentives increase their solar share by 0.2 points, while those with wind-related incentives expand wind capacity by 1.8 points. Columns (4) and (5) present falsification tests using gas and other clean energy (nuclear and waste) capacity shares as outcomes not expected to respond to the incentives, and neither coefficient is statistically significant.

Panel B further restricts the sample to the coal-reliant firms whose generation shares in the incentivized technology (r) were stable or zero three years prior to the air pollution regulations. This narrower group provides a more conservative test by limiting bias from pre-existing renewable investment momentum. Within this sample, solar and wind capacity shares increase significantly in response to their respective incentives with smaller magnitudes in columns (2) and (3). The effect on total renewables in column (1) is weaker, possibly due to pooled technology types blurring the specific targeted incentive–outcome relationships. These findings suggest that targeted financial incentives shape the specific technology choices firms make when redirecting investment after the coal exit.

To further examine the timing of these incentive-driven adjustments, I estimate dynamic treatment effects following the specification in Panel A of Table 6, with results presented across four models in Figure 4. Panels A and B use pre-policy incentive indicators ($Incentive_i^{(r)}$) to define the treated group. Panel A compares changes in renewable generation (the targeted outcome of renewable incentives), to other clean technologies (nuclear and waste, the untargeted outcome), which serve as a falsification benchmark. Targeted firms gradually increase their renewable generation share starting around year 9 after the air pollution regulations, while other clean generation remains unaffected. Panel B decomposes the incentive target and firm response: firms facing wind-specific incentives expand wind capacity starting in year 5; solar-related adjustments appear more gradually, with effects becoming significant around year 7.

Panels C and D turn to a heterogeneous treatment specification in which event time is defined relative to the first year a firm becomes jointly subject to both a non-zero subsidy and the air pollution regulations ($\tau = 0$). For each year, I define a time-variant continuous treatment variable for subsidy intensity, measured as the maximum number of active direct financial support programs (grants, rebates, and performance-based incentives) across a firm's operating states. The estimation follows the method of de Chaisemartin and D'Haultfoeuille (2024), which accommodates evolving subsidy intensity and staggered timing. Panel C again compares renewable (targeted outcome) and other clean generation (untargeted outcome), confirming that effects remain only in the targeted category. Panel D examines solar and wind separately: both technologies show increasing trends after the treatment, with wind responding earlier and more

[19] For consistency with the generation portfolio share categories in $PortfolioShare_{i,t}^{(k)}$, here I define renewable financial incentive programs as the combined category of solar, wind and hydro technologies, though hydro is not analyzed separately due to its limited siting flexibility.

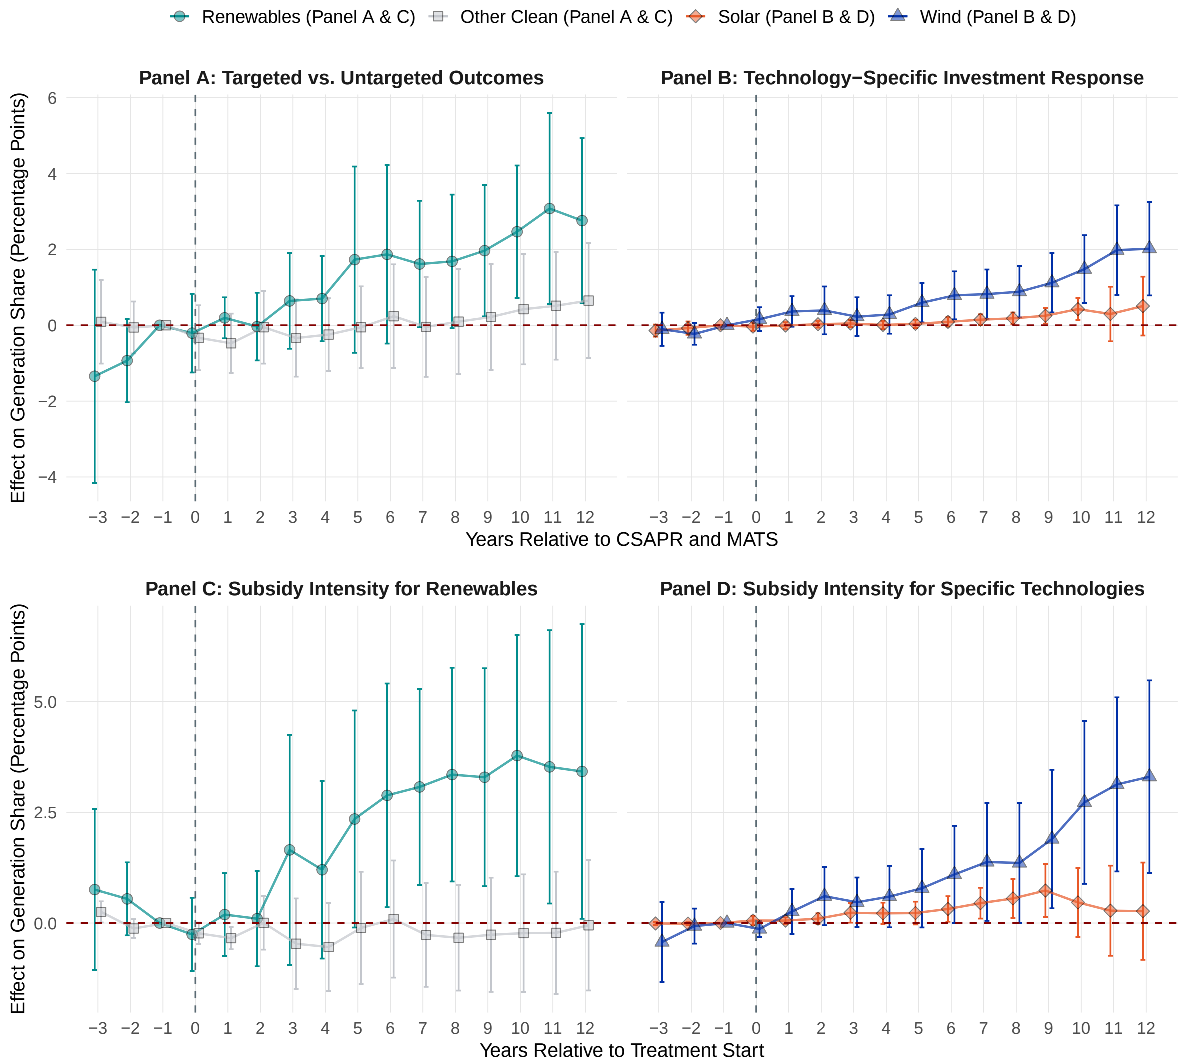

Figure 4. Dynamic Treatment Effects of Financial Incentives on Transition Paths. This figure shows event-study estimates of generation portfolio adjustments following federal air pollution regulations, conditional on their exposure to state-level financial incentives for renewables technologies. The sample covers firms with coal-fired generation before the air pollution regulations. The outcome is $PortfolioShare_{i,t}^{(k)}$, the share of generation type k in firm i's owned capacity in year t. Panels A and B use pre-policy indicators to classify firms based on whether they were exposed to renewable, solar, or wind-related financial incentives prior to the regulations. Panel A compares the impact of renewable incentives on renewables (the targeted outcome), and on other clean sources (nuclear and waste), which serve as a non-targeted falsification outcome. Panel B separately examines how solar and wind generation respond to their respective pre-policy incentives. Panels C and D use time-variant subsidy intensity measures based on the maximum number of technology-specific state-level subsidy programs (direct financial support) a firm is subject to in each year. Estimates are obtained using the method of de Chaisemartin and D'Haultfoeuille (2024). Event time ranges from $\tau = -3$ to $\tau = 12$, with $\tau = -1$ as the reference year. In the specification of Panels A and B, the reference year refers to the year prior to the air pollution regulations. In the specification of Panels C and D, the reference year refers to the last year before a firm is exposed to both a non-zero subsidy and the air pollution regulations. Vertical bars represent 95% confidence intervals based on standard errors clustered at the firm level.

strongly. Together, the results highlight how financial incentives shape not just whether firms transition, but which technologies they adopt and when.

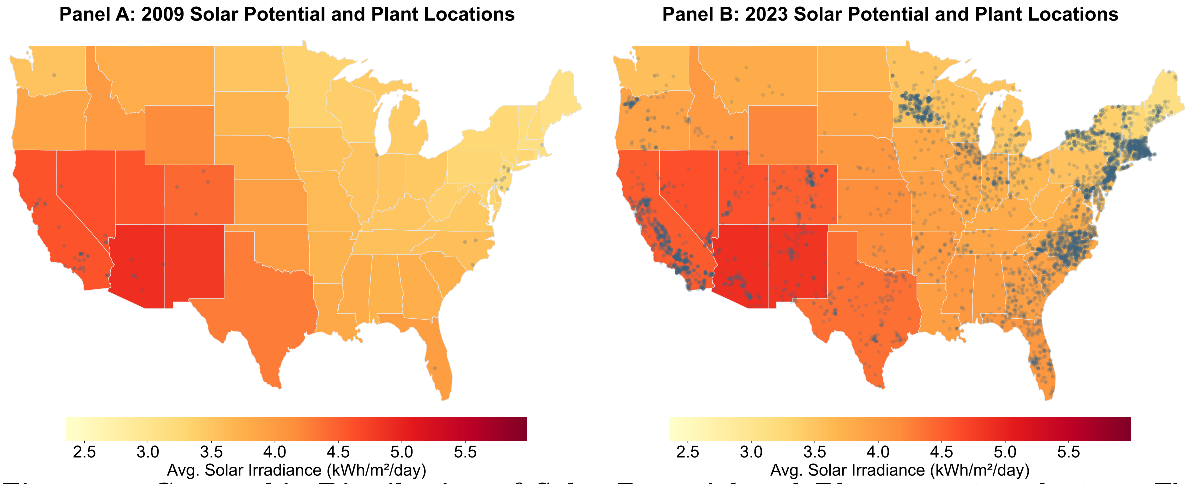

Figure 5. Geographic Distribution of Solar Potential and Plants, 2009 and 2023. This figure presents the distribution of utility-scale solar power plants across the U.S. mainland. Background shading reflects state-level solar potential, measured as average solar irradiance in kilowatt-hours per square meter per day (kWh/m^2/day). Blue dots represent the locations of utility-scale solar power plants in 2009 (Panel A) and 2023 (Panel B).

5 Misaligned Transitions: Geographic Potential vs. Realized Performance

A key distinction of renewable energy technologies is their direct reliance on local geographic conditions, compared to fossil fuels or other clean sources such as nuclear and waste. For example, solar and wind generation depend critically on natural endowments like irradiance and wind speed, which vary substantially across regions. The question is, Section 4 shows that the renewable technology adoption can be shaped by renewable portfolio requirements and financial incentives, but these policies may not always align with local generation potential. From a cross-state perspective, if firms respond to regulations without fully prioritizing geographic comparative advantages in their investment decisions, they may invest in solar or wind in suboptimal but regulated areas, reflecting only local optimization within those areas and weakening the economic returns to the transition. This raises an important question: to what extent do firms adopt renewable technologies that match their location's comparative advantage?

To build intuition, Figure 5 plots the geographic distribution of utility-scale solar PV generation plants in the contiguous United States alongside the underlying state-level average solar irradiance. Panel A visualizes conditions in 2009, before the federal air pollution regulations, while Panel B shows the same map for 2023. The comparison highlights a sharp increase in plant deployment, particularly in eastern and midwestern states with only moderate solar irradiance. While some concentration remains in the high-irradiance southwest, the spatial spread of plant locations suggests that transition activity is not solely determined by geographic potential.

Building on these considerations, this section first develops a static choice model to illustrate the firm's investment tradeoffs among the key external conditions during the post-coal-exit transition stage to either less pollutive fossil fuels or renewables. Second, this section empirically examines the extent of geographic misalignment by testing whether greater geographic potential mechanically leads firms to invest more in the corresponding renewable technology. Third, the section investigates whether geographic misalignment imposes financial costs on firms.

5.1 Static Transition Decision Model and Empirical Implication

Economic Environment. I develop a parsimonious static model of a firm's investment decision in year t following the federal air pollution regulations. For simplicity, the firm must choose between two mutually exclusive technologies: a utility-scale solar power plant (representing renewables, R) or a natural gas power plant (representing fossil fuels, G). This choice frames a central tradeoff in the energy transition. Solar PV is chosen as the representative renewable because its output is mechanically affected by site-specific solar irradiance, and its land-based nature allows the model to focus on geographic site selection, in contrast to other options like offshore wind. Natural gas is chosen as the representative fossil fuel and path dependence choice. As shown in Section 4, gas power is the most prevalent post-coal-exit choice, making it a good benchmark for the transition analysis.

The firm's decision is conditioned on a vector of exogenous state variables, $X_{i,t}$, for firm i in state s at time t:

$$X_{i,t} \equiv \left(\phi^k_{i,t},\ P^e_{s,t},\ \text{Carbon}_{s,t},\ \text{RPS}_{s,t},\ \text{Subsidy}^k_{s,t},\ C^k_{s,t}\right),$$

where k denotes generation technology types with $k \in \{R, G\}$. $\phi^k_{i,t} \in (0, 1]$ is a site-specific proxy for geographic resource quality. For instance, if two solar PV sites have the same installed capacity, the site with the higher $\phi^R_{i,t}$ produces more electricity and has lower grid stability costs. Because the model does not allow gas output to vary with local resource quality, I normalize its resource-quality term to $\phi^G_{i,t} = 1$. The decision is also shaped by market and policy conditions, including the state retail electricity price converted to annual revenue per installed kW at full utilization ($P^e_{s,t}$, in \$/kW-year)[20], a carbon price ($Carbon_{s,t}$, in \$/ton), the state's Renewable Portfolio Standard target ($RPS_{s,t}$, in percentage points), a proxy for state-level investment subsidies ($Subsidy^k_{s,t}$, the number of active state-level direct financial support programs, including grants, rebates, and performance-based incentives), and technology-specific annual O&M costs ($C^k_{s,t}$, in \$/kW-year). Subsidies are assumed to be available only for solar technology, thus $Subsidy^G_{s,t} = 0$.

Investment and Profitability. The upfront capital investment per kW for technology k is given by:

$$I^k_{i,t} \;=\; \underbrace{\kappa^k}_{\text{CapEx}} - \underbrace{\psi\, \text{Subsidy}^k_{s,t}}_{\text{Scaled Subsidy}}. \tag{6}$$

[20] Project-level electricity prices can differ from published wholesale benchmarks due to ISO/RTO locational marginal pricing (LMP), transmission congestion, contract heterogeneity (e.g., bilateral contracts and PPAs), etc. For simplicity, I proxy the expected revenue price using the annual state-level average retail electricity price, which is typically less volatile than wholesale prices at high frequency and focuses on the cross-state variations in electricity value in my model setting.

Here, κ^k represents the technology-specific construction cost per kW. The parameter ψ translates the proxy for state-level solar technology subsidy programs ($Subsidy^R_{s,t}$) into a direct reduction in capital expenditure.

Assuming the economic environment remains fixed, the expected annual operating cash flow per kW in year t is:

$$\pi_i^k \;=\; \underbrace{\left(P^e_{s,t} + \eta_{\text{rps}}\, \text{RPS}_{s,t}\, \mathbb{1}\{k=R\}\right) \phi^k_{i,t}}_{\text{Revenue \& Compliance Benefit}} - \underbrace{C^k_{s,t}}_{\text{O\&M}} - \underbrace{\lambda_{\text{grid}}\,(1-\phi^k_{i,t})}_{\text{Grid Cost}} - \underbrace{\lambda_{\text{carb}}\, \text{CO}_2^k\, \text{Carbon}_{s,t}}_{\text{Carbon Cost}}. \tag{7}$$

The solar project's electric operating revenue scales with the geographic resource quality factor $\phi^R_{i,t}$. Compared to gas, solar generates renewable energy credits (RECs) that are used for RPS compliance and can also be sold for extra revenue. In other words, better geographic resource quality creates more renewable generation, leading to higher revenue and compliance benefit. A concern is that, such benefit from solar project is hard to quantify: firms may comply through other renewable investments or purchases of RECs, unless technology carve-outs require solar specifically. Thus, I treat the RPS target as a compliance pressure and an economic incentive for solar investment through the value of tradable RECs, and parameterize the expected per-MWh premium as $\eta_{rps}\, RPS_{s,t}$, a linear compliance-value response to the stringency of the state target.

Operating costs include three distinct economic forces. First, $C^k_{s,t}$ represents standard, technology-specific O&M expenditures, which also include fuel costs for gas projects. Second, the term $\lambda_{grid}(1-\phi^k_{i,t})$ imposes a cost associated with the interconnection and intermittency of solar power. This specification captures the economic reality that less reliable generation (a lower $\phi^R_{i,t}$) imposes greater costs on the grid system. This cost is a primary economic disadvantage of solar relative to dispatchable gas generation, for which the cost is zero (i.e., $\phi^G_{i,t} = 1$). Finally, the term $\lambda_{carb}\, CO_2^k\, Carbon_{s,t}$ represents the cost of carbon emissions, applicable only to gas ($CO_2^R = 0$). It internalizes the externality of pollution, where CO_2^G is a technological parameter for the emissions rate, and the cost is activated by public policy through the carbon price ($Carbon_{s,t}$).

Decision Rule and Empirical Implication. In a standard corporate finance setup (see, e.g., Brealey et al., 2022), assuming a constant discount rate $r > 0$ and perpetual cash flows, the net present value (NPV) of investing in technology k is:

$$\text{NPV}^k_{i,t} \;=\; -\,I^k_{i,t} + \frac{\pi^k_i}{r}. \tag{8}$$

Equation (8) implies a set of project cost, geographic resource quality, and policy margins that affect the firm's discrete transition choice. Figure 6 illustrates these margins in a two-dimensional cost space, where the horizontal axis is the present-value cost of a gas project ($x \equiv \kappa^G + C^G_{s,t}/r$), and the vertical axis is the present-value cost of a renewable project ($y \equiv \kappa^R + C^R_{s,t}/r$). Each panel divides this space into areas where the firm prefers renewable projects (solar PV; $k = R$), gas projects ($k = G$), or no investment. Panel A shows that geographic resource quality is a first-order margin for renewable adoption. As a counterfactual benchmark, if renewable performance were not tied to site conditions (i.e., $\phi^R_{i,t} = 1$), the renewable break-

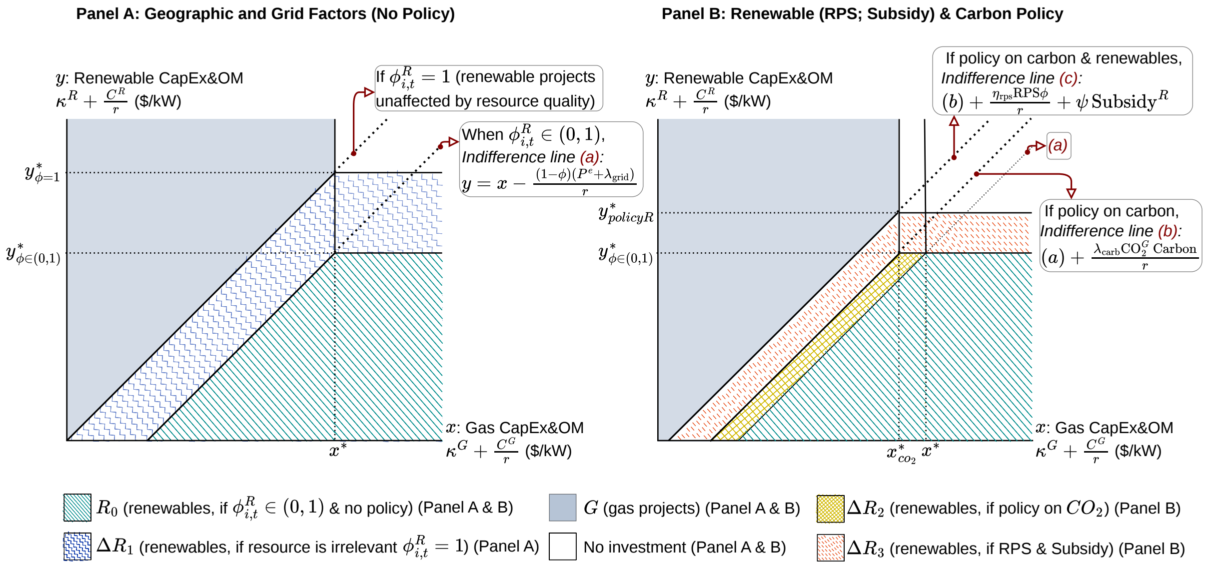

Figure 6. Effect of Resource and Policy Factors on Transition Choices. This figure shows how a firm's transition choice (k) between renewables ($k = R$), gas ($k = G$), and no investment varies with project cost, geographic resource quality, and policy factors. The horizontal axis is the present-value cost of a gas project, x (gas CapEx plus discounted fixed O&M; fuel costs included), and the vertical axis is the present-value cost of a renewable project, y (renewable CapEx plus discounted fixed O&M). All boundaries and indifference lines are implied by the NPV rule in equation (8). Panel A considers a no-policy environment (no carbon policy and no renewable policy such as RPS or subsidy). The cutoff x^* is the break-even boundary for gas investment, defined by $\text{NPV}^G = 0$; to the left of x^* gas projects are profitable. The cutoff $y^*_{\phi\in(0,1)}$ is the break-even boundary for renewable investment when geographic resource quality matters for renewables, capturing both output differences (capacity factor) and grid-related costs from intermittency; below this line renewable projects are profitable. The cutoff $y^*_{\phi=1}$ is the corresponding counterfactual renewable break-even boundary when resource quality does not affect renewables (stable grid supply and no resource-driven variation in effective generation). The line labeled (a) is the technology indifference line where $\text{NPV}^R = \text{NPV}^G$. Area R_0 indicates renewable choices under resource-relevant conditions and no policy (renewables preferred and profitable), while ΔR_1 indicates the additional renewable choices available when resource quality is not a constraint (relative to R_0), based on the corresponding renewable–gas indifference line. Panel B adds policy factors. The cutoff $x^*_{co_2}$ is the gas break-even boundary after incorporating carbon policy (gas is profitable to the left of $x^*_{co_2}$). The cutoff $y^*_{policyR}$ is the renewable break-even boundary after incorporating renewable policy (RPS pressure and investment subsidies). The indifference line shifts from (a) to (b) when carbon policy is introduced, and from (b) to (c) when renewable policy is also introduced. Area ΔR_2 denotes the additional renewable choices induced by carbon policy (relative to the no-policy case), and area ΔR_3 denotes the additional renewable choices induced by renewable policy on top of carbon policy. In both panels, the shaded area labeled G indicates gas choices (gas preferred and profitable), and the blank area indicates no investment (both technologies have negative NPV).

even cutoff would be $y^*_{\phi=1} = \frac{P^e_{s,t}}{r}$. In comparison, when resource quality matters in the realistic setting $\phi^R_{i,t} \in (0,1)$, the same installed capacity produces less electricity (lower capacity factors) and is costlier to integrate because intermittency raises balancing and reliability costs. These frictions: (i) raise the cost of renewable projects, lowering the renewable profitability break-even cutoff to $y^*_{\phi\in(0,1)} = \frac{P^e_{s,t}\phi^R_{i,t} - \lambda_{grid}(1-\phi^R_{i,t})}{r}$, and (ii) make renewables less attractive than gas, shifting the renewable–gas indifference line downward to line (a). On the other hand, the gas break-even threshold x^* (where $x^* = \frac{P^e_{s,t}}{r}$) is unaffected by resource quality, given gas power generation is assumed to be fully dispatchable, with no intermittency penalty and no resource-driven variation in capacity factors.

Panel B of Figure 6 shows how policy partially offsets these resource-induced disadvantages. A carbon price lowers the profitability of gas projects by adding an emissions cost, shifting the gas break-even threshold from x^* to x^*_{co2} (where $x^*_{co2} = \frac{P^e_{s,t} - \lambda_{carb} CO_2^G\, Carbon_{s,t}}{r}$) and shifting the renewable–gas indifference line from (a) to (b), which increases the area where renewables are preferred over gas. Renewable policy then further increases the incentive to invest in renewables through two channels: (i) RPS and investment subsidies raise the renewable break-even cutoff up from $y^*_{\phi \in (0,1)}$ to $y^*_{policyR}$ (where $y^*_{policyR} = \frac{\left(P^e_{s,t} + \eta_{rps} RPS_{s,t}\right)\phi^R_{i,t} - \lambda_{grid}(1-\phi^R_{i,t})}{r} + \psi\, Subsidy^R_{s,t}$); and (ii) these policies also shift the renewable–gas comparison toward renewables, moving the indifference line from (b) to (c) and expanding the renewable-choice area.

Given resource quality is the central margin for renewable adoption, I next derive the threshold level of solar resource needed for solar to outperform gas. Assuming a firm selects solar projects if $\text{NPV}^R_{i,t} \geq \text{NPV}^G_{i,t}$, this inequality defines a break-even threshold for the solar capacity factor, $\phi^\star_{i,t}$, which represents the minimum level of resource quality required for solar to be the more profitable investment:

$$\phi^\star_{i,t} \;=\; \frac{r\left(\kappa^R - \kappa^G - \psi\, \text{Subsidy}^R_{s,t}\right) + \left(C^R_{s,t} - C^G_{s,t}\right) + P^e_{s,t} - \lambda_{\text{carb}}\, \text{CO}_2^G\, \text{Carbon}_{s,t} + \lambda_{\text{grid}}}{P^e_{s,t} + \eta_{\text{rps}}\, \text{RPS}_{s,t} + \lambda_{\text{grid}}}. \tag{9}$$

Since $\phi^R_{i,t}$ is determined by local solar resource quality, equation (9) implies a geographic sorting mechanism: all else equal, firms with access to superior solar resources ($\phi^R_{i,t} \geq \phi^\star_{i,t}$) prefer solar (renewable transition), whereas those in less endowed locations prefer gas (path dependence). Crucially, the model also predicts that policy instruments directly affect this margin. Stricter climate policy (higher $RPS_{s,t}$ or $Carbon_{s,t}$) and more generous subsidies ($Subsidy^R_{s,t}$) lower the $\phi^\star_{i,t}$ threshold. This highlights a key interaction between geography and policy that can lead to potential misaligned transition: sufficiently strong policy can incentivize renewable investment even in regions where natural resource quality is only moderate.

5.2 Do Transitions Reflect Geographic Potential?

To empirically assess the extent of geographic misalignment, I first compare how renewable energy adoption and post-transition performance vary across firms with different levels of geographic potential using a reduced-form estimation. Specifically, I build percentile-based cutoffs for solar and wind potential and assign firms to groups according to their average geographic potential over the three years preceding the air pollution regulations.

Geographic potential is first measured at the state level using pre-policy averages, then aggregated to the firm level based on the states in which each firm operated before the policy. The potential variables are defined as solar irradiance (kWh/m^2/day) for solar and wind speed (m/s) for wind. Starting from the 5th percentile and increasing in 5-percentage-point increments, I define a potential group indicator equal to one if the potential value is greater than the current percentile cutoff and less than or equal to the cutoff plus 20 percentage points, capped at the 100th percentile when the upper bound exceeds it. This time-invariant grouping allows me to test whether firms located in areas with higher solar irradiance or wind speed are more likely to transition correspondingly.

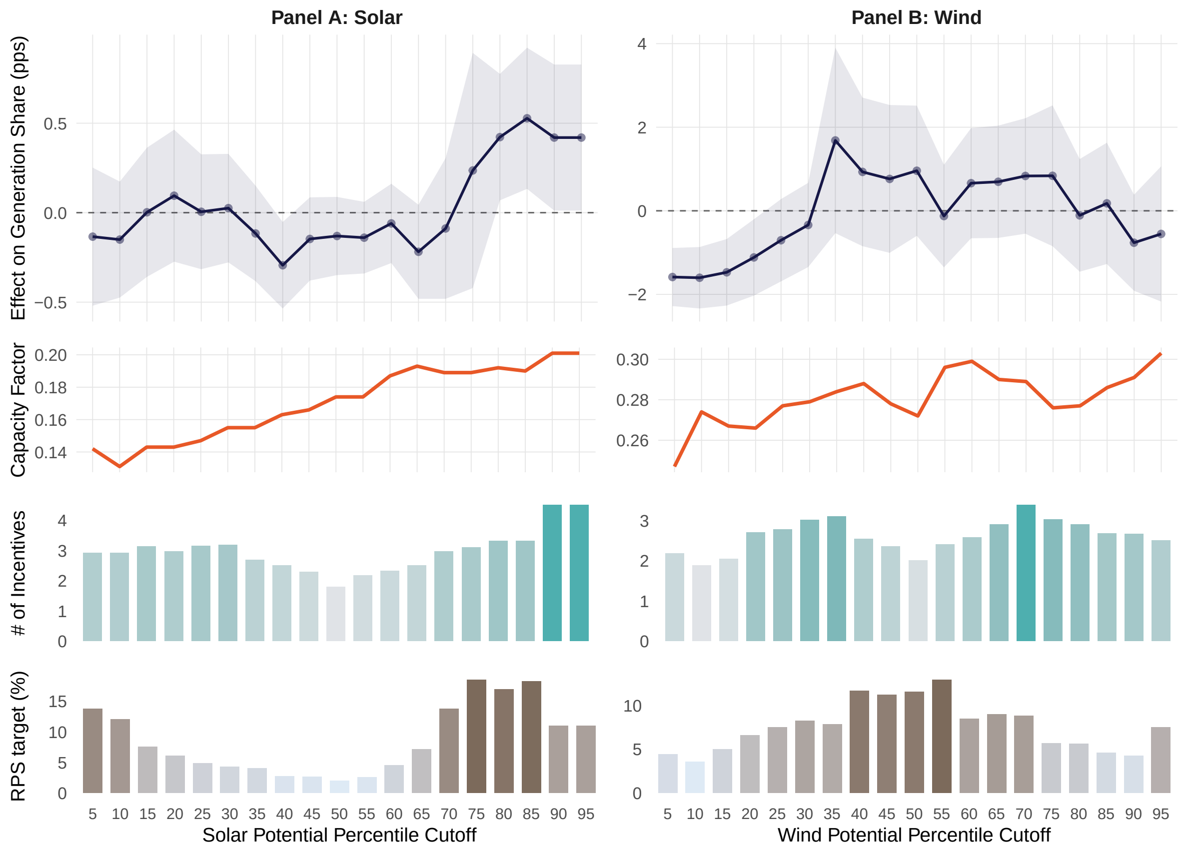

Figure 7. Geographic Potential and Solar/Wind Transitions. This figure presents results examining how firms' generation portfolios vary across levels of geographic potential for solar and wind energy after the air pollution regulations. Panel A presents results for solar, and Panel B for wind. In each panel, the top graph plots the estimated coefficients from the DiD specification based on equation (1), showing changes in generation share (percentage points) across potential percentile cutoffs with 95% confidence intervals. The estimation interacts the post-policy indicator $PostPolicy_{i,t}^{Air}$ with solar or wind potential group indicators. Firms are assigned to potential groups according to their average pre-policy geographic potential, measured as solar irradiance (kWh/m^2/day) for solar and wind speed (m/s) for wind, averaged over the three years preceding the regulations. Each potential group indicator equals one if a firm's average potential is greater than the current percentile cutoff and less than or equal to the cutoff plus 20 percentage points. The lower three graphs in each panel show firm-level averages by potential percentile: (i) capacity factors of the technology (second), number of technology-specific state-level financial incentives (third), and RPS targets (bottom).

The model specification follows the structure of equation (1), interacting the post-policy indicator $PostPolicy_{i,t}^{Air}$ with the list of solar and wind potential indicators defined by these percentile cutoffs. Figure 7 presents the estimation results along with firm-level averages for (i) capacity factors by technology, (ii) the number of state-level financial incentives available for each technology, and (iii) RPS targets across percentile cutoffs. First, the relationship is not linear between geographic potential and transition directions corresponding to such potential. For solar, the 80th to 100th percentile range of solar irradiance appears to represent a threshold that significantly motivates firms to transition toward solar generation. The average solar irradiance of this range is 5.61 kWh/m^2/day, which corresponds to the second-highest of the ten utility-scale PV resource classes defined by NREL.[21] For wind, the trend is not positively significant nor linearly increase across all percentiles, and fall significantly below zero when wind potential is really low.

Second, state-level financial incentives are correlated with the estimated transition effects for each technology. For solar, the correlation between the number of solar-specific incentives and the t-values of the transition estimates across percentiles is 0.78 ($p < 0.01$). Transition toward solar declines from the 15th to the 50th percentile, alongside fewer solar incentives. Beyond the 80th percentile of solar potential, stronger incentives coincide with higher t-values, suggesting that financial support is concentrated in high-solar-resource states. For wind, the transition effect rises from the 10th to the 35th percentile as wind incentives increase, and both flatten afterward. Average RPS targets also correlate positively with the t-values of the transition coefficients for solar (correlation = 0.67, $p < 0.01$) and wind (correlation = 0.63, $p < 0.01$). In the middle percentiles (40th–60th), higher irradiance does not translate into stronger transition effects, consistent with lower RPS targets in these states. For wind, RPS targets are higher in the middle percentiles and lower in the upper percentiles, mirroring the pattern in which estimated coefficients peak in the mid percentiles and decline at higher cutoffs.

5.3 Financial Costs of Misaligned Solar Siting

In this section, I move on to test whether the observed spatial misalignment of renewable transition carries measurable financial outcomes. Here I focus on solar PV plants again, because solar output is determined by site-specific irradiance, and its location choice reflects observable land-based geographic constraints. Specifically, I examine whether siting solar power plants in less favorable locations leads firms to realize lower electricity-related financial performance per unit of installed capacity.

I estimate a firm-plant-level specification that compares post-operation electric revenues across solar and fossil projects. The outcome variable is $Revenue_{i,t}^{perMW}$, the EIA-861 electric operating revenue (in millions USD) per megawatt of owned nameplate capacity for firm i in year t. The electric operating revenue includes utilities' earnings from electricity sales to ultimate customers, delivery and transmission services, sales for resale, and other regulated electric activities under accrual accounting. Dividing the operating revenue by capacity yields a standardized performance proxy that captures revenue-generating efficiency relative to installed generation assets. To isolate post-operation effects at the project level and examine dynamic changes over time, the baseline model employs a stacked event-study specification:

$$
\begin{aligned}
\text{Revenue}_{i,t}^{\text{perMW}} = \alpha_0 &+ \beta_1 \left(\text{InvestType}_p \times \text{PostActive}_{p,t}\right) \\
&+ \beta_2 \left(\text{Solar}_p \times \text{PostActive}_{p,t}\right) \\
&+ \beta_3 \left(\text{InvestType}_p \times \text{Solar}_p \times \text{PostActive}_{p,t}\right) \\
&+ X'_{i,t-1}\gamma + Z'_{s(i),t-1}\delta + \theta_{i,p} + \lambda_\tau + \varepsilon_{i,p,t}.
\end{aligned}
\tag{10}
$$

[21]See National Renewable Energy Laboratory, *Annual Technology Baseline: Utility-Scale PV* (2023), https://atb.nrel.gov/electricity/2023/utility-scale_pv.

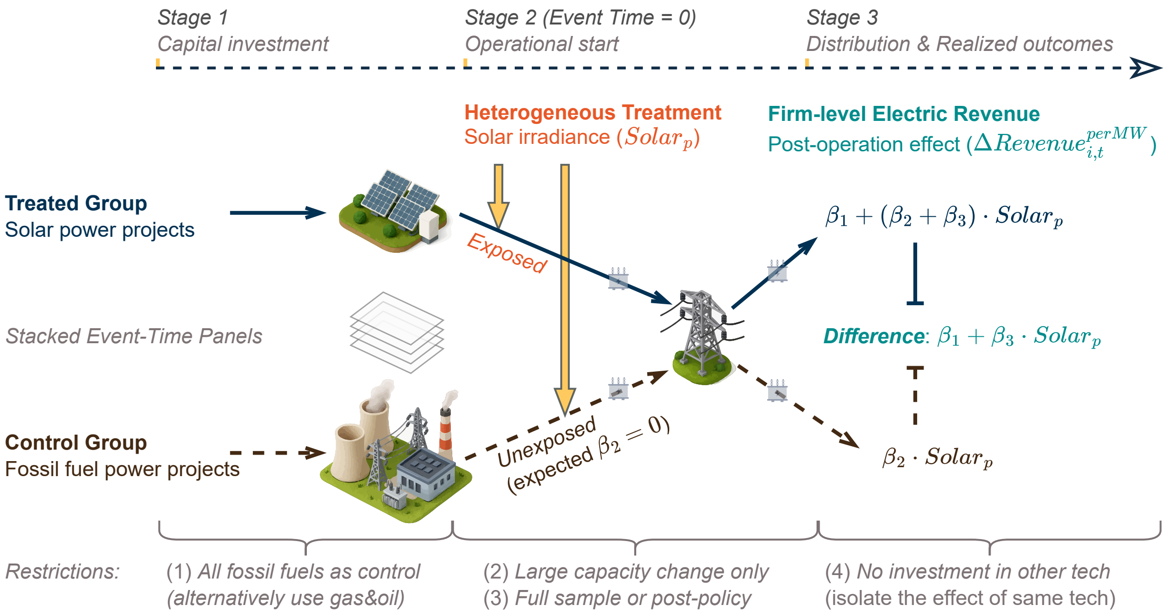

Figure 8. Identification Logic for Misalignment Financial Costs. This figure illustrates the identification strategy used to estimate the relative change in post-operation electric operating revenue for solar power projects, exploiting the heterogeneous treatment effect of solar irradiance based on the technical characteristics of solar versus fossil fuel power generation. At the top, the project timeline is shown from capital investment to post-operation stage, with the operational start defined as event time 0. At the bottom, the figure outlines the sample restrictions applied at each transition stage, designed to ensure valid control groups (unaffected by irradiance), retain only large-capacity projects that signal meaningful transition, and isolate a clean post-estimation window without later investments in other technologies that might interfere with the analysis. Coefficient definitions follow equation (10).

Here, $InvestType_p$ equals one if the plant p is solar-powered and zero if fossil-fueled. The variable $PostActive_{p,t}$ is an event-time indicator equal to one for the five years following the operational start of plant p, and zero in the four years prior. The time-invariant exposure $Solar_p$ measures average solar irradiance at the project site (kWh/m^2/day) during the post-operation period. This acts as a non-zero heterogeneous treatment for solar power plant after the operational start, given the operation of fossil fuel power plants is assumed to be indifferent from solar irradiance. β_3 identifies whether solar projects yield stronger revenue performance in sunnier locations relative to fossil plants. The differential revenue effect between solar and fossil projects at irradiance level $Solar_p$ is $\beta_1 + \beta_3 \cdot Solar_p$, while the slope of the post-operation revenue difference with respect to irradiance is β_2 for fossil projects and $\beta_2 + \beta_3$ for solar. Since irradiance is not expected to affect fossil generation technologies, β_2 is expected to be zero.[22]

Figure 8 illustrates the core identification logic. Solar irradiance functions as a heterogeneous treatment affecting only treated groups, which operate solar projects, and not control groups operating fossil projects. The assumption is that solar irradiance influences the operational efficiency (primarily capacity factor) exclusively for solar projects, whereas fossil projects remain unaffected. This allows solar irradiance to provide exogenous variation in treatment intensity.

[22] In equation (10), the TWFE and state-year-level controls such as average electricity prices account for potential heterogeneity driven by higher solar output in sunnier states depressing wholesale prices relative to cloudier regions. For literature related to the merit order effect, see, e.g., Peura and Bunn (2021), Antweiler and Muesgens (2021), Bushnell and Novan (2021), and Hausman (2025).

The analysis is based on the EIA–FERC matched sample of generation investments and applies several sample restrictions to ensure comparability and interpretability of results: (i) control groups are defined in two ways: either including all fossil fuel projects, or only gas and oil projects (less pollutive options compared to coal). (ii) To focus on economically meaningful investments, I retain only projects whose installed capacity exceeds 5 percent of the firm's total utility-owned generation capacity at the time of investment. (iii) The estimation uses both the full sample period, and a subsample limited to the years following the adoption of federal air pollution regulations. (iv) The projects must have no additional investments in other generation technologies during the post-operation window, ensuring that post-operation outcomes capture the performance of a single clearly defined technology type.

In addition, stacking the panel around event time zero enables a rigorous matching between each treated group (solar projects initiating operations under heterogeneous irradiance levels) and their respective control groups (fossil projects, due to the persistent nature of irradiance eliminating any post-operation "untreated" solar project) at identical event times. I follow de Chaisemartin and D'Haultfoeuille (2024) to implement this matching design in the dynamic specification, which isolates the heterogeneous treatment effects from broader trends affecting similar groups, and allows irradiance level to have time-variation after project operation.

The regression includes firm-plant fixed effects ($\theta_{i,p}$) to control for persistent heterogeneity across projects, and event-time fixed effects (λ_τ) to absorb common patterns in revenue relative to the operational start. Lagged firm controls ($X_{i,t-1}$) include *ConstructCapEx* and *O&M*, controlling for upfront investment and ongoing operating costs; *Capacity*, $PortfolioShare^{Fossil}$, and *FuelCost*, accounting for generation scale and exposure to generation input prices; and *FirmSize*, *Leverage*, *ROA*, *Liquidity*, and *Depreciation*, which control for firm financial characteristics. State level controls ($Z_{s(i),t-1}$) include *StateDemand*, $Subsidy^{Solar}$, *ElectricPrice*, and *CarveOut*, controlling for market demand and policy incentives affecting revenue.

Panel A of Table 7 presents the results. In column (1), based on the projects from the full sample period, β_1 is estimated at -0.868 and β_3 at 0.216. At the sample mean irradiance of 4.54 kWh/m^2/day, this implies a post-operation revenue change differential of \$113 per kW capacity ($-0.868 + 0.216 \times 4.54$) between solar and fossil fuel projects. The positive and statistically significant β_3 indicates that solar revenue performance improves sharply with irradiance: each additional unit of solar resource raises revenue by roughly \$216 per kW capacity relative to fossil benchmarks. Column (2) restricts the sample to projects initiated after the adoption of federal air pollution regulations; the revenue gap at mean irradiance is similarly at \$106 per kW capacity ($-0.820 + 0.204 \times 4.54$). Columns (3) and (4) replicate the specifications in columns (1) and (2), but use gas and oil projects as the control group, which represent less pollutive fossil operations and more realistic transition options than coal projects. The magnitude of the differential effect remains nearly unchanged, reinforcing the interpretation that solar projects' financial outcomes are strongly shaped by local irradiance conditions.

To further examine such differential effect by resource quality, Panel B of Table 7 replaces the solar irradiance measure with a binary high-resource indicator $\mathbb{1}\{Solar_p \geq 5.25\}$, which corresponds to NREL's Class 3 for utility-scale PV resources. The Class 3 threshold marks

Table 7: Financial Payoff of Transition: Geographic Advantage and Revenue

This table reports results examining how firms' investments in solar projects located in areas with differing levels of solar irradiance affect their electric operating revenue, benchmarked against gas projects. The dependent variable $Revenue_{i,t}^{\text{perMW}}$ measures annual electric operating revenue (in million USD) per megawatt of owned nameplate capacity for firm i in year t. The indicator $InvestType_p$ equals one if plant p is solar-powered and zero if fossil-fueled. In columns (1) and (2), fossil generation (control group) includes all fossil fuel types; in columns (3) and (4), the control group is restricted to gas- and oil-fired generation plants (less pollutive fossil fuels). $PostActive_{p,t}$ equals one for the five years following plant p's operational start and zero for the four years prior. $Solar_p$ denotes the average level of solar irradiance at the plant's location, measured in kWh/m^2/day following the operational start. In Panel B, this variable is used to define a high-resource indicator equal to one if $Solar_p \geq 5.25$, the minimum threshold for Class 3 solar zones as classified by NREL for utility-scale PV resources. For both panels, columns (1) and (3) include all investment events in the sample period, while columns (2) and (4) restrict to events dated after the adoption of the air pollution regulations (CSAPR or MATS). All regressions include firm-plant fixed effects ($\theta_{i,p}$) and event-time fixed effects (λ_τ), where $\tau = t - g_p$ denotes the number of years since plant p began operation, with cohort g defined as the first year of operation. Firm-level controls include lagged values of $O\&M$, $ConstructCapEx$, $FirmSize$, $Leverage$, ROA, $Liquidity$, $Depreciation$, $Capacity$, $PortfolioShare^{Fossil}$, and $FuelCost$. State-level controls are averaged at the firm level and include lagged $StateDemand$, $Subsidy^{Solar}$, $ElectricPrice$, and $CarveOut$. All variable definitions are in the Appendix. Standard errors are clustered at the firm level. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Panel A. Electric Operating Revenue: Solar vs. Fossil Fuels

Dep. Var.: Revenue$_{i,t}^{\text{perMW}}$	(1)	(2)	(3)	(4)
InvestType$_p$ × PostActive$_{p,t}$	-0.868***	-0.820***	-0.859***	-0.774***
	(0.188)	(0.150)	(0.206)	(0.164)
Solar$_p$ × PostActive$_{p,t}$	-0.014	-0.023	-0.015	-0.024
	(0.031)	(0.035)	(0.032)	(0.038)
InvestType$_p$ × Solar$_p$ × PostActive$_{p,t}$	0.216***	0.204***	0.214***	0.194***
	(0.040)	(0.034)	(0.044)	(0.037)
Firm-project FE	Yes	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes	Yes
Controls (Firm & State)	Yes	Yes	Yes	Yes
Sample	All cases	Post-policy	All cases	Post-policy
Control Group (InvestType$_p$ = 0)	Fossil Fuels	Fossil Fuels	Gas&Oil	Gas&Oil
Observations	1,274	911	1,150	853
Adjusted within R^2	0.165	0.289	0.168	0.296

Panel B. Solar Resource Quality and Revenue Differences

Dep. Var.: Revenue$_{i,t}^{\text{perMW}}$	(1)	(2)	(3)	(4)
InvestType$_p$ × PostActive$_{p,t}$	0.045	0.045**	0.043	0.047**
	(0.028)	(0.020)	(0.032)	(0.021)
$\mathbb{1}\{\text{Solar}_p \geq \text{C3}\}$ × PostActive$_{p,t}$	-0.043	-0.079	-0.053	-0.080
	(0.099)	(0.102)	(0.100)	(0.105)
InvestType$_p$ × $\mathbb{1}\{\text{Solar}_p \geq \text{C3}\}$ × PostActive$_{p,t}$	0.292***	0.302***	0.300***	0.290***
	(0.101)	(0.098)	(0.104)	(0.099)
Firm-project FE	Yes	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes	Yes
Controls (Firm & State)	Yes	Yes	Yes	Yes
Sample	All cases	Post-policy	All cases	Post-policy
Control Group (InvestType$_p$ = 0)	Fossil Fuels	Fossil Fuels	Gas&Oil	Gas&Oil
Observations	1,274	911	1,150	853
Adjusted within R^2	0.166	0.295	0.170	0.302

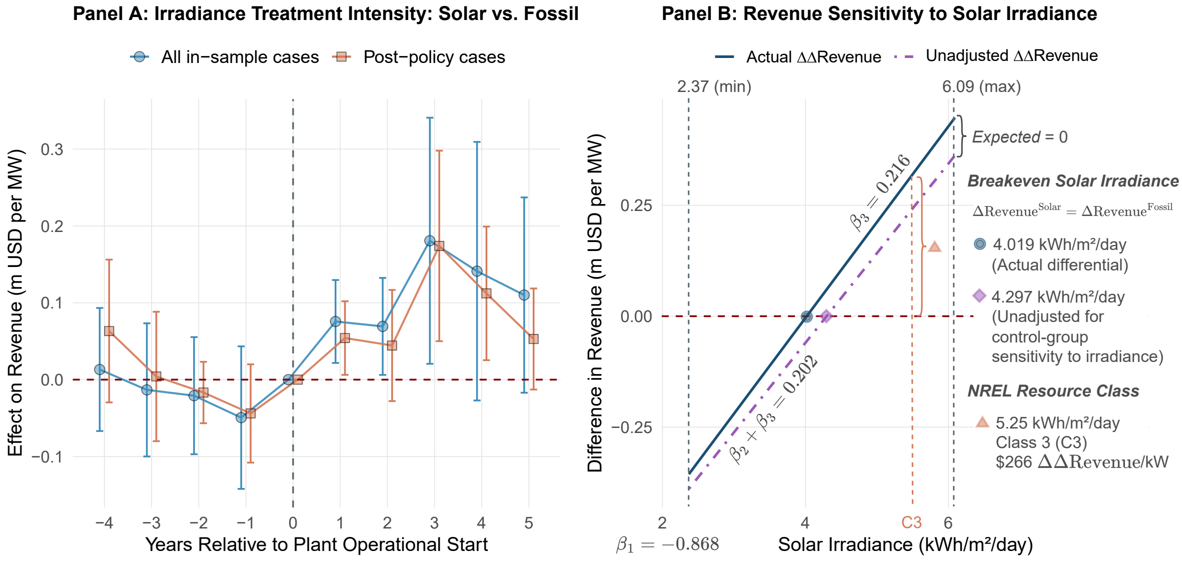

Figure 9. Revenue Differentials: Dynamic Effects and Sensitivity to Irradiance. This figure shows event-study estimates comparing post-operation revenue change differentials between solar and fossil generation projects, conditional on solar irradiance. Panel A plots dynamic differentials in electric operating revenue (million USD per MW of capacity) between solar and fossil plants over event time relative to operational start, with treatment intensity measured by annual solar irradiance around the plant location. The estimates correspond to the baseline specifications in Table 7 Panel A columns (1) and (2), which correspond to the full sample and the post–air pollution regulation period, respectively. Estimates are computed using the method of de Chaisemartin and D'Haultfoeuille (2024). Vertical bars denote 95% confidence intervals. Panel B illustrates the decomposed differential revenue effects (labeled as "$\Delta\Delta Revenue$") of solar versus fossil investment as a function of irradiance, corresponding to the baseline specification in Table 7 Panel A columns (1). The solid line shows the actual differential effects (adjusted for irradiance sensitivity of fossil plants), while the dashed line reflects the unadjusted differential effects (without considering the irradiance sensitivity of fossil plants).

the first resource class with an average solar capacity factor exceeding 30% (see footnote 21). The specification enables a triple-difference (DDD) interpretation of the post-operation revenue differential effects. β_3 remains positive and statistically significant across all specifications (point estimates near 0.29–0.30). This suggests that in Class 3 or better solar irradiance areas, solar projects generate approximately \$290 to \$300 more per kW capacity in post-operation revenue than solar projects in lower irradiance areas, benchmarked against the post-operation revenue difference for fossil fuel projects.

Figure 9 Panel A presents event-study estimates of the time-variant heterogeneous treatment effect using the estimator of de Chaisemartin and D'Haultfoeuille (2024). The results correspond to the baseline specifications in Table 7 Panel A columns (1) and (2), which correspond to the full sample and the post–air pollution regulation period, respectively. The treated status applies only to solar projects after their operational start, allowing identification of the differential revenue effect relative to fossil projects. For both specifications, the pre-treatment estimates confirm the parallel-trends assumption, and the coefficients become positive and statistically significant shortly after operation begins. The differential revenue effects persist over several years, indicating that solar investments in higher-irradiance areas consistently improve firms' electricity-related financial performance after deployment.

Panel B decomposes the baseline results in Table 7 (Panel A, column 1) to illustrate how post-operation revenue change differentials (labeled as "$\Delta\Delta Revenue$") vary continuously with solar irradiance. The estimated differential effects for fossil projects (β_2) is small and statistically indistinguishable from zero, confirming that fossil project revenues are insensitive to irradiance variation and thus provide a stable benchmark for comparison. The actual breakeven irradiance for solar projects at 4.019 kWh/m²/day, derived from β_3 and benchmarked against fossil projects that are expected to be indifferent to solar irradiance (i.e., $\beta_2 = 0$ is expected). This breakeven marks the point at which solar power investments start to generate positive effects on firms' electric operating revenue relative to fossil projects. When not adjusting for the control group's irradiance sensitivity (i.e., using $\beta_2 + \beta_3$), the breakeven shifts slightly to 4.279 kWh/m²/day. The figure further shows that at the lower bound of NREL Class 3 resources (labeled as "C3"), the corresponding differential revenue reaches about \$266 per kW of installed capacity, quantifying how irradiance-driven siting differences translate into meaningful financial outcomes for solar investments.

6 Dilemma: Environmental vs. Social

Sections 5.2 and 5.3 show that renewable investment does not always follow geographic potential and that this misalignment has real financial consequences. These findings suggest that aligning renewable policy more closely with resource quality could improve the transition, as implied by the model in Section 5.1. However, geography can also create a nonmonetary constraint. Conceptually, holding regional grid access, land availability and the regulatory environment fixed, renewable siting must still account for local resource quality, whereas gas siting does not face the same generation constraint.

Thus, relative to the coal-to-gas transition, the renewable path introduces an additional geographic margin into local site choice. Without identifying the frictions it creates, a more geographically aligned policy cannot be assumed to improve all outcomes. Unlike the financial cost above, this concern relates to local and within-firm worker adjustment: where a utility locates its new operations matters for workers affected by coal exit. Because renewable output depends directly on local resources, sites that support renewable generation may be farther from population centers and workers associated with the utility's retiring coal plants.[23] Greater distance between a utility's new sites and the residential geography associated with its retiring coal plants may increase the geographic frictions involved in retraining and transfer.

Therefore, the renewable path may deliver larger emissions reductions than the coal-to-gas path while imposing greater geographic frictions to worker adjustment. This comparison remains an empirical question because pollution and local opposition may also push fossil

[23]Solar power related occupations include research, engineering, manufacturing, development, construction, installation, and operations and maintenance (O&M), each requiring specific credentials and training (Hamilton, 2011). The discussion here focuses on utility-scale solar O&M technicians engaged in regular on-site work, including start-up inspections, repairs following power disruptions, diagnosis of inverter or electrical faults, component maintenance, and early fault detection (see Figure B.4 for an example job posting). System monitoring is also a key part of O&M work but can be performed off-site through centralized control centers. Emerging tools such as drone-based or AI-assisted robotic inspections are outside the scope of this discussion.

plants away from communities. In this section, I examine this tension first through workplace remoteness, then through the project co-location channel with incumbent fossil sites, and finally through accessibility from the residential geography associated with retiring coal plants.

6.1 Site Remoteness in Renewable Siting

I first examine whether utility-scale solar projects are systematically located in less populated areas, both relative to new fossil fuel projects and to the same firms' historical fossil sites. This analysis uses the EIA sample of 1,362 solar and gas- or oil-fired projects whose operation under the focal utility began after the federal air pollution regulations. Surrounding population density provides a consistent measure of proximity to population centers, which is a broad proxy for potential access frictions rather than directly measuring project employment, worker residences, or realized commuting behavior.

For each project p owned by firm i, I measure surrounding population density within concentric rings of radius $b \in \{0\text{–}10, 10\text{–}25, 25\text{–}60, 60\text{–}100\}$ km centered on the project's generation site. I refer to the first two bands as the near range and the latter two as the extended range. The multiple rings show whether differences in proximity to population centers are concentrated around the project site or persist over a wider area.

Next, for each radius b and year t, $Population_{p,t}^{(b)}$ is defined as the mean population per km^2 within the ring around the plant p. Panel A of Figure 10 compares the population density around a utility-scale solar PV plant and a gas power plant in Alabama. Both plants are adjacent to high-voltage ($\geq$ 220 kV) transmission lines. The gas plant is closer to dense settlements from 10 to 60 km, with Birmingham located within its 60–100 km ring. In contrast, the solar plant is surrounded by sparser population across all radii.

A comparison of new solar and fossil sites does not account for where each utility historically operated. Therefore, I benchmark each focal project against the same firm's legacy fossil generation geography. The observed coal exit makes this comparison relevant because utilities may hire locally at the new site or retrain and transfer incumbent fossil workers, although the LandScan measure neither identifies these workers nor assumes that such transfers occur.[24] I construct a project-year measure $CommunityAccess_{i,p,t}^{(b)}$ that uses the average population surrounding the firm's operating fossil plants in each pre-operation year and the population surrounding the focal site from its operation year onward:

$$\text{CommunityAccess}_{i,p,t}^{(b)} = \begin{cases} \dfrac{1}{|\mathcal{F}_{i,t}|} \displaystyle\sum_{j \in \mathcal{F}_{i,t}} \text{Population}_{j,t}^{(b)}/100, & \tau < 0, \\ \text{Population}_{p,t}^{(b)}/100, & \tau \geq 0, \end{cases} \tag{11}$$

where $\tau = t - g_p$ is event time, g_p is the first year in which project p operates under utility firm i, and $\mathcal{F}_{i,t}$ is the set of fossil fuel plants operated by firm i in pre-operation year t. Population density is rescaled from people per km^2 to per hectare (people/ha) by dividing by 100 for clearer

[24]For discussion of workforce transitions in coal-dependent regions, see, for example, Center for the New Energy Economy, *Coal Communities in Transition* (2022), https://cnee.colostate.edu/wp-content/uploads/2022/08/Coal-Communities-in-Transition.pdf.

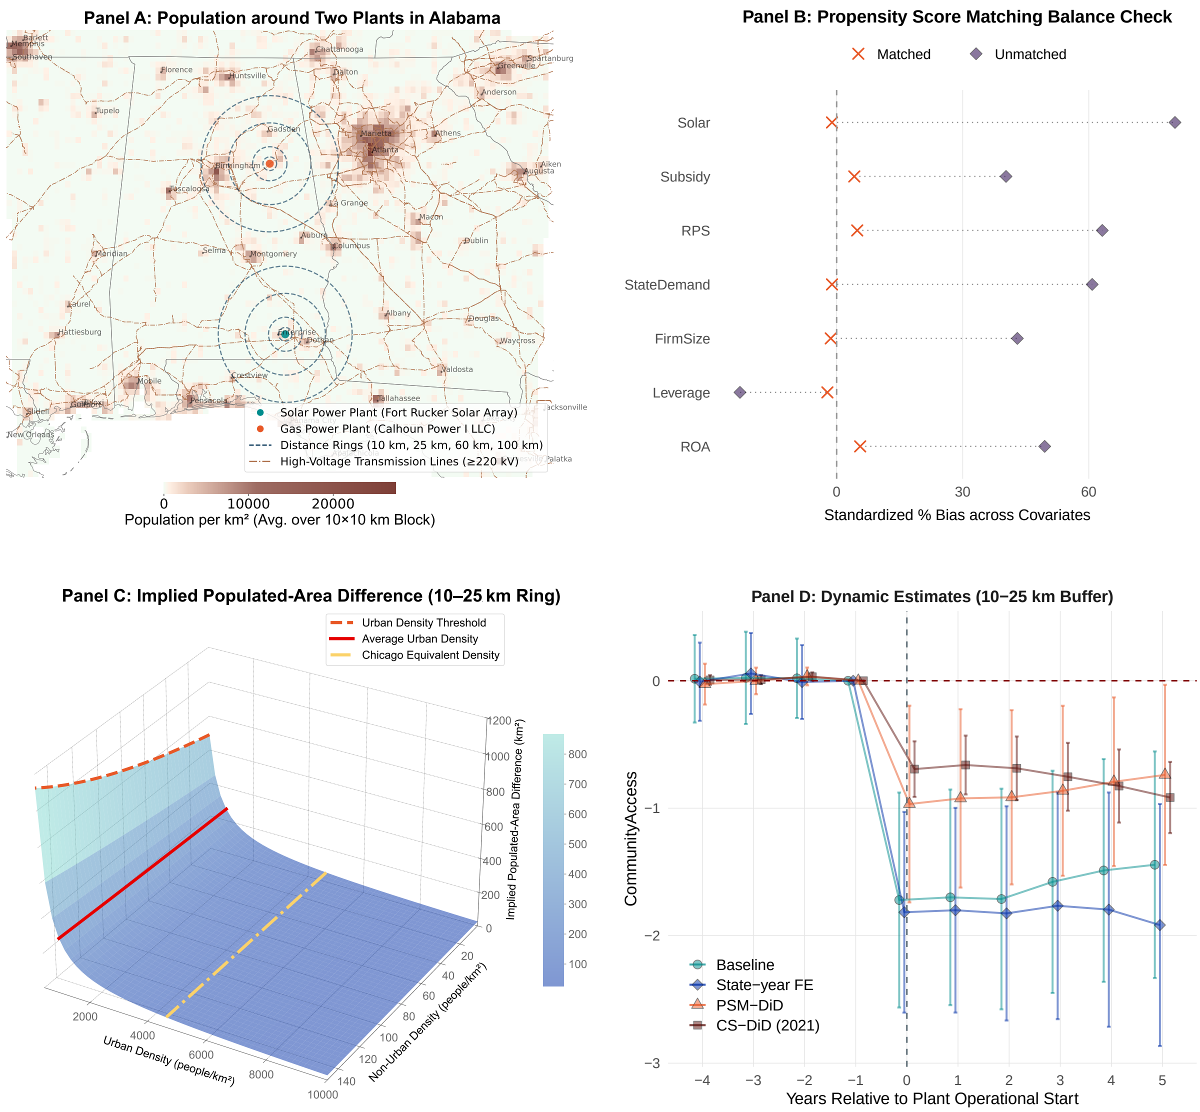

Figure 10. Project Siting Remoteness and Event Study Evidence. This figure illustrates the remoteness measure using plant cases and reports the sample matching, magnitude conversion, and dynamic estimates for the event study design. Panel A plots a utility-scale solar PV plant and a gas power plant in Alabama on a population density map (persons/km^2), with high-voltage transmission lines ($\geq$220 kV) and concentric rings at 10, 25, 60, and 100 km around each plant. Panel B reports propensity-score matching balance from a logit model for the treated group indicator $InvestType_p$ ($= 1$ for utility-scale solar PV projects and 0 for gas- or oil-fired projects) on lagged measures of utility-level solar-resource exposure, state RPS and subsidy exposure, state electricity demand, and firm fundamentals (assets, leverage, and ROA). Projects are then matched on the score to select a balanced sample of treated and control projects for the PSM-DiD model. Panel C converts the absolute solar–gas/oil project difference in population density (people/km^2) estimated in the baseline stacked DiD model in equation (12) for the 10–25 km ring into an implied populated-area difference (km^2). Under the two-zone decomposition, the ring area is multiplied by the absolute density gap and divided by the difference between assumed populated-zone density $\mathcal{P}_c$ and non-urban density $\mathcal{P}_s$. Reference lines correspond to the U.S. urban density threshold (386 people/km^2), the average density of U.S. urbanized areas (985.7 people/km^2), and central Chicago (4,656.3 people/km^2) in 2020. Panel D plots event-time coefficients for the 10–25 km ring from four specifications: baseline stacked DiD, stacked DiD with state–year fixed effects, PSM-DiD, and B. Callaway and Sant'Anna (2021) estimators, with the horizontal axis in event time $\tau \in [-4, 5]$ measured in years relative to the first year in which project p operates under utility i ($\tau = 0$). The dependent variable is $CommunityAccess^{(b)}_{i,p,t}$ with $b =$ 10–25 km. Points are estimates, vertical bars are 95% confidence intervals with firm-clustered standard errors, and the horizontal line marks zero.

coefficient presentation. Because the definition switches from the utility's operating fossil plants to the focal site at $\tau = 0$, $CommunityAccess^{(b)}_{i,p,t}$ measures a change in siting geography rather than a change in population at a fixed location.

As the baseline, I estimate a stacked event-study specification that compares these shifts in siting geography across investment types:

$$\text{CommunityAccess}^{(b)}_{i,p,t} = \alpha_0 + \beta\big(\text{InvestType}_p \times \text{PostActive}_{p,t}\big) + \theta_{i,p} + \lambda_\tau + \varepsilon_{i,p,t}, \qquad (12)$$

where $InvestType_p = 1$ for utility-scale solar PV projects (treated group) and 0 for gas- or oil-fired projects (control group). Two sample restrictions sharpen comparability: both groups consist of utilities that operated fossil plants before the focal project, and all focal utility–project events occur after the federal air pollution regulations. The post indicator $PostActive_{p,t} = 1$ for $1 \leq \tau \leq 5$ and 0 for $-4 \leq \tau \leq -1$; the operation year is excluded from this static specification. The specification includes firm–project fixed effects $\theta_{i,p}$ and event-time fixed effects λ_τ. The coefficient β is the difference between investment types in the change from the same utility's pre-operation fossil plants' surroundings to the focal project site. A negative estimate therefore indicates that solar projects represent a larger shift toward less populated surroundings than gas- or oil-fired projects.

The EIA sample maximizes project coverage but does not contain FERC financial variables for every utility. Thus, I report two complementary specifications. First, state–year fixed effects absorb policy and market conditions common to projects in the same state and year. This specification focuses on within-state-year variation but does not hold project-specific land or interconnection conditions fixed. Second, a propensity-score-matched DiD (PSM-DiD) uses the EIA–FERC matched sample to improve balance between solar and fossil projects on observed firm and state characteristics. Panel B of Figure 10 reports the balance check. Propensity scores come from a logit model for $InvestType_p$ on lagged measures of utility-level solar-resource exposure, state RPS and subsidy exposure, state electricity demand, and firm fundamentals (assets, leverage, and ROA). After matching, the standardized mean differences are small across the included covariates (mean 2.2%, maximum 5.6%). This procedure improves comparability on observables but does not eliminate selection on unobserved siting factors.

Table 8 reports the results. Panel A presents unadjusted comparisons of population density around treated and control projects in the operation year. Compared with gas- or oil-fired projects, solar PV projects have about 355 fewer people/km^2 within 0–10 km and 245 fewer people/km^2 within 10–25 km. The gap declines to about 93 people/km^2 at 25–60 km and is close to zero at 60–100 km.

Panels B and C of Table 8 report the estimates from equation (12). In the near range (Panel B), solar projects represent larger shifts toward less populated surroundings than gas- or oil-fired projects across all specifications. The baseline estimates imply differential declines of about 269 people/km^2 within 0–10 km and 161 people/km^2 within 10–25 km. After adding state–year fixed effects, the magnitudes increase slightly to 279 and 183 people/km^2, respectively. The matched-sample estimates are smaller but remain sizable, at approximately 166 and 86 people/km^2.

Table 8: Siting of Solar Projects and Workplace Remoteness

This table examines whether solar generation investments are sited in more remote locations, as measured by surrounding population density within concentric buffer zones. For pre-operation years, $CommunityAccess^{(b)}_{i,p,t}$ is the average population density in people/ha ($Population^{(b)}_{j,t}/100$) within buffer range b across the fossil fuel plants j operated by firm i in year t; from the operation year onward, it is the population density in people/ha ($Population^{(b)}_{p,t}/100$) within buffer range b of focal project p. The indicator $InvestType_p$ equals one if the project is solar and zero if it is gas- or oil-fired. $PostActive_{p,t}$ equals one in event years 1 through 5 and zero in event years -4 through -1; the operation year is excluded from the regressions in Panels B and C. Panel A reports comparisons of population density between treated (solar) and control (gas or oil) projects in the operation year, with equality of means tested using t-tests. Panel B reports stacked event-study estimates for near-site buffers (b = 0–10 km and 10–25 km), while Panel C reports estimates for extended buffers (b = 25–60 km and 60–100 km). Each buffer-range test includes a baseline specification, a specification with state–year fixed effects, and a PSM-DiD specification. All regressions include firm–project fixed effects ($\theta_{i,p}$) and event-time fixed effects (λ_τ), where $\tau = t - g_p$ and g_p is the first year in which project p operates under firm i. The sample covers utility–project events occurring after the federal air pollution regulations. All variable definitions are in the Appendix. Standard errors are clustered at the firm level. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Panel A. Univariate Comparison: Treated vs. Control Project Population

$Population^{(b)}_{p,t}/100$ (people/ha)	Treated		Control		Treated - Control	
	Mean	*SD*	*Mean*	*SD*	*Difference*	*SE*
b = 0–10 km	2.140	(4.332)	5.688	(9.808)	-3.548***	(0.414)
b = 10–25 km	1.705	(3.371)	4.156	(6.840)	-2.451***	(0.295)
b = 25–60 km	1.242	(1.853)	2.170	(2.905)	-0.928***	(0.133)
b = 60–100 km	1.048	(1.336)	1.053	(1.145)	-0.005	(0.067)
Number of projects	664		698		1,362	

Panel B. Community Accessibility in Near Range

Dep. Var.: $\text{CommunityAccess}^{(b)}_{i,p,t}$						
b =	*0–10 km*			*10–25 km*		
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{InvestType}_p \times \text{PostActive}_{p,t}$	-2.686***	-2.786***	-1.658**	-1.612***	-1.830***	-0.859**
	(0.764)	(0.772)	(0.750)	(0.448)	(0.452)	(0.349)
Firm-project FE	Yes	Yes	Yes	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes	Yes	Yes	Yes
State-year FE	No	Yes	No	No	Yes	No
PSM-DiD	No	No	Yes	No	No	Yes
Observations	10,570	10,570	4,118	10,570	10,570	4,118
Adjusted within R^2	0.037	0.028	0.047	0.031	0.030	0.045

Panel C. Community Accessibility in Extended Range

Dep. Var.: $\text{CommunityAccess}^{(b)}_{i,p,t}$						
b =	*25–60 km*			*60–100 km*		
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{InvestType}_p \times \text{PostActive}_{p,t}$	-0.773*	-1.016***	-0.329*	-0.623	-0.586**	-0.029
	(0.396)	(0.296)	(0.170)	(0.452)	(0.249)	(0.073)
Firm-project FE	Yes	Yes	Yes	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes	Yes	Yes	Yes
State-year FE	No	Yes	No	No	Yes	No
PSM-DiD	No	No	Yes	No	No	Yes
Observations	10,570	10,570	4,118	10,570	10,570	4,118
Adjusted within R^2	0.026	0.034	0.026	0.037	0.032	0.001

Such differences get smaller in the extended ranges in Panel C. In the baseline specification, the differential decline is about 77 people/km^2 at 25–60 km and 62 people/km^2 at 60–100 km, with only the former marginally significant. The within-state-year estimates remain negative at both ranges, whereas the matched estimate is close to zero at 60–100 km. The magnitude decrease indicates that the difference is concentrated nearer the project site. Its persistence with state–year fixed effects further shows that the pattern is not fully explained by common state policy and market conditions, though the specification does not identify the underlying siting mechanism.

To interpret the observed population density gap more intuitively, I adopt a stylized two-zone decomposition. I assume that the 10–25 km ring comprises populated zones with density $\mathcal{P}_c$ and non-urban zones with density $\mathcal{P}_s$, with both density benchmarks held common across solar and fossil fuel projects. Let w_c and w_s denote the respective land shares of populated and non-urban zones within a ring, such that $w_c + w_s = 1$. Average population density is therefore the weighted sum $\bar{\mathcal{P}} = w_c\mathcal{P}_c + w_s\mathcal{P}_s$. Let $\bar{\mathcal{P}}_{\text{gap}}$ denote the difference in average population density between fossil fuel and solar projects. Subtracting their weighted averages shows that this gap reflects the difference in w_c between the two project types, multiplied by $\mathcal{P}_c - \mathcal{P}_s$. The corresponding difference in populated land area is:

$$\text{Implied populated-area difference} = \pi(25^2 - 10^2) \times \frac{|\bar{\mathcal{P}}_{\text{gap}}|}{\mathcal{P}_c - \mathcal{P}_s}.$$

Panel C of Figure 10 visualizes this conversion across plausible values of $(\mathcal{P}_c, \mathcal{P}_s)$ using the absolute value of the baseline stacked DiD estimate, 161.2 people/km^2, in column (4) of Panel B of Table 8. I overlay three reference lines corresponding to the U.S. urban density threshold (386 people/km^2; Ratcliffe et al. (2016)), the average density of U.S. urbanized areas (985.7 people/km^2), and central Chicago (4,656.3 people/km^2) in 2020, retrieved from the U.S. Census Bureau. Setting $\mathcal{P}_s = 10$ people/km^2, the estimated gap corresponds to approximately 707 km^2 at the U.S. urban density threshold (42.9% of the ring), 272 km^2 at the average U.S. urbanized-area density (16.5%), and 57 km^2 at central Chicago's density (3.5%). These quantities are stylized density-equivalent area benchmarks rather than direct estimates of developed land.

Panel D of Figure 10 plots dynamic estimates for the 10–25 km ring from the baseline stacked DiD, the specification with state–year fixed effects, the PSM-DiD, and the estimator of B. Callaway and Sant'Anna (2021). The pre-operation estimates are close to zero and show no systematic differential trend between the solar and control-project event stacks in their legacy fossil benchmarks. The post-operation estimates are negative across all four estimators, consistent with the results in Table 8.

6.2 Siting Continuity and Retirement-Linked Workforce Accessibility

The population evidence identifies a broad difference in the geography of new solar and fossil projects, but does not show whether either investment path continues the utility's existing operating geography. Existing fossil generation sites can lower development costs through reusable interconnection and other infrastructure, creating an incentive to locate new capacity nearby (e.g., Kumar, 2023; Tassi and Kittner, 2024). Therefore, I test whether fossil projects are more likely than solar projects to locate at or near the utility's incumbent fossil sites. A stronger co-location tendency for fossil projects would be consistent with the population evidence, suggesting greater siting continuity along the fossil path dependence, while solar projects may more often use new locations because they must also satisfy geographic constraints.

To measure this channel, I focus on the EIA-sample projects that open as newly commissioned generation under the utility, rather than existing plants that enter its portfolio through acquisition. For each project, I identify incumbent fossil plants in the same state in which the utility owned operating capacity in event year -1, one year before the project opened. An incumbent plant qualifies if the utility owned at least half of its operating fossil capacity. For $b \in \{2, 5\}$ km, $SiteContinuity_{i,p}^{(b)}$ equals one when new project p lies within b km of at least one qualifying incumbent fossil plant. Internet Appendix Section B.3 reports the detailed construction and sample sequence, and Figure B.1 illustrates the measure using actual project sites and the two distance thresholds. This procedure produces a project-level cross-section of 534 projects: 354 solar and 180 gas- or oil-fired projects across 133 utilities.

I then repeat the construction using only incumbent coal plants. This comparison addresses whether the general siting-continuity result reflects only gas projects opening near incumbent gas plants. Again, a coal site qualifies if the utility owned at least half of its operating coal capacity in event year -1. The coal-only specifications include projects with at least one qualifying same-utility coal site in the project state in that year. This sample contains 221 projects: 183 solar and 38 gas- or oil-fired projects across 52 utilities.

Siting continuity does not identify worker adjustment. To examine that narrower margin, I next restrict the sample to projects that can be linked to a coal retirement by the same utility in the same state around the project's opening. This procedure isolates cases in which internal redeployment of incumbent coal workers is plausible and the residential geography associated with the retiring workplaces can be constructed, without establishing that workers transferred to the new project or ruling out new hiring.

I ask two related questions on worker adjustment: (i) how much of the estimated residential geography associated with the retiring coal workplaces lies near the new site, and (ii) how far the new site is from that geography.[25] After requiring a fully observed retirement-link window and valid LODES workplace and residence information, the EIA sample yields 110 retirement-linked projects: 84 solar and 26 gas- or oil-fired projects across 22 utilities. The remaining projects may still represent transition investments, but they either lack a qualifying same-utility same-state retirement in the specified window, or do not have associated LODES residential geography under the workplace screen.

To construct the residential geography of retiring coal plants, I first use the LODES Workplace Area Characteristics (WAC) data to locate utility-sector employment blocks

[25] The baseline links retirements occurring from five years before through three years after project operation. The five-year pre-operation interval allows replacement investment to follow retirement with a lag, while the three-year post-operation interval retains projects that open before formal closure, when retraining or redeployment planning may begin. The project utility must own at least 50% of the linked retiring capacity, and each linked retirement must remove at least 50% of the plant's coal capacity. Internet Appendix Section B.3 details the sample construction.

associated with the linked coal workplaces, and obtain the number of NAICS 22 utility-sector primary jobs in those blocks.[26] WAC identifies utility-sector jobs at the workplace, but it does not report their residence locations. The LODES Origin-Destination (OD) data report residence-to-workplace flows only for a broader category combining trade, transportation, and utilities. I distribute the WAC utility-job total across residence blocks using each block's share of these broader OD flows.[27] This allocation assumes that, within a validated workplace block, utility-sector workers have the same residential distribution as workers in the broader OD category. When a project is linked to multiple retiring coal plants, I pool their estimated residential distributions for the utility-project event. Figure B.2 illustrates the construction using two actual utility-project cases.

For radius $b \in \{25, 100\}$ km, $ResidenceShare_{i,p}^{(b)}$ is the percentage of the pooled estimated residence weight associated with utility firm i's retiring coal workplaces that lies within b km of new project p. Each residence block is weighted by its estimated number of utility-sector jobs, so blocks associated with more jobs receive more weight. I then define $WorkforceDistance_{i,p}$ as the weighted average geographic distance from these residence blocks to the new project. Finally, for each residence-workplace flow, I subtract the distance to the corresponding coal workplace from the distance to the new project, and average these differences using the same weights. The resulting $AddedDistance_{i,p}$ captures the hypothetical increase in geographic separation if the estimated residential distribution remained fixed. The analysis is a project-level cross-section in which each new project contributes one observation. Table 9 reports the siting-continuity and retirement-linked estimates.

Columns (1) and (2) of Panel A show that solar projects are 14.7 percentage points less likely than gas- or oil-fired projects to continue the utility's incumbent fossil geography within 2 km and 16.3 percentage points less likely within 5 km. Columns (3) and (4) sharpen the comparison to incumbent coal sites. Solar projects are 20.2 percentage points less likely to locate within 2 km of such a site and 17.1 percentage points less likely to locate within 5 km. Thus, the siting-continuity result does not reflect only gas projects opening near incumbent gas plants. The evidence suggests a greater tendency toward co-location along the fossil-dependent transition path compared to the renewable path.

Panel B reports the retirement-linked estimates. Relative to gas- or oil-fired projects, the estimated residence shares within 25 and 100 km are 16.1 and 20.7 percentage points lower for solar projects, respectively. The residence-to-project distance is 72.6 km greater for solar, while

[26] The baseline uses the coal-coordinate block when it qualifies and otherwise selects the qualifying block with the most utility-sector jobs within 2 km; a project is retained when at least one linked coal plant is matched. Internet Appendix Section B.3 details the workplace screen. For robustness, Panel B of Table B.4 compares this baseline with specifications that (i) require the coordinate block itself, (ii) require every linked plant to be matched within 2 km, or (iii) extend the search to 5 km only for otherwise unmatched plants.

[27] For example, if WAC reports 100 utility-sector jobs in a workplace block and 20% of the broader-sector OD flow into that workplace originates in a given residence block, I assign that residence block an estimated weight of 20 utility-sector jobs.

Table 9: Siting Continuity and Retirement-Linked Workforce Accessibility

Panel A tests the co-location channel by asking whether new projects continue the utility's incumbent operating geography through development at or near its fossil plants. All incumbent capacity is measured in event year −1, one year before new projects' operation. Columns (1) and (2) require at least one same-state fossil plant for which the utility owned at least 50% of aggregate physical operating fossil capacity. Columns (3) and (4) condition on having at least one qualifying same-utility coal site in the project state in event year −1, where the utility owned at least 50% of the plant's physical operating coal capacity. $SiteContinuity^{(b)}_{i,p}$ equals one when at least one qualifying incumbent site lies within b km of project p, for $b \in \{2, 5\}$ km. Panel B examines whether solar projects are farther from the estimated residential geography associated with same-utility retiring coal workplaces. Its conditional sample retains projects linked to a recorded coal retirement by the same utility and in the same state during event years −5 through +3 around the new project's operation. The project utility must own at least 50% of the linked retiring capacity, and linked retirements must remove at least 50% of each plant's coal capacity. The sample further requires a valid LODES workplace for at least one linked coal plant and a corresponding residence measure. $ResidenceShare^{(b)}_{i,p}$ is the percentage of estimated utility-job residence weight located within b km of new project p. $WorkforceDistance_{i,p}$ is the weighted mean distance from the same residence blocks to the project p. $AddedDistance_{i,p}$ is the weighted mean difference between each residence block's distance to the new project p and its distance to the corresponding coal workplace. $InvestType_p$ equals one for solar and zero for gas or oil. All specifications include operation-year fixed effects, and standard errors are clustered by firm. The control mean is the unadjusted mean among gas- or oil-fired projects. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Panel A. Continuity with Incumbent Fossil Sites

Dep. Var.:	$\text{SiteContinuity}^{(b)}_{i,p}$			
	$b = 2$ km	$b = 5$ km	$b = 2$ km	$b = 5$ km
	(1)	(2)	(3)	(4)
InvestType_p	-0.147***	-0.163**	-0.202**	-0.171**
	(0.047)	(0.066)	(0.076)	(0.078)
Operation-year FE	Yes	Yes	Yes	Yes
Incumbent site set	Any fossil	Any fossil	Coal only	Coal only
Control group	Gas/Oil	Gas/Oil	Gas/Oil	Gas/Oil
Control mean	0.200	0.278	0.237	0.237
Observations	534	534	221	221
Adjusted within R^2	0.041	0.034	0.088	0.048

Panel B. Retirement-Linked Workforce Accessibility

Dep. Var.:	$\text{ResidenceShare}^{(b)}_{i,p}$		$\text{WorkforceDistance}_{i,p}$	$\text{AddedDistance}_{i,p}$
	$b = 25$ km	$b = 100$ km	(km)	(km)
	(1)	(2)	(3)	(4)
InvestType_p	-16.079*	-20.730**	72.563***	66.173***
	(8.583)	(9.658)	(19.139)	(19.028)
Operation-year FE	Yes	Yes	Yes	Yes
Control group	Gas/Oil	Gas/Oil	Gas/Oil	Gas/Oil
Control mean	13.769	45.525	131.427	62.653
Observations	110	110	110	110
Adjusted within R^2	0.144	0.043	0.064	0.054

the added-distance measure is 66.2 km greater. The latter estimate benchmarks separation from the new project against the corresponding coal workplace.[28]

[28]For robustness, Panel A of Internet Appendix Table B.4 reports a lower residence share within 60 km, and positive statistically significant weighted-median estimates for both distance outcomes. Panel

These estimates apply to the conditional sample in which a qualifying retirement and the LODES measure can be observed. They do not identify named utility employees, actual transfers, or worker welfare. Instead, they provide narrower evidence that solar projects are geographically less accessible from the estimated residential distribution associated with same-utility retiring coal workplaces than the fossil transition benchmark.

One possible counterargument is that remote siting may benefit receiving communities by creating employment in new areas. Internet Appendix Section B.3 examines this possibility using nearby utility-sector employment around new project openings, and finds no statistically detectable solar advantage in average employment growth relative to gas projects.

7 Model Estimation: Transition Tradeoff and Policy Counterfactual

The reduced-form results in Sections 5 and 6 show two key patterns on the post-coal-exit renewable transition. First, geographic misalignment is costly, and varies with firms' policy exposure. Second, steering investment toward renewables can raise social costs, i.e., the "E vs. S" dilemma. A hypothetical perspective is: if policy were more aligned with geography, would the transition yield fewer but more productive solar projects, so that both financial and social costs can be lowered? Answering such question requires the benefit–cost margin at which a firm switches technology, not just average effects.

Therefore, I estimate the static adoption model in Section 5.1. The model estimation has two core objectives: first, to translate the reduced-form results into the firm's choice margin by quantifying how changes in RPS targets, subsidies, carbon prices, upfront investment costs, grid intermittency, and local resource quality shift the breakeven between renewable projects and fossil fuel alternatives. Second, to run out-of-sample policy counterfactuals that test whether better geographic alignment raises generation efficiency with lower renewable adoption, providing solutions to mitigate both financial and social costs.

7.1 Probability of Technology Adoption

Let $\mathcal{D}_{i,t} \in \{0,1\}$ indicate whether firm i in state s invests in solar ($\mathcal{D}_{i,t} = 1$) or gas ($\mathcal{D}_{i,t} = 0$) projects in year t. The firm chooses to invest in solar projects if its site quality exceeds the breakeven threshold from equation (9):

$$\mathcal{D}_{i,t} = 1 \quad \Longleftrightarrow \quad \phi^{R}_{i,t} \;\geq\; \phi^{\star}_{i,t}.$$

This is the firm-level breakeven rule when the exogenous state variables including market prices, policies, and O&M costs in equation (9) are held fixed within state×year (s,t). Considering the within-state variation in firm traits in the actual data, I express the rule with observables by (i) mapping the project-level irradiance-converted capacity factor $\widehat{\phi}^{R}_{i,t}$ to economic resource quality, and (ii) allowing firm traits to shift the effective hurdle through path dependence and generation scale:

C re-estimates added distance using $[-2,+2]$ and $[-3,+3]$ retirement windows, and using a gas-only comparison group.

$$
\begin{aligned}
z_{i,t} &= \alpha_\phi\, \widehat{\phi}^R_{i,t} \;-\; \phi^\star_{i,t} \;+\; \beta_{\text{coal}}\,\text{coal}_{i,t} \;+\; \beta_{\text{gas}}\,\text{gas}_{i,t} \;+\; \beta_{\text{solar}}\,\text{solar}_{i,t} \\
&\quad + \beta_{\text{capT}}\,\text{capTot}_{i,t} \;+\; \beta_{\text{capP}}\,\text{capProj}_{i,t} \;+\; \beta_{\text{rem}}\,\text{Rem}_{i,t}.
\end{aligned} \tag{13}
$$

Here, α_ϕ links the capacity factor proxy $\widehat{\phi}^R_{i,t}$ to the economic resource quality that governs the indifference condition $\phi^R_{i,t} \geq \phi^\star_{i,t}$.[29] $coal_{i,t}$ and $gas_{i,t}$ are the proportion of firms' coal and gas generation capacity, capturing the path dependence from existing fossil generation, fuel contracting, and operating expertise. These forces tilt the firm toward gas, implying $\beta_{\text{coal}} \leq 0$ and $\beta_{\text{gas}} \leq 0$. $solar_{i,t}$ is the proportion of solar generation capacity, accounting for the prior solar generation exposure that might affect the adoption costs in supply-chain access, siting, etc. Thus, β_{solar} is expected to be positive. $capTot_{i,t}$ is firms' total generation capacity, capturing the impact of firm generation size in project decision, such as project scale, grid access, and financing capacity. $capProj_{i,t}$ denotes the capacity of the project, as larger required capacity discourages solar PV adoption ($\beta_{\text{capP}} \leq 0$).[30] Finally, $Rem_{i,t}$ is the site remoteness trait, measured as the negative population within the 10–25 km ring of the project site. β_{rem} is identified from variation in remoteness that is orthogonal to site quality and the portfolio traits (Internet Appendix Section B.4.1), therefore it measures whether remote sites attract solar investment for reasons other than resource quality, such as land availability or fewer siting constraints. All generation capacity share and size variables are standardized to mean zero and unit variance in the estimation sample.

The probability of technology adoption can be modeled as a logistic choice with two sources of variation. First, a state-year level multivariate shock σ_{env} to the exogenous vector $X^{emp}_{s,t} = (P^e_{s,t},\ Carbon_{s,t},\ RPS_{s,t},\ Subsidy^R_{s,t},\ C^R_{s,t} - C^G_{s,t})$. The simulated exogenous vector is:

$$
X^{\text{sim}}_{s,t} \;=\; X^{\text{emp}}_{s,t} + \sigma_{\text{env}}\, \upsilon_{s,t}, \qquad \upsilon_{s,t} \sim \mathcal{N}(0, \Sigma_X),
$$

where Σ_X is the sample covariance of $X^{emp}_{s,t}$ across state–years. In simulation, I compute a Cholesky factor L with $\Sigma_X = LL^\top$, draw $Z_{s,t} \sim \mathcal{N}(0, I)$, set $\upsilon_{s,t} = Z_{s,t} L^\top$, and form $X^{sim}_{s,t} = X^{emp}_{s,t} + \sigma_{\text{env}} \upsilon_{s,t}$. This provides $X^{sim}_{s,t}$ with covariance proportional to Σ_X and hence the empirical correlation structure.

[29] Based on the Utility-Scale PV Resource Classes in National Renewable Energy Laboratory, *Annual Technology Baseline: Utility-Scale PV* (2023), https://atb.nrel.gov/electricity/2023/utility-scale_pv, I map firm–year solar irradiance (in kWh/m^2/day) to a capacity-factor proxy using an exponential relationship calibrated on the NREL resource-class data. I construct class-level observations by assigning each resource class the midpoint of its irradiance range and the corresponding mean alternating-current (AC) capacity factor reported by NREL, then estimate parameters (α, β, γ) in the model $\phi^R = \alpha \exp(\beta\, \text{Solar}) + \gamma$ by nonlinear least squares on these class-level data. I apply the fitted mapping $\widehat{\phi}^R_{i,t} = \hat{\alpha} \exp(\hat{\beta}, \text{Solar}_{i,t}) + \hat{\gamma}$ to the site-specific solar irradiance of the plant in which firm i invests in year t, so that $\widehat{\phi}^R_{i,t}$ represents the expected fraction of potential generation relative to installed nameplate capacity for that plant-year.

[30] Institutional details: Typical utility-scale PV plants are much smaller than new natural-gas units, making plants with very large capacity more feasibly built with gas. For example, EIA reports average new plant capacity of photovoltaic plants of about 40 MW in 2023, versus natural-gas combined-cycle plants at about 160 MW; see U.S. Energy Information Administration, *Construction Cost Data for Electric Generators* (2023), https://www.eia.gov/electricity/generatorcosts/.

Next, for a given parameter vector θ:

$$\theta \;=\; \big(\eta_{\text{rps}},\, \lambda_{\text{carb}},\, \psi,\, \alpha_{\phi},\, \sigma,\, \sigma_{\text{env}},\, \beta_{\text{solar}},\, \beta_{\text{coal}},\, \beta_{\text{gas}},\, \beta_{\text{capT}},\, \beta_{\text{capP}},\, \beta_{\text{rem}}\big),$$

and environment draw b, the rule of adoption $z_{i,t}^{(b)}$ can be estimated following equation (13).[31] To further translate $z_{i,t}^{(b)}$ into a choice, I allow for an unobserved project–specific disturbance and assume a logistic distribution, so adoption is determined by a smooth choice probability:

$$p_{i,t}^{(b)}(\theta) \;=\; \Lambda\!\left(\frac{\kappa_0 + z_{i,t}^{(b)}(\theta)}{\sigma}\right), \qquad \Lambda(u) = \frac{1}{1+e^{-u}}. \tag{14}$$

Here, κ_0 is a common intercept that shifts the index for all firms and absorbs level differences not captured elsewhere. $\Lambda(\cdot)$ is the logistic cumulative distribution function (CDF). $\sigma > 0$ governs how sharply firms sort around the breakeven: a smaller σ makes choices concentrate near the indifference cutoff, while a larger σ reflects more dispersed private information and internal frictions.

The logistic form is convenient: the choice probability stays in $[0,1]$, rises monotonically with the decision index, and can be turned into a 0–1 decision in a single step. To generate a firm's choice, I draw one uniform random number and record adoption when that number falls below $p_{i,t}^{(b)}(\theta)$. For each firm i, environment draw b, and replication r,

$$\mathcal{D}_{i,t}^{(b,r)}(\theta) \;=\; \mathbb{1}\Big\{\, U_{i,t}^{(b,r)} \le p_{i,t}^{(b)}(\theta) \,\Big\}, \quad U_{i,t}^{(b,r)} \sim \text{Unif}[0,1], \tag{15}$$

so that $\Pr\big(\mathcal{D}_{i,t}^{(b,r)}(\theta) = 1\big) = p_{i,t}^{(b)}(\theta)$ because $\Pr(U \le p) = p$ for a uniform draw. The intercept κ_0 is set so that the model's average adoption probability equals the empirical adoption rate $\bar{\mathcal{D}}$:

$$\frac{1}{NB}\sum_{b=1}^{B}\sum_{i=1}^{N}\Lambda\!\left(\frac{\kappa_0 + z_{i,t(i)}^{(b)}(\theta)}{\sigma}\right) = \bar{\mathcal{D}}.$$

Equivalently, the simulated adoption rate matches the data by construction:

$$\frac{1}{NBR}\sum_{b=1}^{B}\sum_{r=1}^{R}\sum_{i=1}^{N}\mathcal{D}_{i,t(i)}^{(b,r)}(\theta) \;\approx\; \bar{\mathcal{D}}.$$

[31] Calibration: I fix four objects rather than estimate them. The discount rate is set to $r = 0.07$ per year, and the capital-cost gap $\Delta\kappa_t = \kappa^R - \kappa^G$ is the observed series of utility-scale solar minus natural-gas construction costs by the year the generator was installed, from the EIA generator construction-cost data. The cost difference falls from about \$4,400 to \$800 per kW over the sample. The other two are normalizations rather than empirical inputs: the grid access term is set to $\lambda_{\text{grid}} = 0.5$ and emissions intensity to $CO_2^G = 1$. Neither is separately identified, and each fixes only a level or scale against which the estimated coefficients are read. Internet Appendix Section B.4.3 describes the two normalizations, and Section B.4.4 re-estimates the model under alternative values of r and $\Delta\kappa_t$.

7.2 Estimation Results

I estimate the parameter vector θ by simulated method of moments (SMM).[32] As in generalized method of moments (GMM), estimation matches a vector of empirical moments to model–implied moments. SMM differs in how the model moments are obtained: they are simulated rather than derived in closed form. SMM is appropriate here because several target moments, such as within–state-year partial covariances, dispersion components, and tail indicators built from a logistic choice with a breakeven threshold ratio, have no tractable closed forms under the state–year shocks.[33]

The SMM sample consists of 768 utility-scale solar PV and gas power generation projects from the EIA-FERC sample across the post federal air pollution regulations period. I restrict to projects with nameplate capacity ≥ 40 MW, which is the median gas plant capacity in the sample. The cutoff keeps project scale comparable across technologies, because PV plant can be built at much smaller and modular scales compared to gas plant, which makes solar–gas choices not effectively a multi-factor tradeoff in my model specification at very low capacities (see footnote 30). For example, when the required nameplate capacity is very small, the project–capacity term in the decision equation (13) dominates and mechanically favors solar, thus the model's choice margin is no longer informative.

The estimator chooses θ to minimize the weighted distance between the $q = 15$ moments computed from the data, m^{emp}, and their model counterparts, $m^{\text{sim}}(\theta)$, built from the same firm–state–year panel: $Q(\theta) = (m^{\text{sim}}(\theta) - m^{\text{emp}})^\top W(m^{\text{sim}}(\theta) - m^{\text{emp}})$ (see, e.g., Hansen, 1982; Duffie and Singleton, 1993). I simulate the state–year environment $B = 200$ times, reusing the same random draws at every θ so that changes in fit across parameter values are not an artifact of the draws (on simulators, see McFadden, 1989; Pakes and Pollard, 1989). Rather than draw whether each firm adopts, I use the adoption probability $p_{i,t}^{(b)}(\theta)$ directly, so that the model's fit responds smoothly to small changes in θ. I estimate the covariance of these moments by bootstrapping over the 45 state clusters (state-level block bootstrap, 1,000 replications).[34] The test of overidentifying restrictions does not reject the model ($p = 0.49$; Internet Appendix Section B.4.2).[35]

Table 10 summarizes the SMM estimates, with moment fit discussed in Section B.4.2. Regarding path dependence and scale effect, the slopes on legacy portfolios have the expected signs: $\widehat{\beta}_{\text{coal}} = -0.018$ and $\widehat{\beta}_{\text{gas}} = -0.031$ point toward gas for fossil-reliant firms, and $\widehat{\beta}_{\text{solar}} = 0.050$ suggests prior solar exposure facilitates adoption. Firm scale pushes toward

[32] For SMM, see, for example, McFadden (1989), Pakes and Pollard (1989), Duffie and Singleton (1993), and Rust (1994). For corporate finance research with implementations of SMM, see, for example, Hennessy and Whited (2007), Taylor (2010), Taylor (2013), Bakke and Gu (2017), and Wu (2018).

[33] See Internet Appendix Section B.4.1 for the selection procedure and the definitions of the 15 moments.

[34] Weighting and inference: a 15×15 covariance estimated from 45 state clusters is too imprecise to invert, so the weighting matrix W uses the inverse of its diagonal. Inference is based on the full bootstrap covariance rather than the diagonal: standard errors take the sandwich form, and the test uses the null distribution that the full covariance implies.

[35] Robustness: estimates are robust to $r \in \{0.05, 0.09\}$, to $\pm 20\%$ scaling of $\Delta\kappa_t$, to the simulation seed and depth, and to estimating the 2011–2017 and 2018–2024 subsamples separately (Internet Appendix Section B.4.4).

Table 10: Structural SMM Estimates

This table reports SMM estimates of 12 parameters using 15 moments. Standard errors are sandwich standard errors based on the full covariance matrix of the moments, estimated using 1,000 state-cluster bootstrap replications. The test of overidentifying restrictions does not reject the model ($p = 0.49$; moment fit is reported in Internet Appendix Section B.4.2). Calibrated objects are described in Internet Appendix Section B.4.3.

Description	Notation	Coefficient	*SE*
Solar share slope	β_{solar}	0.050	0.167
Coal share slope	β_{coal}	-0.018	0.062
Gas share slope	β_{gas}	-0.031	0.044
Total capacity slope	β_{capT}	0.047	0.080
Project capacity slope	β_{capP}	-0.068	0.051
Remoteness site trait	β_{rem}	0.025	0.042
Subsidy pass-through	ψ	0.356	0.131
Carbon-price leverage on gas	λ_{carb}	0.099	0.013
RPS level effect	η_{rps}	0.380	0.213
Irradiance slope	α_{ϕ}	0.576	0.411
Idiosyncratic dispersion	σ	0.005	0.000
Environment dispersion	σ_{env}	0.046	0.022

solar ($\widehat{\beta}_{\text{capT}} = 0.047$) and large project size discourages it ($\widehat{\beta}_{\text{capP}} = -0.068$), consistent with the smaller typical PV plant scale.

The policy levers also enter with the expected signs: the RPS level raises the return to solar ($\widehat{\eta}_{\text{rps}} = 0.380$), additional state subsidy programs encourage adoption ($\widehat{\psi} = 0.356$), and the carbon price penalizes gas ($\widehat{\lambda}_{\text{carb}} = 0.099$). Internet Appendix Section B.4.4 further examines how these policy responses vary across sample periods. The economically notable contrast is over time: the subsidy and carbon-price coefficients are about 2.7 and 2.4 times as large in 2011–2017 as in 2018–2024, while the RPS coefficient is also about one-third larger in the earlier period. This pattern is consistent with policy incentives playing a larger role earlier in the transition, when solar faced a greater capital-cost disadvantage.

The irradiance–to–economics link is positive ($\widehat{\alpha}_{\phi} = 0.576$), consistent with more efficient generation at more irradiant sites. The idiosyncratic dispersion is estimated near its lower bound allowed in estimation ($\widehat{\sigma} = 0.005$): conditional on site quality, calibrated costs and prices, technology choices are close to deterministic, so the model attributes adoption to fundamentals rather than to unobserved project-specific shocks. The remoteness site trait is statistically insignificant ($\widehat{\beta}_{\text{rem}} = 0.025$, SE 0.042, and equally null using an alternative 10–60 km buffer ring): within the choice model, remoteness does not independently attract solar investment once irradiance and portfolio traits are held fixed. Thus, the evidence that policy-driven siting raises workplace remoteness rests on the reduced-form results of Section 6, and the model treats remoteness as a by-product of siting rather than a driver of it.

7.3 Counterfactual Analysis: Aligning Renewable Policy with Geography

Given the identified financial costs of geographic misalignment and social costs of engineering optimization during the renewable transition, the counterfactual analysis asks a narrow question that follows directly: if policy pressure is reallocated toward high–resource locations while holding the installed total capacity fixed, is there a scenario where the system produces more electricity even with fewer solar projects, thereby lowering financial costs from misalignment and mitigating the social costs from "E" vs. "S" dilemma?

I reallocate state–year renewable policy pressure using a proportional rule from the existing RPS allocation to a geography-aligned benchmark. Let $\widehat{\phi}^{R}_{s,t}$ be the state–year average irradiance–converted capacity factor (see footnote 29 for the calculation). Given $RPS_{s,t} \in [0, 100]$, the aligned benchmark $RPS^{\star}$ distributes the same total policy pressure in proportion to resource quality, i.e., $RPS^{\star}_{s,t} \propto \widehat{\phi}^{R}_{s,t}$ with $\sum_{s,t} RPS^{\star}_{s,t} = \sum_{s,t} RPS_{s,t}$. Thus, the changes reflect reallocation rather than overall stringency, and the proportional benchmark spreads policy pressure across all sunnier states rather than concentrating it only in the sunniest.

Then, I trace a path from the observed allocation to the aligned benchmark, with the path indexed by $\lambda \in [0, 1]$. Let $\Delta RPS_{s,t} = RPS^{\star}_{s,t} - RPS_{s,t}$ denote the change a state–year undergoes over the full path (λ moves from 0 to 1), and let $\mathit{Reallocation} = \frac{1}{2} \sum_{s,t} |\Delta RPS_{s,t}|$ denote the total policy pressure that reallocates from cloudy state–years to sunny ones. λ fixes a fraction of that total reallocation, sequenced to maximize solar creation where renewable policy tightens and to minimize solar loss where it relaxes: among gaining state–years ($\Delta RPS_{s,t} > 0$), those with the most gas capacity are realigned to $RPS^{\star}_{s,t}$ first, and among losing state–years ($\Delta RPS_{s,t} < 0$), those with the least solar capacity first. Thus, when λ starts to move from 0, a gaining state–year is realigned only if $\lambda\, \mathit{Reallocation} > \sum_{\text{ahead}} \Delta RPS$, in which case

$$RPS_{s,t}(\lambda) = RPS_{s,t} + \min\left\{\Delta RPS_{s,t},\ \lambda\, \mathit{Reallocation} - \textstyle\sum_{\text{ahead}} \Delta RPS\right\}, \tag{16}$$

and a losing state–year only if $\lambda\, \mathit{Reallocation} > \sum_{\text{ahead}} |\Delta RPS|$, in which case

$$RPS_{s,t}(\lambda) = RPS_{s,t} - \min\left\{|\Delta RPS_{s,t}|,\ \lambda\, \mathit{Reallocation} - \textstyle\sum_{\text{ahead}} |\Delta RPS|\right\}, \tag{17}$$

where $\sum_{\text{ahead}}$ sums over the state–years that precede (s, t) in its own queue, with gaining and losing state–years ranked separately as defined above. For instance, gaining state–year reaches $RPS^{\star}_{s,t}$ in full whenever the difference $\lambda\, \mathit{Reallocation} - \sum_{\text{ahead}} \Delta RPS$ covers $\Delta RPS_{s,t}$, and moves only partway otherwise. Thus, at a given λ, at most one gaining state–year and one losing state–year lie strictly between $RPS_{s,t}$ and $RPS^{\star}_{s,t}$. Intermediate values of λ represent partial alignment, and $\lambda = 1$ delivers $RPS^{\star}$ everywhere.

The counterfactual relies on two assumptions. First, the capacity the project must deliver is determined before the technology is chosen and does not respond to renewable policy. In other words, the firm serves the same local demand whether it builds solar or gas, so a reallocation changes which technology is built and where, but not how much capacity the firm installs. Second, the firm need not build on exactly the same site once the technology changes. In practice, a location chosen to site a gas project is not the location a firm would choose for a solar project serving the same demand, therefore the counterfactual allows one local relocation when, and only when, the technology switches.

The assumed alternative location is a representative site rather than a specific coordinate, described by two geographic proxies: the irradiance-implied capacity factor for site resource quality, and the population density of 10–25 km buffer for workforce proximity. Specifically, if a firm builds solar rather than the gas project observed in the data, the firm is assumed to seek out resource quality and locate at a site whose two traits are the averages across the utility-scale solar sites with irradiance in the top quartile in the same state. On the other hand, if a firm builds gas rather than the observed solar project, resource quality no longer matters and the firm locates at a site whose two traits are the averages across all utility-scale gas sites in the same state. Everything else the firm faces is held at its observed value.

Given the two locations, I evaluate the adoption index in equation (13) twice at each λ, once at the project's actual site and once at the representative site. The relocation option is a response to the reallocated policy rather than an independent opportunity, so a project may exercise the option only in the direction in which $RPS_{s,t}(\lambda)$ has moved, which keeps every switch in the exercise attributable to the policy reallocation itself. When the two evaluations agree, the project is built as observed. When they disagree, the comparison of $RPS_{s,t}(\lambda)$ with the observed $RPS_{s,t}$ decides the outcome: a gas project for which the model predicts solar at the representative site is built as solar if $RPS_{s,t}(\lambda) > RPS_{s,t}$, and a solar project for which the model predicts gas at the representative site is built as gas if $RPS_{s,t}(\lambda) < RPS_{s,t}$. In every other case the project is built as observed, including in any state–year with $RPS_{s,t}(\lambda) = RPS_{s,t}$.

I track four outcomes along the path. The first is the solar adoption rate: the share of the 768 projects that are solar. The second is the solar generation share: solar generation divided by solar plus gas generation. Solar generation is each project's capacity times its site's capacity factor, summed over solar projects, with relocated projects evaluated at the new site. Gas generation is the same sum at the sample average gas capacity factor of 0.26. The third is total annual emissions of all projects: annual gas generation times 0.37 tonnes of CO_2 per MWh.[36] The fourth is workforce proximity: the average population density in the 10–25 km ring around each project's site, again at the new site for relocated projects. Total installed capacity $K_{tot} = \sum_p capProj_p$ is fixed along the path, with only its composition moving. Thus, every net change in these outcomes can be decomposed into two directions, one from projects switching from gas to solar, and one from projects switching from solar to gas. The two directions cannot coexist in a same state–year, since a project switches to solar only where $RPS_{s,t}(\lambda) > RPS_{s,t}$ and to gas only where $RPS_{s,t}(\lambda) < RPS_{s,t}$.

Figure 11 reports the four outcomes along the reallocation path. The core result is an interior optimum: the path improves on the observed allocation and then reverses. The improvement comes in two parts, matching the two costs in Section 6. The "environmental" gain comes first, at $\lambda \approx 0.60$, where emissions reach their lowest level on the path. Panel B shows that every state–year that has given up policy pressure up to the point holds almost no solar capacity to lose. The solar generation share rises (30.97% to 31.65%), total generation rises (376.4 to 376.9 TWh from the same installed capacity), and total emissions fall (96.13 to 95.31 $MtCO_2$ per

[36]The rate is that of a combined-cycle gas plant, which describes most gas projects in the sample. It multiplies EIA's carbon dioxide emissions factor for natural gas, 116.65 pounds per million Btu, by a heat rate of about 7,000 Btu per kWh. This heat rate is the weighted-average design value EIA reports for combined-cycle capacity added over the sample period. Solar generation is treated as emissions-free.

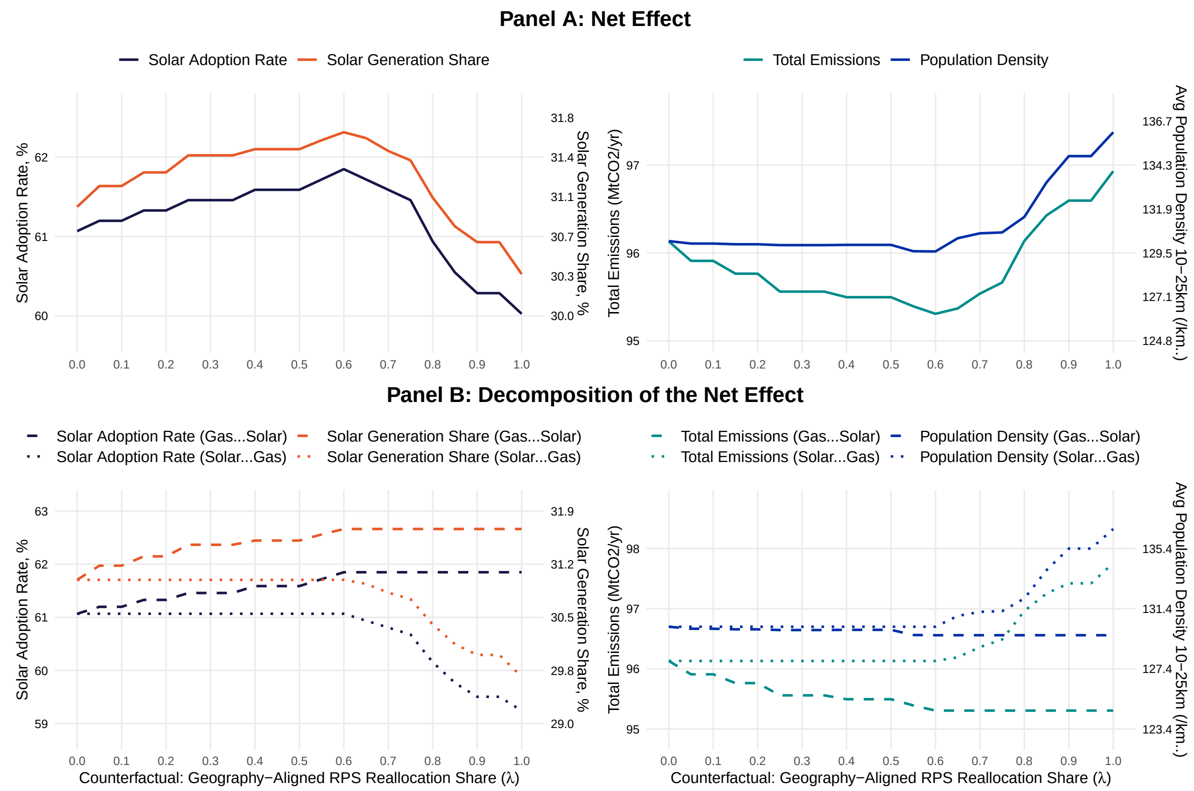

Figure 11. Counterfactual Analysis: Geography-Aligned RPS Reallocation. This figure traces the path from the observed RPS allocation ($\lambda = 0$) to the proportional geography-aligned allocation ($\lambda = 1$). λ fixes a fraction of that total reallocation, sequenced to maximize solar creation where renewable policy tightens and to minimize solar loss where it relaxes. Panel A reports the net effect: the left plot shows the solar adoption rate (left axis) and the solar generation share (right axis); the right plot shows total emissions (left axis, applying 0.37 tCO_2 per gas MWh) and the average population density in the 10–25 km ring around project sites (right axis). Panel B decomposes each net path into the gas-to-solar switches (dashed) and the solar-to-gas switches (dotted). Total installed capacity is fixed along the reallocation path.

year). Siting moves slightly away from population with average density around project sites falling from 130.2 to 129.6 per km^2.

Past $\lambda \approx 0.60$, more solar projects begin to switch to gas, and the "social" gain follows by $\lambda \approx 0.80$ without emissions rising above where the reallocation starts. At $\lambda \approx 0.80$, the solar adoption rate is slightly lower (61.07% to 60.94%), yet both the generation share (30.97% to 31.05%) and total generation (376.4 to 376.9 TWh) are still higher. Emissions are almost unchanged (96.13 to 96.14 $MtCO_2$ per year), and siting is closer to populated areas (130.2 to 131.5 per km^2).

At full alignment ($\lambda = 1$), solar-to-gas switches dominate, and the solar adoption rate falls below its starting level (61.07% to 60.03%). The siting continues toward populated areas (130.2 to 136.1 per km^2), but the solar generation share falls slightly (30.97% to 30.36%) and emissions rise (96.13 to 96.93 $MtCO_2$ per year). After $\lambda \approx 0.80$, reduced workplace remoteness comes at the cost of emissions rather than alongside them.

Overall, I find two scenarios along the policy reallocation that improve on the observed allocation at no cost to the other margin: the first is pro-environmental, where higher solar adoption and a higher generation share come with lower emissions and little change in workplace remoteness; the second is pro-social, where roughly the same solar adoption comes with essentially no additional emissions and lower workplace remoteness.

7.4 Extended Discussion and Practical Considerations

Transmission Constraints. The counterfactual analyses suggest that reallocating renewable policy toward the high-resource states can raise total generation efficiency. Notably, the gains arise without any project moving across state lines: the realignment works through where policy pressure is applied rather than through relocating projects between regions. However, the renewable end use may still remain regionally concentrated when long-distance delivery is limited. In the United States, the Eastern Interconnection, the Western Interconnection, and the Electric Reliability Council of Texas (ERCOT) operate largely as separate systems with limited transfer capability across seams. In practice, these deliverability constraints already show up as rising curtailment when local renewable supply exceeds what the grid can absorb or export, such as the wind curtailment in Midwest regions including the Southwest Power Pool (SPP) and Midcontinent Independent System Operator (MISO) footprints, and the wind/solar curtailment in the California Independent System Operator (CAISO).[37] At the same time, the interconnection queue process and required network upgrades have become constraints for bringing new clean projects online.[38] Therefore, a practical extension of the counterfactual is transmission investment that can move large blocks of low-cost clean power across regions. For instance, high-voltage direct current can be one candidate for long-distance transfer and for linking asynchronous systems through controllable interconnections.[39]

Renewable Intermittency. In the model, renewable grid integration costs enter through the calibrated wedge λ_{grid} and through the resource-quality margin itself: a lower capacity factor directly reduces the revenue side of the solar hurdle, so intermittency is priced primarily through geography. In practice, intermittency and system constraints exist as higher curtailment risk and output restrictions that reduce realized generation and raise operation uncertainty. Existing solutions are not perfect: short-duration battery storage can smooth within-day fluctuations, but it does not fully address long-term balancing needs as wind and solar expand. In turn, the intermittency costs may be reduced by investment and innovation in longer-duration storage and other longer-horizon system balancing options.[40]

Workforce and Community Impacts. The empirical results also point to a practical "just transition" issue that is inherently geographic. Coal retirements concentrate job losses in legacy communities, while renewable buildout shifts some construction and operations work

[37] See U.S. Energy Information Administration, *Why Are Midwest Grid Operators Turning Away Wind Power?* (2024), available at https://www.eia.gov/todayinenergy/detail.php?id=62406; and U.S. Energy Information Administration, *Solar and Wind Power Curtailments Are Increasing in California* (2025), available at https://www.eia.gov/todayinenergy/detail.php?id=65364.

[38] See Energy Technologies Area, Lawrence Berkeley National Laboratory, *Queued Up: Characteristics of Power Plants Seeking Transmission Interconnection* (2024), available at https://emp.lbl.gov/queues.

[39] See U.S. Department of Energy, Office of Electricity, *Connecting the Country with HVDC*, available at https://www.energy.gov/oe/articles/connecting-country-hvdc.

[40] See U.S. Department of Energy, Office of Electricity, *Achieving the Promise of Low-Cost Long Duration Energy Storage* (2024), available at https://www.energy.gov/media/328822.

toward areas that are often rural. Two practical considerations follow. First, concentrating renewable buildout only in the highest-resource areas, as the counterfactual analyses suggest, can improve generation efficiency but may intensify geographic job dislocation along the energy transition. This tension is not an artifact of the reallocation rule: in a complementary probe in which tradable renewable credits, rather than state targets, equalize the shadow price of compliance across states while holding national clean generation fixed, solar migrates to sunnier but more remote sites, so the environmental-versus-social tradeoff survives even under a pure market-design solution. Second, regarding the inspection and maintenance costs related to workplace remoteness, the development and adoption of automated inspection tools such as drone-based inspection for renewable assets and related electrical infrastructure might reduce on-site workload in hard-to-reach areas.[41]

8 Conclusion

This paper studies how U.S. electric utilities transition away from coal-fired generation in response to federal air pollution regulations, and examines the pathways, constraints and consequences of redirecting investment toward renewable energy. The analysis delivers three main takeaways. First, coal exit takes different forms, with plant closures becoming more likely where a substitution effect favors gas over coal. Second, RPS requirements and targeted incentives lead firms to increase investment in renewables, affecting both the likelihood and technology adoption of their transition. Third, geography shapes two transition frictions: firms earn less revenue per kilowatt of capacity from solar investments in low-irradiance regions, and the tendency to site projects in less populated areas raises social costs by increasing workplace remoteness.

These findings suggest that while air pollution regulations initiate the transition away from coal generation, the economic and social efficiency of the renewable buildout depends on how policy design interacts with geographic and institutional constraints. For policymakers, the counterfactual simulations show that aligning renewable mandates more closely with resource quality can mitigate these transition frictions, allowing policy to reduce emissions or workplace remoteness without materially worsening the other outcome.

[41]See, for example, U.S. Department of Energy, *An Operations and Maintenance Roadmap for U.S. Offshore Wind* (2024), available at https://www.energy.gov/sites/default/files/2024-05/operations-maintenance-roadmap-us-offshore-wind.pdf.

Appendix: Variable Definitions

A. Generation Capacity and Portfolio:

- $GeneratorSize_{j,t}$: Nameplate generation capacity (MW) of generator j in year t. Source: EIA Form 860.
- $Ownership_{i,j,t}$: Ownership share (0–1) of firm i in generator j in year t. Source: EIA Form 860.
- $PortfolioShare^{(k)}_{i,t}$: Share (percentage points) of owned generation capacity in type $k \in \{coal, gas, renewables, other\ clean\}$ for firm i in year t; the calculation follows equation (2). Source: EIA Form 860.
- $Capacity_{i,t}$: Ownership-weighted generation capacity owned by firm i in year t, computed as $\sum_j GeneratorSize_{j,t} \times Ownership_{i,j,t}$. As a control variable it is the total across all generation types, in log(1+MW). As an outcome variable it is the level for a single type k, denoted $Capacity^{(k)}_{i,t}$ and reported in GW for readability of the estimated coefficients. Source: EIA Form 860.

B. Policy Exposure and Timing Variables:

- $CoalReliance_i$: Ownership-weighted coal generation capacity share for firm i one year before the introduction of the federal air pollution regulations CSAPR and MATS. The calculation follows equation (2) without $\times 100$. Source: EIA Form 860.
- $InRPS_i$: Indicator equal to one if firm i operated in any state with an adopted renewable portfolio standards (RPS) one year before the introduction of the federal air pollution regulations CSAPR and MATS. Source: Lawrence Berkeley National Laboratory.
- $Incentive^{(r)}_i$: Indicator equal to one if firm i operated in any state offering financial incentives one year before the introduction of the federal air pollution regulations CSAPR and MATS, for $r \in \{Renewables, Solar, Wind\}$ (i.e., any renewables, solar-specific, and wind-specific). Source: Database of State Incentives for Renewables & Efficiency (DSIRE) for financial incentives.
- $FuelCostGap^{Coal\text{-}Gas}_i$: Standardized difference (USD/MWh) between firm i's coal fuel cost and the same-year industry average gas cost, measured in one year after the introduction of the federal air pollution regulations CSAPR and MATS. Source: EIA Form 923.
- $PostPolicy^{Air}_{i,t}$: Indicator equal to one if firm i is subject to the adoption of CSAPR (from 2011) or MATS (from 2012) in year t. Source: EPA CSAPR/MATS (Federal Register).

C. Firm Financials:

- $FirmSize_{i,t}$: log(1+total assets) of firm i at time t. Source: FERC Form 1.
- $Leverage_{i,t}$: Total debt over total assets of firm i at time t. Source: FERC Form 1.
- $ROA_{i,t}$: Net income divided by total assets of firm i at time t. Source: FERC Form 1.
- $Liquidity_{i,t}$: (current assets − current liabilities)/total assets of firm i at time t. Source: FERC Form 1.

- $Depreciation_{i,t}$: Accumulated depreciation divided by gross plant assets of firm i at time t. Source: FERC Form 1.
- $OM_{i,t}$: Operation and maintenance expenses of firm i at time t. Source: FERC Form 1.
- $ConstructCapEx_{i,t}$: log(1+construction-related capital expenditures (m$)) of firm i at time t. Source: FERC Form 1.
- $FuelCost_{i,t}$: Average fuel cost (USD/MWh) of firm i at time t. Source: EIA Form 923.

D. State-Level Market Controls:

- $CarveOut_{s,t}$: Technology-specific minimum requirements within the RPS (e.g., solar set-aside) in state s in year t. Source: Lawrence Berkeley National Laboratory.
- $ElectricPrice_{s,t}$: Average retail electricity price (cents/kWh) in state s in year t. Source: EIA Electricity Retail Sales.
- $StateDemand_{s,t}$: Total electricity demand (TWh) in state s in year t; used as log(1+$StateDemand_{s,t}$) and averaged across a firm's states when needed. Source: FERC Form 714.

E. Exit Channels and Compliance Outcomes:

- $ExitType^{(m)}_{i,t}$: Exit outcome for firm i in year t: no exit (m=0), divestiture without closure (m=1), or closure (m=2); if both occur in the same year, coded as closure. Source: EIA Form 860.
- $Divestiture_{i,t}$: Indicator equal to one if firm i reduces ownership in any coal-fired generator without an EIA retirement flag in year t. Source: EIA Form 860.
- $Closure_{i,t}$: Indicator equal to one if firm i retires any coal-fired generator in year t. Source: EIA Form 860.
- $AbatementCost^{log}_{i,t}$: Pollution-control/environmental equipment expenditures aggregated to the firm–year; defined as the natural logarithm of one plus abatement expenditures (m$). Source: EIA Form 860.
- $SO2_{i,t}$: Sulfur dioxide emissions of firm i in year t (k lbs). Source: EPA Continuous Emissions Monitoring System (CEMS).
- $NOx_{i,t}$: Nitrogen oxides emissions of firm i in year t (k lbs). Source: EPA Continuous Emissions Monitoring System (CEMS).
- $CO2_{i,t}$: Carbon dioxide emissions of firm i in year t (k tons). Source: EPA Continuous Emissions Monitoring System (CEMS).

F. RPS Discontinuity Constructs:

- $PortfolioShare^{Renewables}_{i,forward}$: Average renewable portfolio share for firm i in the forward window (either t+1 or mean of t+1 to t+3) used as the second-stage outcome of the fuzzy regression discontinuity design. Source: EIA Form 860.
- $d_{i,t}$: Distance to the binding annual RPS threshold for firm i in year t, defined as installed renewable capacity share minus the binding state–year Renewable Portfolio Standard (RPS) annual target.

- $BelowRPS_{i,t}$: Indicator equal to one if firm i's installed renewable capacity share is below the binding state–year Renewable Portfolio Standard (RPS) annual target in year t; for multi-state firms, the target is the maximum across their operating states.

G. Project-Level Variables:

- $Revenue^{perMW}_{i,t}$: Electric operating revenue (m\$) per MW of owned nameplate capacity for firm i in year t. The electric operating revenue includes utilities' earnings from electricity sales to ultimate customers, delivery and transmission services, sales for resale, and other regulated electric activities under accrual accounting. Source: EIA Form 861 (revenue) and EIA Form 860 (capacity).
- $InvestType_p$: Indicator equal to one for utility-scale solar PV projects and zero for fossil projects. Source: EIA Form 860.
- $PostActive_{p,t}$: Event-time indicator equal to one for $1 \leq \tau \leq 5$ after project p becomes operational and zero for $-4 \leq \tau \leq -1$; the operation year $\tau = 0$ is excluded from the regressions. Source: EIA Form 860.
- $Solar_{p,t}$: Global horizontal irradiance (kWh/m^2/day) measured at the site location of plant p in year t. Source: ECMWF ERA5.
- $Wind_{p,t}$: Wind speed (m/s) measured at the site location of plant p in year t. Source: ECMWF ERA5.
- $Population^{(b)}_{p,t}$: The average population density (people/km^2) measured within ring $b \in \{0\text{–}10, 10\text{–}25, 25\text{–}60, 60\text{–}100\}$ km at the site location of project p in year t; rescaled to people/ha by dividing by 100 in regressions. Source: Oak Ridge National Laboratory's LandScan.
- $CommunityAccess^{(b)}_{i,p,t}$: Equals firm i's historical average fossil-plant population density within ring $b \in \{0\text{–}10, 10\text{–}25, 25\text{–}60, 60\text{–}100\}$ km for $\tau < 0$ and project p's ring-b density for $\tau \geq 0$, comparing new project siting to legacy fossil siting. Source: Oak Ridge National Laboratory's LandScan; EIA Form 860.
- $SiteContinuity^{(b)}_{i,p}$: Indicator equal to one when new project p lies within $b \in \{2, 5\}$ km of a qualifying same-utility incumbent site in the same state. Incumbent capacity is measured one year before project operation. The candidate set includes all operating fossil capacity in the first specification and only operating coal capacity in the coal-only specification; the utility must own at least 50% of the corresponding physical capacity at the plant. Source: EIA Form 860.
- $ResidenceShare^{(b)}_{i,p}$: Percentage of the estimated residence weight associated with utility-sector workplaces at same-utility retiring coal plants that lies within b km of new project p, for $b \in \{25, 60, 100\}$ km. Source: Census LEHD Origin-Destination Employment Statistics (LODES); EIA Form 860.
- $WorkforceDistance_{i,p}$: Estimated utility-job-weighted mean distance from the residence blocks associated with same-utility retiring coal workplaces to new project p. Source: Census LODES; EIA Form 860.

- $AddedDistance_{i,p}$: Estimated utility-job-weighted mean difference between each residence block's distance to new project p and its distance to the corresponding retiring coal workplace. Source: Census LODES; EIA Form 860.
- $UtilityJobs^{(b)}_{p,\tau}$: Number of NAICS 22 primary jobs reported by LODES WAC in the fixed set of Census blocks within b km of physical project p in event year τ. Source: Census LODES.
- $UtilityGain^{(b)}_{p}$: Indicator equal to one if the mean of $UtilityJobs^{(b)}_{p,\tau}$ in event years $+1$ through $+3$ exceeds its mean in event years -3 through -1. Source: Census LODES.
- $UtilityEntry^{(b)}_{p}$: Indicator equal to one if mean nearby utility-sector employment is zero in event years -3 through -1 and positive in event years $+1$ through $+3$. Source: Census LODES.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012). "The Environment and Directed Technical Change". *American Economic Review* 102.1, pp. 131–166.

Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr (2016). "Transition to Clean Technology". *Journal of Political Economy* 124.1, pp. 52–104.

Aghion, Philippe, Antoine Dechezleprêtre, David Hémous, Ralf Martin, and John Van Reenen (2016). "Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry". *Journal of Political Economy* 124.1, pp. 1–51.

Altonji, Joseph G. and Lewis M. Segal (1996). "Small-Sample Bias in GMM Estimation of Covariance Structures". *Journal of Business & Economic Statistics* 14.3, pp. 353–366.

Andonov, Aleksandar and Joshua Rauh (2025). "The Shifting Finance of Electricity Generation". *Available at SSRN*.

Antweiler, Werner and Felix Muesgens (2021). "On the Long-Term Merit Order Effect of Renewable Energies". *Energy Economics* 99, p. 105275.

Bakke, Tor-Erik and Tiantian Gu (2017). "Diversification and Cash Dynamics". *Journal of Financial Economics* 123.3, pp. 580–601.

Becker, Randy and Vernon Henderson (2000). "Effects of Air Quality Regulations on Polluting Industries". *Journal of Political Economy* 108.2, pp. 379–421.

Blonz, Joshua A., Cooper Howes, and Joakim A. Weill (2026). "Money and Power". *Available at SSRN*.

Brealey, Richard A., Stewart C. Myers, Franklin Allen, and Alex Edmans (2022). *Principles of Corporate Finance*. 14th. McGraw-Hill Education.

Bushnell, James and Kevin Novan (2021). "Setting with the Sun: The Impacts of Renewable Energy on Conventional Generation". *Journal of the Association of Environmental and Resource Economists* 8.4, pp. 759–796.

Callaway, Brantly and Pedro HC Sant'Anna (2021). "Difference-in-Differences with Multiple Time Periods". *Journal of Econometrics* 225.2, pp. 200–230.

Callaway, Duncan S., Meredith Fowlie, and Gavin McCormick (2018). "Location, Location, Location: The Variable Value of Renewable Energy and Demand-Side Efficiency Resources". *Journal of the Association of Environmental and Resource Economists* 5.1, pp. 39–75.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik (2014). "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs". *Econometrica* 82.6, pp. 2295–2326.

Darmouni, Olivier, Clemens Lehner, and Parinitha Sastry (2026). "Financing Investment in Electricity". *Available at SSRN*.

Darmouni, Olivier and Yuqi Zhang (2024). "Brown Capital (Re)Allocation". *Available at SSRN*.

Davis, Lucas W. (2011). "The Effect of Power Plants on Local Housing Values and Rents". *Review of Economics and Statistics* 93.4, pp. 1391–1402.

Davis, Lucas W., Catherine Hausman, and Nancy L. Rose (2023). "Transmission Impossible? Prospects for Decarbonizing the US Grid". *Journal of Economic Perspectives* 37.4, pp. 155–180.

de Chaisemartin, Clément and Xavier D'Haultfoeuille (2024). "Difference-in-Differences Estimators of Intertemporal Treatment Effects". *Review of Economics and Statistics*, pp. 1–45.

Duffie, Darrell and Kenneth J. Singleton (1993). "Simulated Moments Estimation of Markov Models of Asset Prices". *Econometrica* 61.4, pp. 929–952.

Fabra, Natalia, Eduardo Gutiérrez, Aitor Lacuesta, and Roberto Ramos (2024). "Do Renewable Energy Investments Create Local Jobs?" *Journal of Public Economics* 239, p. 105212.

Fell, Harrison and Daniel T. Kaffine (2018). "The Fall of Coal: Joint Impacts of Fuel Prices and Renewables on Generation and Emissions". *American Economic Journal: Economic Policy* 10.2, pp. 90–116.

Fell, Harrison, Daniel T. Kaffine, and Kevin Novan (2021). "Emissions, Transmission, and the Environmental Value of Renewable Energy". *American Economic Journal: Economic Policy* 13.2, pp. 241–272.

Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur (2012). "What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NOx Trading Program". *American Economic Review* 102.2, pp. 965–993.

Gelman, Andrew and Guido Imbens (2019). "Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs". *Journal of Business & Economic Statistics* 37.3, pp. 447–456.

Gong, Haozhou, Chen Lin, Zacharias Sautner, and Thomas Schmid (2025). "Biodiversity Impacts of Renewable Energy". *Available at SSRN*.

Gowrisankaran, Gautam, Ashley Langer, and Wendan Zhang (2025). "Policy Uncertainty in the Market for Coal Electricity: The Case of Air Toxics Standards". *Journal of Political Economy* 133.6, pp. 1757–1795.

Gowrisankaran, Gautam, Stanley S. Reynolds, and Mario Samano (2016). "Intermittency and the Value of Renewable Energy". *Journal of Political Economy* 124.4, pp. 1187–1234.

Greenstone, Michael (2002). "The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufactures". *Journal of Political Economy* 110.6, pp. 1175–1219.

Hamilton, James (2011). "Careers in Solar Power". *U.S. Bureau of Labor Statistics*.

Hansen, Lars Peter (1982). "Large Sample Properties of Generalized Method of Moments Estimators". *Econometrica*, pp. 1029–1054.

Hansen, Lars Peter, John Heaton, and Amir Yaron (1996). "Finite-Sample Properties of Some Alternative GMM Estimators". *Journal of Business & Economic Statistics* 14.3, pp. 262–280.

Hausman, Catherine (2025). "Power Flows: Transmission Lines, Allocative Efficiency, and Corporate Profits". *American Economic Review* 115.8, pp. 2574–2615.

Hennessy, Christopher A. and Toni M. Whited (2007). "How Costly Is External Financing? Evidence from a Structural Estimation". *The Journal of Finance* 62.4, pp. 1705–1745.

Isen, Adam, Maya Rossin-Slater, and W. Reed Walker (2017). "Every Breath You Take—Every Dollar You'll Make: The Long-Term Consequences of the Clean Air Act of 1970". *Journal of Political Economy* 125.3, pp. 848–902.

Jarvis, Stephen (2025). "The Economic Costs of NIMBYism: Evidence from Renewable Energy Projects". *Journal of the Association of Environmental and Resource Economists* 12.4, pp. 983–1022.

Jha, Akshaya and Gordon Leslie (2025). "Start-Up Costs and Market Power: Lessons from the Renewable Energy Transition". *American Economic Review* 115.2, pp. 690–724.

Kumar, Mayank (2023). "Dirty Assets, Clean Opportunities: Private Equity Ownership of Fossil Fuels and the Energy Transition". *Available at SSRN*.

Kumar, Mayank (2025). "Private Equity and Creative Destruction: Evidence from the U.S. Power Sector". *Available at SSRN*.

Kumar, Mayank and Amiyatosh Purnanandam (2023). "Carbon Emissions and Shareholder Value: Causal Evidence from the US Power Utilities". *Available at SSRN*.

Lebakula, Viswadeep, Kelly Sims, Andrew Reith, Amy Rose, Jake McKee, Phil Coleman, Jason Kaufman, Marie Urban, Chris Jochem, and Carrie Whitlock (2025). "Landscan Global 30 Arcsecond Annual Global Gridded Population Datasets from 2000 to 2022". *Scientific Data, Nature* 12.1, p. 495.

Lin, Chen, Thomas Schmid, and Michael S. Weisbach (2026). "Climate Change, Demand Uncertainty, and Firms' Investments: Evidence from Planned Power Plants". *The Journal of Finance*.

Linn, Joshua and Kristen McCormack (2019). "The Roles of Energy Markets and Environmental Regulation in Reducing Coal-Fired Plant Profits and Electricity Sector Emissions". *The RAND Journal of Economics* 50.4, pp. 733–767.

McFadden, Daniel (1989). "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration". *Econometrica*, pp. 995–1026.

Pakes, Ariel and David Pollard (1989). "Simulation and the Asymptotics of Optimization Estimators". *Econometrica*, pp. 1027–1057.

Peura, Heikki and Derek W. Bunn (2021). "Renewable Power and Electricity Prices: The Impact of Forward Markets". *Management Science* 67.8, pp. 4772–4788.

Ratcliffe, Michael, Charlynn Burd, Kelly Holder, and Alison Fields (2016). "Defining Rural at the US Census Bureau". *American Community Survey and Geography Brief* 1.8, pp. 1–8.

Rauh, Joshua and Brandon Garcia (2026). "The Energy Balance Sheet". *Substack (Public Notes for Stanford GSB Finance 338)*.

Rust, John (1994). "Structural Estimation of Markov Decision Processes". *Handbook of Econometrics* 4, pp. 3081–3143.

Shapiro, Joseph S. and W. Reed Walker (2018). "Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade". *American Economic Review* 108.12, pp. 3814–3854.

Tassi, Francesco and Noah Kittner (2024). "Repurposing Coal Plants—Regional Economic Impacts from Low Carbon Generation". *Renewable and Sustainable Energy Reviews* 199, p. 114467.

Taylor, Lucian A. (2010). "Why Are CEOs Rarely Fired? Evidence from Structural Estimation". *The Journal of Finance* 65.6, pp. 2051–2087.

Taylor, Lucian A. (2013). "CEO Wage Dynamics: Estimates from a Learning Model". *Journal of Financial Economics* 108.1, pp. 79–98.

Walker, W. Reed (2013). "The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce". *The Quarterly Journal of Economics* 128.4, pp. 1787–1835.

Wu, Yufeng (2018). "What's Behind Smooth Dividends? Evidence from Structural Estimation". *The Review of Financial Studies* 31.10, pp. 3979–4016.

Zhang, Yuqi (2025). "Clean Energy Subsidies and Capital Misallocation". *Working Paper*.

Internet Appendix for
Burnout to Buildout: Coal Exit and the Geography of Transition Costs

Internet Appendix – Sections

Internet Appendix – Tables

Internet Appendix – Figures

B.1 Timing Sensitivity in Measuring the Coal–Gas Fuel Cost Gap

This section examines whether the estimates in Table 4 depend on when the time-invariant coal–gas fuel cost gap is measured. In the main specification, $FuelCostGap_i^{Coal\text{-}Gas}$ is measured in the year following the introduction of the air pollution regulations. That timing is deliberate: event year 0 marks the proposal of the regulations, and the main measurement timing in event year 1 captures the firm's coal exit decision in a cost environment right after the policy introduction. A question is whether the cost gap, measured after policy treatment, could itself respond to the regulation, e.g., through the retirement of high-cost units, renegotiated fuel contracts, or a shift in the coal a firm burns. If so, the coal cost disadvantage would be partly an outcome rather than a pre-existing cost condition, and the interaction would be difficult to interpret cleanly.

I address this by reconstructing the gap at a sequence of event times and re-estimating Panels B and C at each. For event time τ, I take firm i's own coal fuel cost in calendar year $g_i + \tau$, where g_i is the firm's cohort year (the year of initial policy exposure), and subtract the cross-firm average gas fuel cost in that same calendar year, computed over all firms with gas generation. This is the construction used in the main specification, evaluated at τ rather than only at $\tau = +1$. Each series is then standardized over the coal-reliant sample, so the coefficient in every column is the effect of a one-standard-deviation increase in the gap as measured at that event time. Because the dispersion of the gap changes over the window, one standard deviation corresponds to a different dollar amount in each column.

Table B.1 reports the results, with the window from $\tau = -1$ to $\tau = +3$. The effect is present at every event time. Closure remains positive and statistically significant throughout, ranging from 4.684 at $\tau = -1$ to 7.505 at $\tau = +2$. Divestiture is statistically insignificant throughout, so the asymmetry between closure and sale in the main text does not depend on measurement timing. The portfolio results in Panel B hold in both shares and levels at every event time, with gas rising and coal falling in each case. The $\tau = -1$ column matters most for the concern above: measured before the regulations were proposed, the gap produces the same pattern as the main specification, which indicates the measure captures a persistent heterogeneity in firms' coal cost position rather than a consequence of the regulation.

Table B.1: Measurement Timing of the Coal–Gas Fuel Cost Gap

This table re-estimates Panels B and C of Table 4 with $FuelCostGap_i^{Coal\text{-}Gas}$ measured at event time τ, where $\tau = t - g_i$ and g_i is the year firm i first became subject to the air pollution regulations. Each row is a separate estimation. The gap at τ is firm i's coal fuel cost in year $g_i + \tau$ minus the cross-firm average gas fuel cost in that year, standardized over the coal-reliant sample; one standard deviation equals \$17.65, \$20.27, \$21.13, \$24.26, and \$21.08 per MWh at $\tau = -1, 0, +1, +2, +3$ respectively. The row $\tau = +1$ is the main specification reported in Table 4. All regressions include firm and event-time fixed effects and the firm- and state-level controls used in Table 4. Standard errors are clustered at the firm level. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Panel A. Exit Outcomes

Gap measured at	$Closure_{i,t}$	$Divestiture_{i,t}$	$Abatement_{i,t}^{log}$	Obs.
	(1)	(2)	(3)	
$\tau = -1$	4.684***	1.222	1.670***	1,825
	(1.743)	(1.614)	(0.481)	
$\tau = 0$	5.820***	1.451	1.991***	1,822
	(1.987)	(1.734)	(0.569)	
$\tau = +1$	5.468***	1.368	1.922***	1,820
	(2.038)	(1.740)	(0.619)	
$\tau = +2$	7.505***	-0.314	3.011***	1,811
	(2.660)	(1.694)	(0.276)	
$\tau = +3$	6.412***	-1.999	2.348***	1,764
	(1.762)	(1.448)	(0.249)	

Panel B. Generation Portfolio

	$\text{PortfolioShare}_{i,t}^{(k)}$ (%)		$\text{Capacity}_{i,t}^{(k)}$ (GW)		
Gap measured at	$k = Gas$	$k = Coal$	$k = Gas$	$k = Coal$	Obs.
	(1)	(2)	(3)	(4)	
$\tau = -1$	3.531***	-2.961***	0.314***	-0.244***	1,825
	(1.062)	(0.930)	(0.109)	(0.080)	
$\tau = 0$	4.345***	-3.578***	0.342***	-0.270***	1,822
	(1.350)	(1.215)	(0.129)	(0.099)	
$\tau = +1$	5.199***	-4.873***	0.314**	-0.248**	1,820
	(1.525)	(1.366)	(0.130)	(0.099)	
$\tau = +2$	5.742***	-4.574***	0.595***	-0.423***	1,811
	(1.386)	(1.378)	(0.135)	(0.092)	
$\tau = +3$	7.194***	-6.463***	0.403***	-0.325***	1,764
	(1.338)	(1.318)	(0.098)	(0.078)	

B.2 Emission Reduction Compliance

Section 4.2 shows that coal-reliant firms leave coal through closure rather than sale, and that a larger coal cost disadvantage accelerates that exit. This section asks whether those exits produce measurable emission reductions, and how much of any reduction the coal-to-gas substitution accounts for. The main text carries the headline magnitudes; what follows reports the specification, both panels, and the decomposition of the CO_2 decline.

Since coal-reliant firms respond to federal air pollution regulations by closing or divesting coal assets, the key question is whether these exits translate into measurable emission reductions, which is the ultimate goal of regulatory intervention. I assess firm-level changes in sulfur dioxide (SO_2), nitrogen oxides (NO_x), and carbon dioxide (CO_2) emissions. The first two pollutants are directly targeted under CSAPR. While MATS focuses on mercury and acid gases, its overlapping compliance timeline and targets with CSAPR are expected to reduce these pollutants at the firm level through reductions in coal generation.

CO_2, while not directly regulated, serves as an indicator of whether the broader transition away from coal also mitigates greenhouse gas emissions (GHG). A decline in CO_2 can be decomposed into three main channels: (i) the indirect effect of the observed coal-to-gas substitution, as gas generation emits substantially less CO_2 per unit of output;[42] (ii) enhanced abatement efforts (e.g., post-combustion carbon capture); (iii) a broader shift toward renewable generation sources, which will be analyzed in subsequent sections.

Table B.2 reports the results. Panel A compares firms by pre-policy coal reliance. Coal-reliant firms show significantly larger reductions in SO_2 and NO_x, consistent with CSAPR's regulatory focus. SO_2 emissions fall by 8.7 thousand pounds, which exceeds the pre-policy mean of 7.39 and corresponds to close to 60% of the standard deviation. NO_x emissions decline by 2.0 thousand pounds, approximately 77% of the mean and 52% of a standard deviation. CO_2 falls by 0.63 thousand tons, close to 45% of its pre-regulation average.

Panel B isolates the coal-to-gas substitution effect by focusing only on firms that were coal-reliant before the policy. The analysis interacts the post-policy indicator with the standardized coal–gas fuel cost gap, capturing how differences in relative fuel prices translate into shifts in generation portfolios from coal to gas, as previously observed in Table 4 Panel C. Columns (1) and (2) show that a one–standard-deviation increase in the fuel cost gap reduces SO_2 and NO_x emissions by 7.8 and 1.2 thousand pounds, respectively, indicating that the coal-to-gas substitution played a significant role in reducing pollutants directly targeted by regulation.

Column (3) isolates the indirect effect of this substitution on CO_2 emissions (the first CO_2 reduction channel as mentioned above), showing that the shift toward gas also lowers carbon emissions. Comparing the coefficient magnitudes across panels clarifies the relative contribution of this mechanism. For SO_2 and NO_x, the estimated effects in Panel A correspond to 1.1 and 1.7 standard deviation changes in the coal–gas cost gap in Panel B, while for CO_2 the aggregate effect in Panel A equals roughly a 3.0 standard deviation change in the coal–gas gap. This

[42] Coal-fired generation emits more CO_2 per kWh than natural gas. In 2023, U.S. utility-scale plants averaged about 2.31 lb CO_2/kWh for coal versus 0.96 lb CO_2/kWh for natural gas. Source: U.S. Energy Information Administration (EIA), *How Much Carbon Dioxide Is Produced per Kilowatthour of U.S. Electricity Generation?*, https://www.eia.gov/tools/faqs/faq.php?id=74&t=11.

Table B.2: Pollution Compliance Outcomes Following Air Regulation

This table examines the changes in air pollution outcomes among electric firms in response to federal regulations targeting sulfur dioxide (SO_2) and nitrogen oxides (NO_x). Panel A compares pollution changes for firms with high ex-ante reliance on coal-fired generation. Panel B assesses whether compliance strength is moderated by the relative cost of coal versus gas. The dependent variables are SO_2 emissions (k lbs), NO_x emissions (k lbs), and CO_2 emissions (k tons). $PostPolicy_{i,t}^{Air}$ equals one if firm i's state has adopted the air pollution regulations (CSAPR or MATS) by year t. $CoalBased_i$ is an indicator for firms that operated coal-fired generation prior to regulation. $FuelCostGap_i^{\text{Coal-Gas}}$ is the standardized difference between coal and gas fuel costs measured in the year following regulation. All regressions include firm fixed effects (θ_i) and event-time fixed effects (λ_τ), where $\tau = t - g_i$ denotes the number of years since firm i first became subject to the regulations, with cohort g defined as the year of initial exposure. Firm-level controls include lagged values of *FirmSize*, *Leverage*, *ROA*, *Liquidity*, *Depreciation*, *Capacity*, and *FuelCost*. State-level controls are averaged at the firm level and include lagged *StateDemand*, *ElectricPrice*, and *CarveOut*. All variable definitions are in the Appendix. Standard errors clustered at the firm level. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Panel A. Compliance Outcomes for Coal-Reliant Firms

Dep. Var.:	SO_2 (k lbs)	NO_x (k lbs)	CO_2 (k tons)
	(1)	(2)	(3)
$\text{CoalBased}_i \times \text{PostPolicy}_{i,t}^{\text{Air}}$	-8.665***	-2.013***	-0.628***
	(1.836)	(0.380)	(0.219)
Firm FE	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes
Controls (Firm & State)	Yes	Yes	Yes
Sample (pre-policy defined)	All firms	All firms	All firms
Observations	2,158	2,158	2,158
Adjusted within R^2	0.104	0.088	0.233

Panel B. Substitution Effect: Fuel Relative Prices and Compliance Outcomes

Dep. Var.:	SO_2 (k lbs)	NO_x (k lbs)	CO_2 (k tons)
	(1)	(2)	(3)
$\text{PostPolicy}_{i,t}^{\text{Air}} \times \text{FuelCostGap}_i^{\text{Coal-Gas}}$	-7.792***	-1.219***	-0.211**
	(2.616)	(0.454)	(0.095)
Firm FE	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes
Controls (Firm & State)	Yes	Yes	Yes
Sample (pre-policy defined)	Coal-based	Coal-based	Coal-based
Observations	1,820	1,820	1,820
Adjusted within R^2	0.199	0.129	0.249

comparison indicates that the substitution effect explains only a limited share of the observed carbon reductions. The result implies that, while gas combustion is cleaner than coal, both remain fossil-based sources, thus deeper CO_2 reductions likely stem from additional abatement or a broader transition toward non-fossil generation.

B.3 Siting Continuity, Retirement-Linked Geography, and Nearby Employment

This section provides construction details and supporting evidence for the two project-level analyses in Section 6.2. I first describe how I measure proximity to the utility's incumbent fossil sites. I then document how coal retirements, workplace blocks, and residence weights are linked, report the sample construction sequence, and examine the sensitivity of the retirement-linked estimates. The final subsection examines all utility-sector employment near new project openings. The three analyses measure different margins: continuity with incumbent operating sites, geographic separation from the estimated residential distribution associated with retiring coal workplaces, and employment changes near the receiving site.

B.3.1 Continuity with Incumbent Fossil Sites

The site-continuity test asks whether newly developed generation locates at or near the utility's incumbent fossil infrastructure, where existing grid access may provide a siting advantage. Based on the EIA sample, the population test in Table 8 dates each utility–project event by the first year in which the project operates under the focal utility. This definition supports the site-remoteness analysis because it identifies the project coordinate associated with the utility, but it can include an existing plant entering the utility's operating portfolio, such as through an acquisition.

Therefore, for the site-continuity test, I require a physical opening at the reported site rather than only the start of operation under the focal utility. A project qualifies when a generator using the project's fuel begins operation at that EIA plant in the stated utility–project event year and the focal utility reports a positive ownership share. This restriction links 792 of the 1,362 projects in the EIA sample of Table 8 to an observable physical opening at the reported coordinate.

For each verified opening, I first use EIA Form 860 ownership records from event year -1, one year before project operation, to identify incumbent coal-, gas-, or oil-fired plants in the same state in which the focal utility owned operating capacity. I aggregate ownership across all operating fossil generators at each candidate plant. The plant qualifies when the focal utility's owned capacity accounts for at least 50% of the plant's physical operating fossil capacity. Measuring the incumbent portfolio in event year -1 prevents the new project itself from creating the comparison site. The same-state restriction keeps the measure focused on the utility's existing operating geography around the project rather than distant assets elsewhere in its portfolio.

The first specification yields 534 projects, comprising 354 solar and 180 gas- or oil-fired projects across 133 utilities. I then construct a coal-only specification to determine whether the result reflects only gas projects opening near incumbent gas plants. Again, a candidate site must contain operating coal capacity in event year -1, and the focal utility must own at least 50% of the plant's physical operating coal capacity. I condition on the project having at least one such same-utility coal site in the same state in that year, yielding 221 projects: 183 solar and 38 gas- or oil-fired projects across 52 utilities.

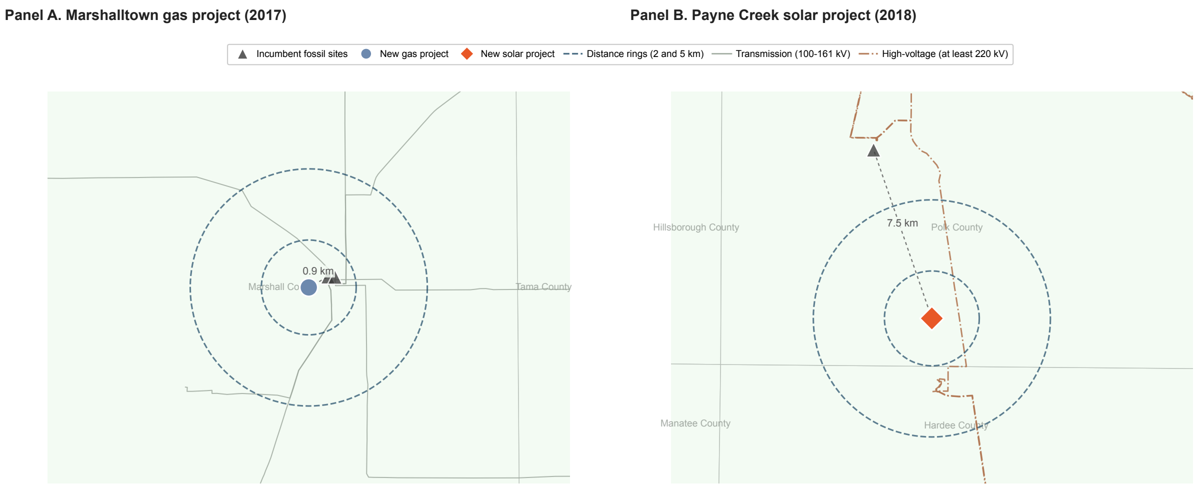

Figure B.1. Siting Continuity: Construction Examples. This figure illustrates the construction of $SiteContinuity_{i,p}^{(b)}$ using actual project openings. Dark triangles identify qualifying incumbent fossil sites in the project state, measured one year before project operation, at which the project utility owned at least 50% of aggregate physical operating fossil capacity. The blue circle identifies the new gas project, and the orange diamond identifies the new solar project. Dashed rings mark distances of 2 and 5 km from each new project. The dotted segment connects each project to its nearest qualifying incumbent site, with the distance reported beside the segment. Solid gray-green lines show mapped transmission corridors rated 100–161 kV, while brown dash-dotted lines show corridors rated at least 220 kV. Thin gray lines identify county boundaries. Panel A presents Interstate Power and Light's Marshalltown gas project, which opened in Iowa in 2017; the nearest qualifying incumbent fossil site, measured in 2016, is 0.93 km away. Panel B presents Tampa Electric's Payne Creek solar project, which opened in Florida in 2018; the nearest qualifying incumbent fossil site, the Polk gas plant measured in 2017, is 7.55 km away.

For $b \in \{2, 5\}$ km, $SiteContinuity_{i,p}^{(b)}$ equals one when at least one qualifying site from the relevant incumbent-site set lies within b km of project p. I estimate project-level linear probability models of $SiteContinuity_{i,p}^{(b)}$ on $InvestType_p$, operation-year fixed effects, and an error term, with standard errors clustered by utility. The measure captures geographic co-location or near-co-location with the utility's incumbent operating sites.

Figure B.1 makes the threshold construction visible at a common scale. The gas project in Panel A lies within both continuity radii, whereas the solar project in Panel B lies outside both. The cases are illustrations rather than a matched comparison, and the transmission corridors provide geographic context but do not enter $SiteContinuity_{i,p}^{(b)}$.

B.3.2 Retirement Links and Residence Weights

I begin with the EIA sample of 1,362 solar and gas- or oil-fired projects in Table 8. A project enters the retirement-linked sample when the same utility records a coal-generator retirement in the same state from five years before through three years after the project's operation year. The window does not require coal retirement to precede the new project: the five-year pre-operation side allows replacement investment to follow retirement with a lag, while its three-year post-operation side retains cases in which the utility commissions replacement capacity before formal closure, when planning for retraining or redeployment could also begin. Requiring retirement first would exclude these anticipatory transition cases.

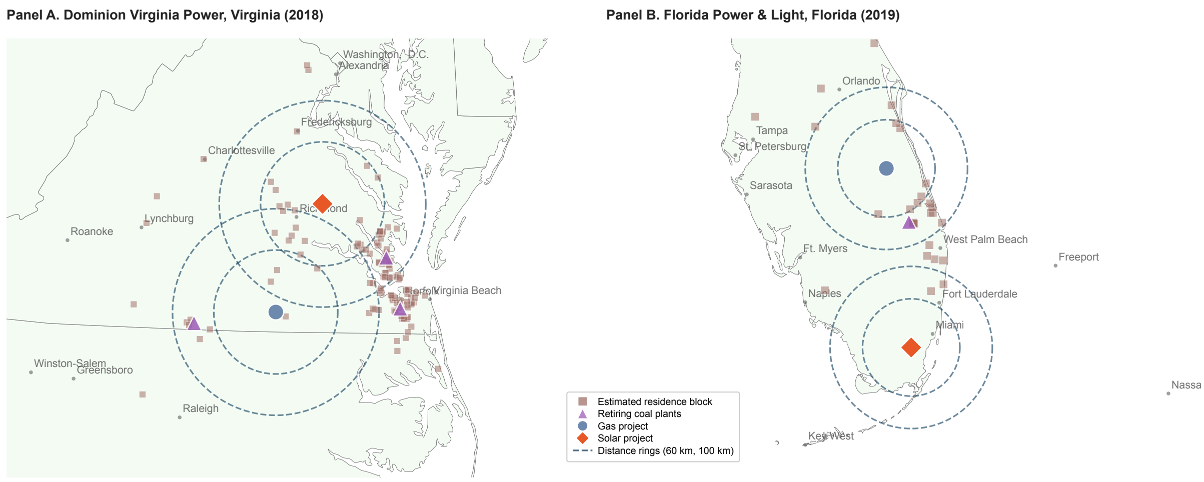

Figure B.2. Retirement-Linked Workforce Accessibility: Construction Examples. This figure illustrates the construction using actual utility-project cases. Squares mark the internal points of estimated residence blocks linked by LODES OD flows to validated utility-sector workplace blocks at the retiring coal plants. Larger squares indicate greater estimated utility-job weight. Purple triangles identify the EIA coordinates of the retiring coal plants, the blue circle identifies the gas project, and the orange diamond identifies the solar project. Dashed rings show distances of 60 and 100 km around each new project. Panel A presents Dominion Virginia Power projects opening in Virginia in 2018, the unique utility-state-operation-year cell in the main sample containing exactly one gas and one solar project. Both projects are evaluated against the same three retiring coal workplaces and the same residence weights. Panel B presents Florida Power & Light projects opening in Florida in 2019. The empirical measures use all qualifying residence blocks.

A retirement must follow an observed operating status, be followed by a terminal retired status, and have no subsequent operating observation. The project utility must own at least 50% of the linked retired capacity, and the linked retirements must remove at least 50% of the plant's coal capacity. These requirements identify a timing-consistent physical retirement rather than a decline in generation or capacity that may be temporary.

For each linked coal plant, the LODES workplace snapshot is the earlier of the year before project operation and the year before the first linked retirement. This ensures that workplace employment is measured before both the new project and the coal retirement. I first examine the Census block containing the coal-plant coordinate. A block qualifies when WAC reports positive NAICS 22 primary employment and utility-sector jobs account for at least one-half of employment in the corresponding trade, transportation, and utilities aggregation. If the coordinate block does not qualify, the baseline selects the qualifying block with the largest NAICS 22 job count within 2 km of the plant coordinate. The screen reduces the risk of assigning a nearby commercial or industrial block to the coal workplace.

Let J_q denote the WAC NAICS 22 job count in validated coal-workplace block q, and let $F_{h,q}$ denote the public OD primary-job flow from residence block h to workplace block q for the broader trade, transportation, and utilities category. The estimated utility-job weight assigned to residence block h is

$$w_{h,q} = J_q \frac{F_{h,q}}{\sum_h F_{h,q}},$$

where the denominator sums flows from all residence blocks connected to workplace block q. Thus, WAC determines the utility-job total at the workplace, while OD determines how that total is distributed across residence blocks. For example, if WAC reports 100 utility-sector jobs and a residence block supplies 20% of the broader OD flow to that workplace, the residence block receives a weight of 20. The identifying assumption is that utility-sector workers in a validated workplace block have the same residential distribution as workers in the broader OD category. I retain both within-state and cross-state residence flows.

When several retiring coal plants qualify for the same utility-project event, I pool their weighted residence blocks. A workplace block shared by colocated coal plants is counted once. Let $d(a, c)$ denote the geographic distance between locations a and c over the earth's surface, and let $q(h)$ denote the coal-workplace block associated with a residence-flow observation. The three outcomes are

$$\begin{aligned}
\text{ResidenceShare}^{(b)}_{i,p} &= 100 \times \frac{\sum_{h,q} w_{h,q} \mathbb{1}\{d(h,p) \leq b\}}{\sum_{h,q} w_{h,q}}, \\
\text{WorkforceDistance}_{i,p} &= \frac{\sum_{h,q} w_{h,q} d(h,p)}{\sum_{h,q} w_{h,q}}, \\
\text{AddedDistance}_{i,p} &= \frac{\sum_{h,q} w_{h,q} \left[d(h,p) - d(h,q(h))\right]}{\sum_{h,q} w_{h,q}}.
\end{aligned}$$

The first outcome is a weighted residence share, the second is distance from the estimated residential distribution to the new project, and the third measures the hypothetical increase relative to the corresponding coal workplace.

B.3.3 Sample Construction and Robustness

Table B.3 reports the sample sequence for both panels of Table 9. Panel A requires an observable project opening and at least one incumbent fossil site against which proximity can be measured; its coal-only columns further require an incumbent coal site. Panel B requires a same-utility same-state coal retirement, a complete event time $[-5, +3]$ retirement-link window to allow workforce redeployment before formal closure, and a validated LODES workplace and residence distribution.

As robustness checks, Table B.4 separates additional outcomes from changes in the construction. Panel A reports the 60 km residence share and weighted-median versions of the workforce- and added-distance measures. Panels B and C focus on $AddedDistance_{i,p}$ because it directly benchmarks the new project against the linked coal workplace while keeping the estimated residential distribution fixed.

The sequence clarifies why the retirement-linked sample is much smaller than the broad remoteness sample. Table 8 requires an observed EIA sample utility–project event and its site coordinate. The retirement-linked analysis additionally requires a qualifying same-utility same-state coal retirement in the event window and a validated LODES residence distribution. This procedure isolates the 110 projects for which within-utility workforce accessibility can be measured at this level of geographic detail. Again, it does not imply that the remaining projects are unrelated to utilities' broader coal transition.

Table B.4 examines whether the results are sensitive to distant residence blocks, workplace assignment, retirement timing, or the composition of the comparison group. Column (1) of

Table B.3: Sample Construction: Siting Continuity and Retirement-Linked Accessibility

This table reports the sequential construction of the samples used in Table 9. The starting sample contains all solar and gas- or oil-fired projects in the EIA sample of Table 8. In Panel A, a verified “project opening” requires a same-fuel generator owned by the focal utility to begin operation at the reported project plant in the stated year. A qualifying incumbent fossil site is a same-state plant in which the utility owned at least 50% of physical operating fossil capacity in event year −1. The “coal-only” specifications condition on having at least one same-utility coal site in the project state in event year −1 for which the utility owned at least 50% of physical operating coal capacity. In Panel B, a qualifying retirement is a same-utility same-state physical coal retirement in event years −5 through +3 that satisfies the ownership and retired-capacity requirements described in this section. The next stage requires this complete event window to be observable through 2023. The final stage requires a validated LODES workplace at or within 2 km of at least one linked coal plant and a corresponding OD residence distribution. “Retained” is the percentage of projects remaining from the preceding row.

Panel A. Continuity with Incumbent Fossil Sites

Stage	Solar	Gas/Oil	Total	Retained (%)
EIA sample in population tests	664	698	1,362	–
Project opening verified at focal EIA site	534	258	792	58.1
At least one qualifying incumbent fossil site	354	180	534	67.4
At least one qualifying incumbent coal site	183	38	221	41.4

Panel B. Retirement-Linked Workforce Accessibility

Stage	Solar	Gas/Oil	Total	Retained (%)
EIA sample in population tests	664	698	1,362	–
Qualifying same-utility, same-state retirement	159	52	211	15.5
Complete $[-5, +3]$ retirement-link window	94	39	133	63.0
At least one workplace measured within 2 km	84	26	110	82.7

Panel A adds the residence share within 60 km, which is 17.0 percentage points lower for solar projects than for gas- or oil-fired projects. Columns (2) and (3) replace the weighted-mean distance measures with weighted medians. The weighted-median residence-to-project distance is 82.1 km greater for solar projects, and the weighted-median added distance is 76.3 km greater.

Panel B varies how the workplace block is assigned to each retiring coal plant. Column (1) requires the Census block containing the exact coal-plant coordinate itself to qualify. Column (2) applies the baseline rule: it uses the coordinate block when eligible, and otherwise selects the qualifying block with the most utility-sector jobs within 2 km, retaining a project when at least one linked coal plant is matched. Column (3) directly applies the same 2 km search, but requires every linked coal plant to be matched. Column (4) retains valid matches within 2 km and extends the search to 5 km only for otherwise unmatched plants. Across these definitions, the estimated difference in added distance between solar and gas- or oil-fired projects ranges from 58.1 to 75.0 km, and is statistically significant at the 1% level in every specification.

Panel C shows a 63.2 km difference under the $[-2, +2]$ retirement window, a 70.6 km difference under the $[-3, +3]$ window, and a 66.1 km difference when oil-fired projects are excluded from the control group. All three estimates are statistically significant at the 1% level.

Table B.4: Robustness: Retirement-Linked Workforce Accessibility

This table reports solar-minus-gas/oil coefficients from project-level cross-sectional regressions using variations of the Table 9 specifications as robustness checks. Panel A reports the 60 km residence share in the first column and weighted-median versions of workforce distance and added distance in the final two columns. Panel B varies workplace measurement. "Exact block" requires the coal-coordinate block itself to qualify. "2 km: any" requires a valid workplace for at least one linked coal plant. "2 km: all" requires a valid workplace for every linked coal plant. "2–5 km: any" preserves valid 2 km matches and searches from 2 to 5 km only for otherwise unmatched plants. Panel C narrows the retirement window or excludes the three oil-fired control projects. Each timing specification requires its complete window to be observable through 2023. A narrower window can contain fewer projects because it also requires a qualifying retirement to occur within that narrower interval. All specifications include operation-year fixed effects and cluster standard errors by firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Panel A. Additional Outcomes Under the Main Construction

Dep. Var.:	$\text{ResidenceShare}_{i,p}^{(b)}$ $b = 60$ km (p.p.)	$\text{WorkforceDistance}_{i,p}$ Median (km)	$\text{AddedDistance}_{i,p}$ Median (km)
	(1)	(2)	(3)
InvestType_p	-17.0*	82.1***	76.3***
	(9.7)	(22.2)	(23.4)
Operation-year FE	Yes	Yes	Yes
Observations	110	110	110
Adjusted within R^2	0.047	0.064	0.051

Panel B. Workplace-Measurement Robustness

Dep. Var.: $\text{AddedDistance}_{i,p}$	Exact block	2 km: any	2 km: all	2–5 km: any
	(1)	(2)	(3)	(4)
InvestType_p	58.135***	66.173***	75.017***	68.023***
	(18.410)	(19.028)	(23.213)	(14.613)
Operation-year FE	Yes	Yes	Yes	Yes
Observations	88	110	95	121
Adjusted within R^2	0.023	0.054	0.073	0.073

Panel C. Retirement Window and Comparison Group

Dep. Var.: $\text{AddedDistance}_{i,p}$	$[-2, +2]$	$[-3, +3]$	Gas-only control
	(1)	(2)	(3)
InvestType_p	63.208***	70.612***	66.083***
	(21.120)	(20.400)	(19.650)
Operation-year FE	Yes	Yes	Yes
Observations	109	103	107
Adjusted within R^2	0.013	0.057	0.043

B.3.4 Nearby Utility-Sector Employment Around Project Openings

One remaining question is whether greater workplace remoteness should be interpreted as a "social cost": if remotely sited solar projects bring greater employment growth to their surrounding communities, this local benefit could offset the adjustment cost of remoteness. I test this possibility by comparing nearby utility-sector employment growth after solar openings with that after gas openings, the investment path associated with continued fossil dependence. If the proposed offsetting benefit is present, employment should increase more around solar projects than around gas projects.

Table B.5: Nearby Utility-Sector Employment Around New Project Openings
This table compares LODES WAC NAICS 22 utility-sector primary employment around physical solar and gas project openings. WAC counts all utility-sector jobs in the surrounding Census blocks but does not identify the named project's employees. $InvestType_p$ equals one for solar and zero for gas. In Panel A, the dependent variable is the number of utility-sector jobs within 5 km of project p in event year τ. Each project contributes event years -3, -2, -1, $+1$, $+2$, and $+3$; event year zero is excluded. Thus, the coefficient on $InvestType_p \times Post_{p,\tau}$ is the solar-minus-gas difference in the change from the three pre-opening years to the three post-opening years. Column (1) includes project and event-time fixed effects, with no capacity adjustment. Column (2) uses the same 528 projects and additionally interacts log project capacity with $Post_{p,\tau}$, allowing the post-opening change to vary with project size. Column (3) applies the column (2) specification to the 411 projects between 1.185 and 90.115 MW, the overlap of the solar and gas 5th–95th percentile capacity ranges. Panel B uses one observation per project and restricts the sample to the 527 projects of at least 1 MW. In column (1), $UtilityGain_p^{(2)}$ equals one when average utility-sector employment within 2 km is higher in event years +1 through +3 than in event years -3 through -1. In column (2), $UtilityGain_p^{(5)}$ records the same increase within 5 km. In column (3), $UtilityEntry_p^{(5)}$ equals one when average utility-sector employment within 5 km is zero before the opening and positive afterward. The coefficient on $InvestType_p$ in each column is therefore the solar-minus-gas difference in the probability of that outcome. All three columns control for log project capacity and include operation-year fixed effects. Standard errors in both panels are clustered by connected 10 km site groups. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Panel A. Average Utility-Sector Employment Within 5 km

Dep. Var.: $\text{UtilityJobs}_{p,\tau}^{(5)}$	Baseline	Capacity adjusted	Common capacity range
	(1)	(2)	(3)
$\text{InvestType}_p \times \text{Post}_{p,\tau}$	-11.259	-12.008	-18.216
	(11.218)	(13.920)	(18.492)
Project FE	Yes	Yes	Yes
Event-time FE	Yes	Yes	Yes
Capacity × Post	No	Yes	Yes
Project-year observations	3,168	3,168	2,466
Adjusted within R^2	0.001	0.001	0.003

Panel B. Extensive-Margin Employment Outcomes

Dep. Var.:	$\text{UtilityGain}_p^{(b)}$		$\text{UtilityEntry}_p^{(b)}$
	$b = 2$ km	$b = 5$ km	$b = 5$ km
	(1)	(2)	(3)
InvestType_p	-0.209***	-0.128**	0.076***
	(0.054)	(0.053)	(0.025)
Capacity control	Yes	Yes	Yes
Operation-year FE	Yes	Yes	Yes
Adjusted within R^2	0.057	0.010	0.021

This test requires the physical commissioning year because employment must be measured around the date when generation begins at the site. I collapse jointly owned records into one physical project and require LODES WAC observations in every year from three years before through three years after commissioning. This yields 528 projects opening from 2010 through 2020: 390 solar and 138 gas projects. For each project and year, I count all WAC NAICS 22 utility-sector primary jobs in a fixed set of Census blocks whose internal points lie within 5 km of the project coordinate, always including the coordinate-containing block.

Table B.5 examines two related margins. Panel A compares utility-sector employment in event years -3 through -1 with employment in years $+1$ through $+3$. I exclude the commissioning year so that a partial opening year is not assigned to either period. Project fixed effects absorb permanent differences across sites, while event-time fixed effects absorb the average employment pattern around commissioning. Panel B instead asks whether nearby employment increases at all and whether utility-sector employment first appears in an area where none was detected before the opening. This distinction allows a small entry into a previously empty area to be separated from a broader local employment gain. Because nearby project catchments can overlap, standard errors are clustered across projects connected within 10 km.

Panel A provides no evidence that solar openings generate larger nearby utility-sector employment gains than gas openings. In column (1), the estimated solar-minus-gas difference in the pre-to-post change is -11.3 jobs within 5 km. The difference is -12.0 jobs after allowing the change to vary with project capacity in column (2), and -18.2 jobs among projects in the common capacity range in column (3). The negative point estimates run against a solar employment advantage, but they are not statistically precise enough to establish that gas openings increase average job counts by more.

Panel B shows a clearer difference on the extensive margin. Relative to gas openings, solar openings are 20.9 percentage points less likely to be associated with any increase in utility-sector employment within 2 km and 12.8 percentage points less likely within 5 km. At the same time, solar openings are 7.6 percentage points more likely to mark a change from zero to positive utility-sector employment within 5 km. This first-entry result is expected under remoteness and should not be interpreted as evidence of a solar employment advantage. Consistent with the population evidence, 45.2% of solar catchments have no utility-sector employment within 5 km before opening, compared with 15.9% of gas catchments. Solar projects have more scope to register a first entry, whereas gas projects tend to open nearer existing population and employment. Thus, the first-entry advantage is consistent with the lower probability of an overall employment gain: solar projects more often start from zero, but are less likely than gas projects to experience any increase in nearby utility-sector jobs. Therefore, the evidence does not indicate that the potential worker-adjustment costs associated with solar projects' greater siting remoteness are offset by stronger nearby utility-sector employment growth.

B.4 Model Estimation: Moments, Fit, and Robustness

This Internet Appendix section provides additional information on moment selection, moment fit, and estimation diagnostics for Section 7.

B.4.1 Moment Selection

The moments summarize the two economic margins in the model: (i) *within a state–year*, how firms make decisions when facing the same external environments; and (ii) *across state–years*, how adoption co–moves with policy and market conditions as the breakeven shifts. This design ensures that moments are transparent about which forces they discipline and isolate parameters through near one–to–one moment-parameter mappings. For the 12 estimated parameters in vector θ, I use $q = 15$ moments: 12 moments for the relationships between adoption and firm traits, site characteristics, and policy and market conditions, and 3 moments for the distribution of renewable project adoption. The overall mean adoption rate is matched by construction through the intercept κ_0, and serves as a diagnostic rather than an objective moment.

Given $\mathcal{D}_{i,t} \in \{0,1\}$ is the adoption, let $\bar{\mathcal{D}}_{s,t}$ be its state–year mean. For any firm–level variable $y_{i,t}$, define the within demeaning transformation:

$$\widetilde{y}_{i,t} = y_{i,t} - \overline{y}_{s,t}, \quad \text{where } \overline{y}_{s,t} \text{ averages } y \text{ over firms observed in } (s,t).$$

Within–State–Year Partial Covariances with Firm Traits. First, to measure how firm-level generation capacity traits affect the decision index in equation (13) with the state–year environment fixed, I use within–state–year partial covariances that rule out collinearity among traits. Let trait vector $T_{i,t} = \big(\text{coal}_{i,t},\ \text{gas}_{i,t},\ \text{solar}_{i,t},\ \text{capTot}_{i,t},\ \text{capProj}_{i,t}\big)$. For trait variable j, I run an OLS of $\widetilde{T}_{j,i,t}$ on a constant and the other four within–demeaned traits $\{\widetilde{T}_{\ell,i,t} : \ell \neq j\}$; denote the residual by $T^{\perp}_{j,i,t}$. The moment is:

$$\text{M1–M5:} \quad \text{pCov}_{\text{within}}\big(\mathcal{D}, T_j\big) = \frac{1}{N} \sum_{i,t} \widetilde{\mathcal{D}}_{i,t}\, T^{\perp}_{j,i,t}, \quad j \in \{\text{coal}, \text{gas}, \text{solar}, \text{capTot}, \text{capProj}\}, \tag{B.1}$$

which primarily maps to the slope coefficients $(\beta_{\text{coal}}, \beta_{\text{gas}}, \beta_{\text{solar}}, \beta_{\text{capT}}, \beta_{\text{capP}})$ in equation (13).

Within–State–Year Remoteness Covariance. The remoteness site trait enters the decision index in equation (13) through β_{rem}, and its moment mirrors the trait construction. I residualize the within-demeaned standardized remoteness $\widetilde{\text{Rem}}_{i,t}$ on a constant, the five within-demeaned traits $\{\widetilde{T}_{j,i,t}\}$, and the within-demeaned site quality $\widetilde{\widehat{\phi}^{R}}_{i,t}$; denote the residual by $\text{Rem}^{\perp}_{i,t}$. The moment is:

$$\text{M6:} \quad \text{pCov}_{\text{within}}\big(\mathcal{D}, \text{Rem}\big) = \frac{1}{N} \sum_{i,t} \widetilde{\mathcal{D}}_{i,t}\, \text{Rem}^{\perp}_{i,t}, \tag{B.2}$$

which maps to β_{rem}. Partialling on site quality ensures that the moment reflects remoteness variation beyond what irradiance already captures.

Across-State–Year Policy Sensitivity. For the state–year level policy variables, the moments summarize their partial covariances with the state–year mean adoption. For variable k in vector $X_{s,t} = \big(\text{Carbon}_{s,t},\ \text{Subsidy}^{R}_{s,t},\ \text{RPS}_{s,t}\big)$, let $X^{\perp}_{k,s,t}$ be the residual from an OLS of $X_{k,s,t}$ on a constant, $P^e_{s,t}$, $\Delta C_{s,t}$, and the other two policy variables $\{X_{\ell,s,t} : \ell \neq k\}$. Again, $\widetilde{y}_{s,t} = y_{s,t} - \overline{y}$ denote the demeaned value for the state-year variables, where $\overline{y}$ is the average of $y_{s,t}$ across all (s,t) in the sample. For variable k, the moment is:

$$\text{M7–M9:}\quad \text{pCov}\big(\overline{\mathcal{D}}, X_k\big) = \frac{1}{\mathcal{S}} \sum_{(s,t)\in\mathcal{S}} \widetilde{\mathcal{D}}_{s,t}\, X^{\perp}_{k,s,t}, \quad k \in \{\text{Carbon},\ \text{Subsidy}^R,\ \text{RPS}\}, \tag{B.3}$$

which map to λ_{carb}, ψ, and η_{rps}, respectively. RPS enters the decision index linearly (see Section 7), so its level moment is sufficient.

Dispersion and Site–Quality Sorting. The following moments discipline the two dispersion parameters (σ and σ_{env}) and the strength of geographic sorting. M10 is the average within–state–year variance of adoption, M11 is the variance of state–year adoption means (between–state–years), and M12 is the correlation between adoption and the site–quality proxy after demeaning both by their state–year means:

$$\text{M10:}\quad \text{Var}_{\text{within}}[\mathcal{D}] \;=\; \frac{1}{\mathcal{S}} \sum_{(s,t)\in\mathcal{S}} \text{Var}\big(\mathcal{D}_{i,t} \mid s,t\big), \tag{B.4}$$

$$\text{M11:}\quad \text{Var}_{\text{between}}\big[\bar{\mathcal{D}}_{s,t}\big] \;=\; \text{Var}\big(\bar{\mathcal{D}}_{s,t}\big), \tag{B.5}$$

$$\text{M12:}\quad \text{Corr}_{\text{within}}\big(\mathcal{D}, \widehat{\phi}^R\big) \;=\; \text{Corr}\big(\widetilde{\mathcal{D}}_{i,t},\ \widetilde{\widehat{\phi}^R_{i,t}}\big). \tag{B.6}$$

Distributional Moments for Adoption. These moments summarize the adoption by (i) the share of state–years with adoption below 0.10 and above 0.90, and (ii) the mean adoption among firms in the top tercile of the resource–quality proxy $\widehat{\phi}^R$.

$$\text{M13:}\quad \Pr\big(\overline{\mathcal{D}}_{s,t} < 0.10\big) = \frac{1}{\mathcal{S}} \sum_{(s,t)\in\mathcal{S}} \mathbb{1}\{\overline{\mathcal{D}}_{s,t} < 0.10\}, \tag{B.7}$$

$$\text{M14:}\quad \Pr\big(\overline{\mathcal{D}}_{s,t} > 0.90\big) = \frac{1}{\mathcal{S}} \sum_{(s,t)\in\mathcal{S}} \mathbb{1}\{\overline{\mathcal{D}}_{s,t} > 0.90\}, \tag{B.8}$$

$$\text{M15:}\quad \mathbb{E}\big[\mathcal{D}_{i,t} \mid \widehat{\phi}^R_{i,t} \text{ in top tercile}\big] = \frac{\sum_{i,t} \mathbb{1}\{\widehat{\phi}^R_{i,t} \text{ in top tercile}\}\, \mathcal{D}_{i,t}}{\sum_{i,t} \mathbb{1}\{\widehat{\phi}^R_{i,t} \text{ in top tercile}\}}. \tag{B.9}$$

B.4.2 Moment Fit

Test of Overidentifying Restrictions. The estimator uses 15 moments to estimate 12 parameters, leaving 3 overidentifying restrictions. Because there are more moments than parameters, the 12 parameters must jointly fit all 15 moments and generally cannot eliminate every model–data discrepancy. The overidentification test asks whether the discrepancies that remain after estimation are consistent with sampling variation or collectively too large to be consistent with the model.

Let the vector of moment discrepancies at the estimate be:

$$g(\widehat{\theta}) = m^{\text{sim}}(\widehat{\theta}) - m^{\text{emp}}, \tag{B.10}$$

where m^{emp} contains the empirical moments and $m^{\text{sim}}(\widehat{\theta})$ their model-implied counterparts. Because the moments include covariances, variances, correlations, and shares, their differences are measured in different units. I make them comparable by dividing each difference by its standard error s_j, estimated from 1,000 bootstrap samples that resample the 45 states with replacement. With only 45 states, estimates of all pairwise correlations among the 15 moments are noisy, and inverting the resulting covariance matrix can produce unstable weights (see Altonji and Segal, 1996; Hansen et al., 1996). Thus, I use a diagonal weighting matrix based on the 15 estimated variances, excluding the estimated correlations when choosing the parameters:

$$W = \operatorname{diag}\bigl(1/s_1^2, \ldots, 1/s_{15}^2\bigr), \qquad Q(\widehat{\theta}) = g(\widehat{\theta})^\top W g(\widehat{\theta}) = \sum_{j=1}^{15} \left(\frac{g_j}{s_j}\right)^2. \tag{B.11}$$

Under correct specification, the diagonal-weighted estimator remains consistent: it converges to the true parameter values as the sample grows. It can, however, be less statistically efficient than well-estimated full-covariance weighting and yield different estimates in a sample of this size. The reported sandwich standard errors still use the full bootstrap covariance $\widehat{\Omega}$, including the estimated correlations among moments.

Because diagonal weighting does not account for correlations among the moments, Q does not follow the usual $\chi^2(3)$ distribution. Thus, I simulate the distribution benchmark directly. In each of 400,000 simulations, I generate hypothetical moment discrepancies using their full bootstrap covariance, approximate how the parameter estimates would adjust, and calculate the resulting Q to form its distribution. These simulated values show the range of Q expected from sampling variation when the model is correctly specified.

The observed value is $Q(\widehat{\theta}) = 0.84$. The simulated 95 percent critical value is 3.53, and 49 percent of the simulated values are at least as large as the observed Q (simulation-based $p = 0.49$). Therefore, the overidentification test does not reject the model. Still, this calculation is a large-sample approximation: it may be less accurate here because $\widehat{\sigma} = 0.005$ is nearly equal to its imposed lower bound, and I treat the result as a specification diagnostic rather than proof that the model is absolutely correct.

No model moment in Table B.6 differs from its empirical counterpart by even one bootstrap standard error. The largest difference is 0.63 standard errors: the within-state–year correlation between adoption and site quality is 0.318 in the model and 0.408 in the data. The model also overstates the partial covariance between state–year adoption and RPS (0.261 versus 0.157), but the difference is only 0.29 standard errors.

M12 contributes to identification of the site-quality coefficient by comparing projects within the same state–year. The within-state–year variation available for this comparison is limited: it accounts for about 10 percent of the total variance in $\widehat{\phi}^R$ (standard deviation 0.011 within state–years, compared with 0.037 in the full sample). Limited within-state–year variation does not itself explain the model–data discrepancy, so I interpret both this discrepancy and the estimate of α_ϕ cautiously.

Table B.6: Structural SMM Moment Fit

This table compares the 15 empirical moments used in Table 10 with their model-implied counterparts at the estimated parameters. “Model” moments are averaged over $B = 200$ simulated state–year environments, using the same draws at every parameter value. The overall mean adoption rate is matched by construction through κ_0 and is not one of the 15 moments used to estimate the parameters.

Moment (notation)	Empirical	Model
M1: Within partial covariance with coal share $\left(\text{pCov}_{\text{within}}(\mathcal{D}, \text{coal})\right)$	-0.017	-0.015
M2: Within partial covariance with gas share $\left(\text{pCov}_{\text{within}}(\mathcal{D}, \text{gas})\right)$	-0.043	-0.042
M3: Within partial covariance with solar share $\left(\text{pCov}_{\text{within}}(\mathcal{D}, \text{solar})\right)$	0.017	0.017
M4: Within partial covariance with total capacity $\left(\text{pCov}_{\text{within}}(\mathcal{D}, \text{capTot})\right)$	0.034	0.035
M5: Within partial covariance with project capacity $\left(\text{pCov}_{\text{within}}(\mathcal{D}, \text{capProj})\right)$	-0.062	-0.063
M6: Within partial covariance with remoteness $\left(\text{pCov}_{\text{within}}(\mathcal{D}, \text{Rem})\right)$	0.025	0.023
M7: State–year partial covariance with carbon price $\left(\text{pCov}(\overline{\mathcal{D}}, \text{Carbon})\right)$	0.085	0.072
M8: State–year partial covariance with subsidy $\left(\text{pCov}(\overline{\mathcal{D}}, \text{Subsidy}^R)\right)$	-0.006	-0.010
M9: State–year partial covariance with RPS $\left(\text{pCov}(\overline{\mathcal{D}}, \text{RPS})\right)$	0.157	0.261
M10: Within variance of adoption $\left(\text{Var}_{\text{within}}[\mathcal{D}]\right)$	0.039	0.040
M11: Between variance of mean adoption $\left(\text{Var}_{\text{between}}[\bar{\mathcal{D}}_{s,t}]\right)$	0.209	0.208
M12: Within correlation with site quality $\left(\text{Corr}_{\text{within}}(\mathcal{D}, \widehat{\phi}^R)\right)$	0.408	0.318
M13: Tail share: $\bar{\mathcal{D}}_{s,t} < 0.10$	0.476	0.463
M14: Tail share: $\bar{\mathcal{D}}_{s,t} > 0.90$	0.341	0.348
M15: Mean adoption, top $\widehat{\phi}^R$ tercile	0.693	0.718
Diagnostic:		
Overall mean solar project adoption ($\bar{\mathcal{D}}$), matched by construction	0.611	0.611
Test of overidentifying restrictions	$Q(\widehat{\theta}) = 0.84$, $p = 0.49$	
95% critical value	3.53	

B.4.3 Parameter Units and Normalizations

Table 10 reports the estimated parameters. Two features of the estimation determine how those numbers should be read.

Units of the Policy Variables. Three variables enter the breakeven condition in equation (9) after being divided by their standard deviations across state–years: the RPS target (standard deviation 10.0 percentage points), the carbon allowance price (4.5), and the count of state subsidy programs (1.2). Thus, a one-unit increase in any of these rescaled variables means a one-standard-deviation increase in the original variable. For example, $\eta_{\text{rps}} = 0.380$ means that a 10-percentage-point increase in the RPS target raises the effective electricity-price term in the denominator of equation (9) by \$0.380 per kilowatt-year. The same one-standard-deviation interpretation applies to λ_{carb} and ψ. The electricity price and operating-cost terms are not rescaled and remain in dollars per kilowatt-year, so $\phi^\star$ remains a capacity factor.

Two Normalizations. The grid access term λ_{grid} enters equation (9) as a constant in both the numerator and the denominator. Its effect differs slightly across state–years because the other terms in the ratio vary. At the chosen value of 0.5, however, it is small relative to the electricity price, which averages about \$1,165 per kilowatt-year, and changes the breakeven threshold by well under one tenth of one percent. Most of this small effect shifts all projects in the same direction, while the intercept κ_0 already adjusts the model to match the overall adoption rate. The remaining differences across state–years provide too little variation to identify λ_{grid} separately with useful precision, so I fix it at 0.5.

Emissions intensity and the carbon-price coefficient are not separately identified because only the product $\lambda_{\text{carb}}\,\text{CO}_2^G$ enters the model. Thus, I set the emissions intensity of an average gas generator to $\text{CO}_2^G = 1$ and estimate the entire combined effect as λ_{carb}. Under this normalization, λ_{carb} measures the response to a one-standard-deviation increase in the carbon price for an average gas generator. Choosing a different numerical value for CO_2^G would change the reported scale of λ_{carb} by the opposite amount, but would not change model fit or the counterfactual results.

B.4.4 Estimation Settings and Robustness Checks

This subsection documents the numerical and calibrated choices behind the reported estimates, and asks whether moment fit and the main conclusions depend on them. Table B.7 shows the baseline specifications. The same $B = 200$ simulated environments are used at every parameter value so that changing simulation draws does not add noise to the objective function. The 40 starting values reduce the risk of selecting a local rather than global solution. The state-cluster bootstrap preserves dependence among observations within a state. For numerical stability, estimation uses the diagonal weighting matrix described in Section B.4.2; standard errors and the simulation-based overidentification test use the full bootstrap covariance.

The robustness checks in Table B.8 address three concerns. The first three alternatives test sensitivity to the size of the numerical search, the number of simulated environments, and the random seed. The next four change the discount rate and capital-cost series. The final three change the sample period and the proxy of workplace remoteness. Across all 11 estimates, Q ranges from 0.67 to 2.93 and every simulation-based p-value exceeds 0.05 (range 0.30–0.69). Thus, none of the overidentification tests rejects the model at the 5 percent level, and moment fit is not driven by one numerical setting, calibrated input, sample period, or alternative remoteness measure.

Table B.9 shows that the subsidy coefficient remains positive and significant when the discount rate rises from 0.05 to 0.09: it increases modestly from 0.507 (SE 0.086) to 0.604 (SE 0.094). The carbon-price coefficient is nearly unchanged, at 0.167 and 0.171, although it has higher statistical significance under the higher discount rate. By contrast, the RPS estimates of 0.708 and 1.109 have standard errors above 2, so their difference cannot be interpreted as a response to the discount rate.

Table B.10 shows the three policy effects on a common full-sample scale. The subsidy coefficient is 1.407 (SE 0.006) in 2011–2017 and 0.527 (SE 0.068) in 2018–2024, making the earlier response about 2.7 times as large. The carbon-price coefficient similarly falls from 0.492 (SE 0.022) to 0.202 (SE 0.096), a decline of about 59 percent. The RPS coefficient is also about one-third larger in the earlier period, at 0.806 (SE 0.106) compared with 0.600 (SE 0.671) later. Overall, the estimates suggest that policy incentives had greater leverage earlier in the transition, when solar faced a larger capital-cost disadvantage.

Table B.7: SMM Estimation Settings

This table records the settings used for the SMM estimates in Table 10.

Setting	Value
Simulated state–year environments	$B = 200$, with the same draws used at every parameter value
Random-number seed for environment draws	123
Moment covariance	1,000 state-cluster bootstrap replications (states resampled)
Estimation weighting matrix W	Diagonal inverse-variance weighting (estimated correlations excluded)
Inference	Sandwich standard errors using the full bootstrap covariance; simulation-based overidentification test in Section B.4.2
Numerical search	Two-step L-BFGS-B bounded optimizer; 40 starting values
Parameter bounds	$\sigma \geq 0.005$; $\alpha_\phi \leq 20$
Calibrated inputs and normalizations	$r = 0.07$; $CO_2^G = 1$; $\lambda_{\text{grid}} = 0.5$; $\Delta\kappa_t$ EIA series, \$4,430–\$800/kW (model year t uses generators installed in $t-1$)

Table B.8: SMM Robustness Checks

Each row reports a separate estimate. The first row is the estimate used in Table 10, based on 40 starting values for the numerical search. The other 10 rows are robustness checks. The second row reduces the search to 12 starting values; the remaining rows retain 12 starting values and change the listed input. Q is the overidentification statistic, so smaller values indicate closer moment fit. The reported p is the simulation-based p-value defined in Section B.4.2: it is the share of the 400,000 simulated Q values at least as large as the value shown.

Run	Change	Q	p
Reported estimate	40 starting values	0.84	0.49
Fewer starting values	12 starting values	0.78	0.66
More simulation draws	Environment draws $B = 250$	0.68	0.60
Different random draws	Different environment-draw seed	2.93	0.33
Lower discount rate	$r = 0.05$	0.90	0.61
Higher discount rate	$r = 0.09$	2.73	0.30
Lower capital costs	$\Delta\kappa_t$ scaled by 0.8	0.86	0.67
Higher capital costs	$\Delta\kappa_t$ scaled by 1.2	0.80	0.66
Earlier projects	2011–2017 subsample	0.67	0.69
Later projects	2018–2024 subsample	1.05	0.48
Wider remoteness ring	Population ring 10–60 km	1.55	0.51

Table B.9: SMM Estimates under Alternative Discount Rates

This table reports the full parameter estimates when the annual discount rate is set to 0.05 or 0.09.

		$r = 0.05$		$r = 0.09$	
Description	Notation	Coefficient	*SE*	Coefficient	*SE*
Solar share slope	β_{solar}	0.331	1.579	0.197	1.137
Coal share slope	β_{coal}	-0.041	0.111	-0.221	2.329
Gas share slope	β_{gas}	-0.087	0.353	-0.182	1.855
Total capacity slope	β_{capT}	0.121	0.473	0.397	3.583
Project capacity slope	β_{capP}	-0.124	0.416	-0.341	3.047
Remoteness site trait	β_{rem}	0.059	0.242	0.131	1.462
Subsidy pass-through	ψ	0.507	0.086	0.604	0.094
Carbon-price leverage on gas	λ_{carb}	0.167	0.167	0.171	0.026
RPS level effect	η_{rps}	0.708	2.071	1.109	2.492
Irradiance slope	α_{ϕ}	1.021	4.148	2.659	18.615
Idiosyncratic dispersion	σ	0.005	0.000	0.007	0.010
Environment dispersion	σ_{env}	0.248	6.590	0.106	1.113

Table B.10: SMM Estimates by Sample Period

This table reports the full parameter estimates for the 2011–2017 and 2018–2024 subsamples.

		2011–2017		2018–2024	
Description	Notation	Coefficient	*SE*	Coefficient	*SE*
Solar share slope	β_{solar}	0.110	0.278	0.354	8.535
Coal share slope	β_{coal}	-0.204	0.552	-0.077	0.812
Gas share slope	β_{gas}	-0.200	0.576	-0.161	0.853
Total capacity slope	β_{capT}	0.268	0.488	0.234	1.582
Project capacity slope	β_{capP}	-0.117	0.255	-0.253	1.289
Remoteness site trait	β_{rem}	0.094	0.649	0.075	0.486
Subsidy pass-through	ψ	1.407	0.006	0.527	0.068
Carbon-price leverage on gas	λ_{carb}	0.492	0.022	0.202	0.096
RPS level effect	η_{rps}	0.806	0.106	0.600	0.671
Irradiance slope	α_{ϕ}	2.604	14.906	2.122	12.759
Idiosyncratic dispersion	σ	0.006	0.000	0.006	0.001
Environment dispersion	σ_{env}	0.111	0.193	0.810	1.985

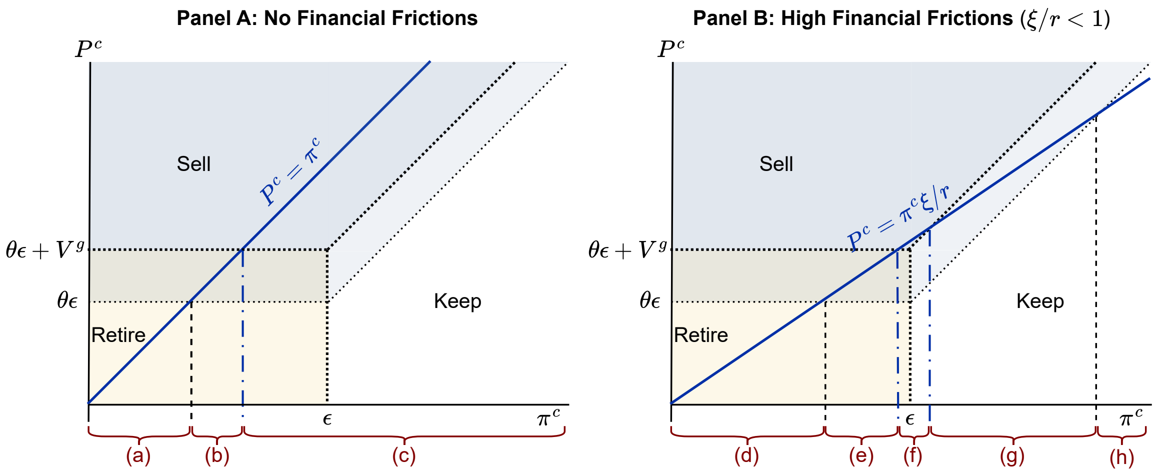

Owners' allocation in equilibrium (Panel A)

Buyer constraints:	*Baseline*	*With $V^g(\Delta)$*
Coal units turned off	(a)	(a) + (b)
Coal units sold	(b) + (c)	(c)
Coal units kept on	*None*	*None*

Owners' allocation in equilibrium (Panel B)

Buyer constraints:	*Baseline*	*With $V^g(\Delta)$*
Coal units turned off	(d)	(d) + (e)
Coal units sold	(e) + (f) + (g)	(f)
Coal units kept on	(h)	(g) + (h)

Figure B.3. Equilibrium Allocation of Coal Assets with a Gas Alternative for Buyers. This figure adapts and extends the framework in Darmouni and Zhang (2024), from Section "The Economics of Brown Capital (Re)Allocation". The horizontal axis is a coal unit's operating profitability π^c, and the vertical axis is the coal asset price P^c. In baseline model, an owner chooses among (i) *Keep/operate*, (ii) *Sell*, or (iii) *Retire/turn off*, with payoffs $\pi^c - \epsilon$, $P^c - \theta\epsilon$, and 0, respectively, where ϵ is the per-unit environmental externality cost faced by mandate owners and $\theta \in [0, 1]$ measures mandate stringency. Panel A assumes no financial frictions for buyers, and the equilibrium pricing line is $P^c = \pi^c$. Panel B assumes high financial frictions, and the equilibrium pricing line is $P^c = (\xi/r)\pi^c$ with $\xi/r < 1$, where ξ is the pledgeable share of profits and r is the required gross return to financiers. For both panels, the extension to the baseline models allows buyers to allocate capital to gas-fired generation as an alternative to acquiring coal assets, captured by $V^g \geq 0$, where V^g is a function of Δ which measures coal's fuel-cost disadvantage relative to gas. This term enters the owner's sell payoff, which becomes $P^c - \theta\epsilon - V^g(\Delta)$; accordingly, the sell-versus-retire threshold shifts from $P^c = \theta\epsilon$ to $P^c = \theta\epsilon + V^g(\Delta)$, shown by the higher horizontal cutoff in both panels. Brackets (a)–(h) denote ranges of π^c along the pricing equilibrium lines; the tables below summarize the corresponding allocation of coal units across *turned off*, *sold*, and *kept on* under the baseline and when incorporating $V^g(\Delta)$.

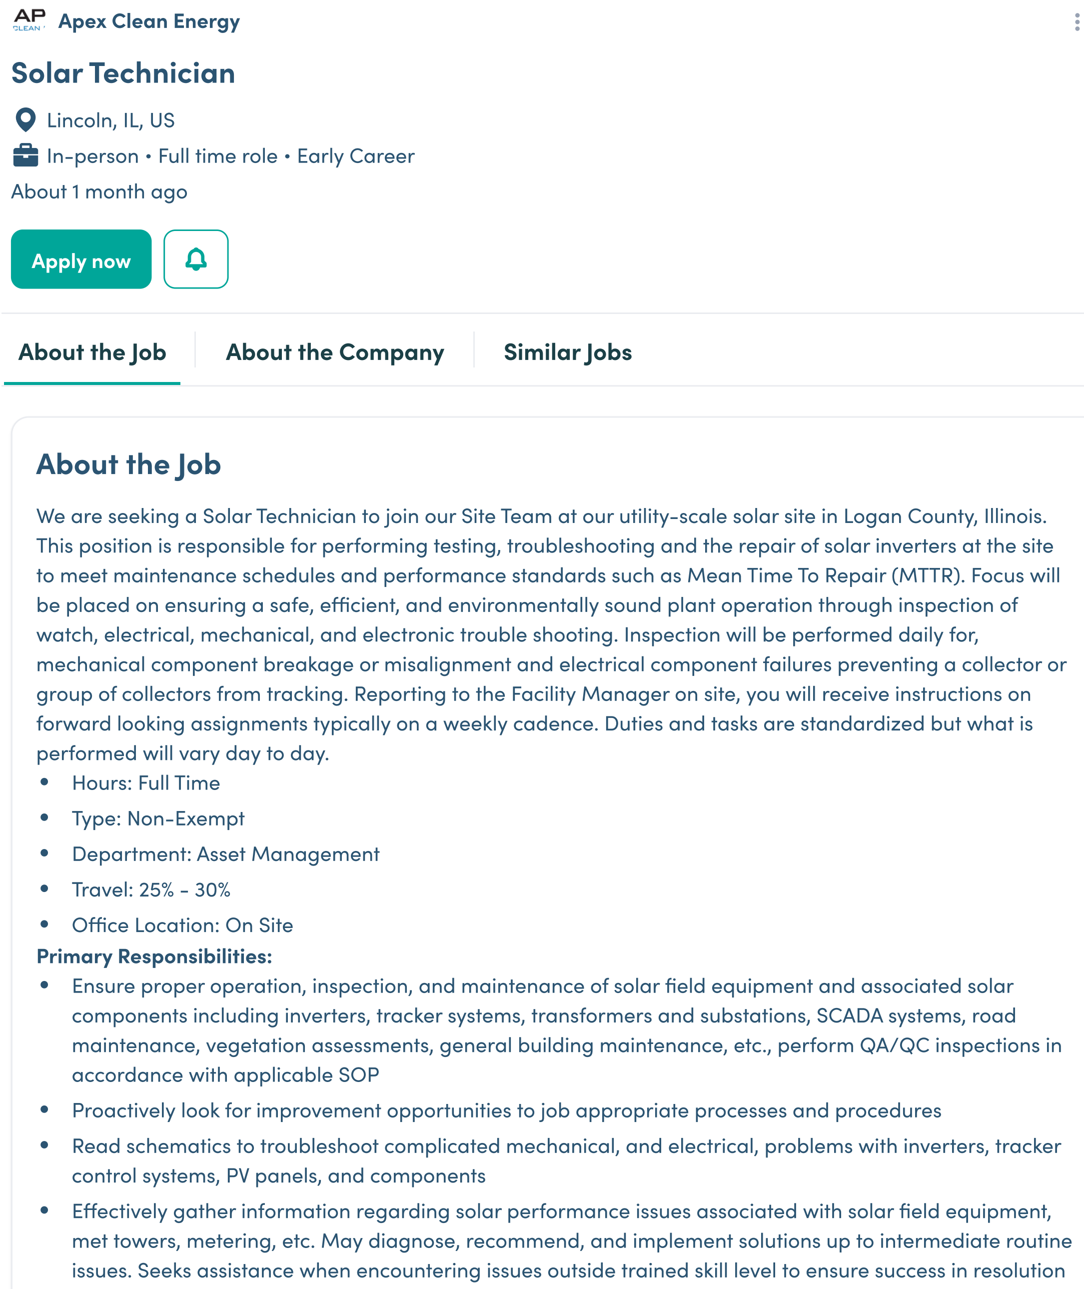

Apex Clean Energy

Solar Technician

Lincoln, IL, US

In-person • Full time role • Early Career

About 1 month ago

Apply now

About the Job | About the Company | Similar Jobs

About the Job

We are seeking a Solar Technician to join our Site Team at our utility-scale solar site in Logan County, Illinois. This position is responsible for performing testing, troubleshooting and the repair of solar inverters at the site to meet maintenance schedules and performance standards such as Mean Time To Repair (MTTR). Focus will be placed on ensuring a safe, efficient, and environmentally sound plant operation through inspection of watch, electrical, mechanical, and electronic trouble shooting. Inspection will be performed daily for, mechanical component breakage or misalignment and electrical component failures preventing a collector or group of collectors from tracking. Reporting to the Facility Manager on site, you will receive instructions on forward looking assignments typically on a weekly cadence. Duties and tasks are standardized but what is performed will vary day to day.

- Hours: Full Time
- Type: Non-Exempt
- Department: Asset Management
- Travel: 25% - 30%
- Office Location: On Site

Primary Responsibilities:

- Ensure proper operation, inspection, and maintenance of solar field equipment and associated solar components including inverters, tracker systems, transformers and substations, SCADA systems, road maintenance, vegetation assessments, general building maintenance, etc., perform QA/QC inspections in accordance with applicable SOP
- Proactively look for improvement opportunities to job appropriate processes and procedures
- Read schematics to troubleshoot complicated mechanical, and electrical, problems with inverters, tracker control systems, PV panels, and components
- Effectively gather information regarding solar performance issues associated with solar field equipment, met towers, metering, etc. May diagnose, recommend, and implement solutions up to intermediate routine issues. Seeks assistance when encountering issues outside trained skill level to ensure success in resolution

Figure B.4. Example Job Posting for Utility-Scale Solar Field Technician. This figure shows a job posting from a utility-scale solar facility operated by Apex Clean Energy in Logan County, Illinois, sourced from *climatebase.org*. The position involves on-site inspection, testing, and repair of solar inverters and tracker systems, as well as routine maintenance and fault diagnostics.