Organizational Structure, Environmental Justice, and the Geography of the Green Transition in the U.S. Power Sector

by

Peter Ludwig Suechting

A dissertation accepted and approved in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in Environmental Sciences, Studies, and Policy

Dissertation Committee:
Craig Kauffman, Chair
Leigh Johnson, Core Member
Richard York, Core Member
Raoul Lievanos, Institutional Representative

University of Oregon

Summer 2026

© 2026 Peter Ludwig Suechting

This work is openly licensed via
Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International
https://creativecommons.org/licenses/by-nc-nd/4.0/

Abstract

Peter Ludwig Suechting

Doctor of Philosophy in Environmental Sciences, Studies, and Policy

Title: Organizational Structure, Environmental Justice, and the Geography of the Green Transition in the U.S. Power Sector

Who bears the burden of fossil-fuel electricity generation, and who receives the benefits of what is replacing it? This dissertation poses that distributive question at the generating facility and resolves it into three: who installs emissions control equipment; who operates it, and how cleanly; and who hosts the capacity a decarbonizing fleet builds and retires. Three analyses answer them with federal data: generator censuses and fuel records from Energy Information Administration (EIA) Forms 860 and 923; emissions from the Environmental Protection Agency's Emissions & Generation Resource Integrated Database (eGRID) and Clean Air Markets Program (CAMD) monitors; and community demographics from the American Community Survey (ACS). The first two analyses observe coal- and gas-fired plants of at least 25 megawatts at three snapshots (2009, 2018, and 2021) spanning three regulatory regimes, 4,818 plant-years in all; the third follows the full fleet, renewables included, annually from 2010 through 2023, 13,063 plants and 112,782 plant-years.

Both traditions the design draws on, organizational political economy and environmental justice, predicted disparities in who installs control. Neither prediction holds; both reappear downstream, in how plants run once equipped. Rate-regulated utilities and cooperatives emit more nitrogen oxides (NO_x) per unit of electricity than merchant generators exposed to market discipline, a penalty that survives accounting for the equipment carried. By 2021, plants in more nonwhite communities likewise run measurably dirtier, independent of income and poverty. The burden is thus concentrating on the more nonwhite communities beside the aging urban plants that rising demand keeps in service,

while the buildout's benefits land in whiter states, as solar disperses between states rather than within them. Coal's retirements, against expectation, have whitened its surviving fleet: the plants that closed sat in more nonwhite communities than those that remain.

Organizational structure and community demographics leave their mark downstream of the decision to install, in how equipment already installed is run, and fleet turnover is what redistributes the burden. Encouraging the transition therefore requires regulation that encourages the entry of renewables, discourages the entry of new fossil capacity, and encourages the retirement of old.

Acknowledgments

I'd like to acknowledge the work of my committee chair, Dr. Craig Kauffman, who effectively and efficiently navigated the requirements of a student going through an unprecedented, difficult interdisciplinary program and handling it with ease. His positivity kept me oriented and moving through a complicated process. My thanks as well to my committee, Dr. Leigh Johnson, for believing in and inspiring my work, Dr. Raoul Liévanos, for critical expertise and helpful direction, and Dr. Richard York, for patience, mentorship, and wisdom.

My thanks to Kathy Furrer of the Environmental Sciences, Studies, and Policy program staff, who offered reassurance and guidance through the complicated, stressful process of defending this dissertation.

To my friends Julian Ricardo, Ethan Corey, Michel Legendre, Taggart Wass, Eric Johnson, Neil Gupta, Isaiah McGuire, Nakiessa Abassi, Mel Keller, and Rainer Lempert. You all know what you did.

To my parents, Ginny and Marty Suechting, your love and support made this possible. To my brother Max, who always had wise words and time for a phone call to help me keep my head screwed on straight. And to Ellen, Clark, and Bodhi (and let's not forget Bridget and Fred, too), my love and appreciation for providing the moral support and the motivation to finish the project strong.

To Mowgli and Juniper, my best buds and constant companions, I look forward to spending more time with you (and I thank you for your patience throughout).

And to Amanda. My love. This is as much your accomplishment as it is mine.

Dedication

To Zachary Hamilton Suechting (1991–2019).

Table of contents

1 Introduction 19

1.1 The distributive question 19

1.2 How electricity gets made, and by whom 22

1.3 Nitrogen oxides: formation, control, and measurement 24

1.4 Literature review 29

1.5 Questions and expectations by chapter 32

1.5.1 The Organization of Pollution Control 32

1.5.2 Environmental Justice at the Operating Margin 33

1.5.3 The Unequal Geography of the Green Transition 34

1.6 Design and data 35

1.7 Plan of the dissertation 37

2 The Organization of Pollution Control 39

2.1 Introduction 40

2.1.1 Research Questions 44

2.1.2 Research Design 44

2.1.3 Hypotheses 51

2.2 Data & Methods 62

2.2.1 Data sources 63

2.2.2 Analytical Strategy 64

2.3 Results 78

2.3.1 Descriptive Statistics 78

2.3.2 Primary Results 80

2.4 Discussion . . . 90
2.5 Limitations . . . 95
2.6 Conclusion . . . 97
Appendix — Variable inventory . . . 98
2.6.1 Dependent Variables . . . 98
2.6.2 Primary IV: Organizational Structure . . . 98
2.6.3 Capacity to emit . . . 101
2.6.4 Regulatory controls . . . 101
2.6.5 Political-economic context . . . 103
2.6.6 NO_x Control — Functional Mediator (Primary) . . . 104
2.7 Appendix — Coal-boiler deep-dive . . . 105
2.8 Appendix — Full regression tables . . . 106
2.8.1 Distributional diagnostics . . . 109
2.8.2 Collinearity diagnostics . . . 110
Computational environment . . . 112

3 The Social Incidence of Pollution . . . 113
3.1 Introduction . . . 114
3.1.1 Research Questions . . . 116
3.2 Theory and Conceptual Framework . . . 116
3.2.1 The environmental justice tradition: race, class, and distributive disparity 117
3.2.2 From proximity to abatement: relocating the environmental justice outcome 118
3.2.3 Path of least resistance: from siting to abatement investment . . . 119
3.2.4 Two margins of abatement: bridging Chapter 2 . . . 123
3.3 Hypotheses . . . 124
3.3.1 H1 — Adoption . . . 125
3.3.2 H2 — Mechanism . . . 126
3.3.3 H3 — Race net of class . . . 127
3.4 Data & Methods . . . 128
3.4.1 Data sources . . . 128

3.4.2 Unit of Analysis . . . 129
3.4.3 Scope of Analysis . . . 129
3.4.4 Dependent Variables . . . 130
3.4.5 Primary Independent Variable: Community Demographics . . . 131
3.4.6 Mediator: Functional NO_x-Control Adoption . . . 134
3.4.7 Control Variables . . . 134
3.4.8 Analytical Strategy . . . 135
3.4.9 Estimating Equations . . . 137
3.4.10 Identification and Causal Language . . . 138
3.4.11 Spatial Autocorrelation . . . 140
3.4.12 Estimation Procedure . . . 143
3.5 Results . . . 144
3.5.1 Descriptive Statistics . . . 144
3.5.2 Primary Results . . . 148
3.6 Discussion . . . 158
3.7 Limitations . . . 164
3.8 Conclusion . . . 166
Computational environment . . . 167
Appendix — Variable inventory . . . 168
3.8.1 Dependent Variables . . . 168
3.8.2 Primary Independent Variable: Community Demographics . . . 168
3.8.3 Mediator: Functional NO_x-Control Adoption . . . 173
3.8.4 Control Variables . . . 174
Appendix — Full regression tables . . . 177

4 The Unequal Geography of the Green Transition . . . 183
4.1 Introduction . . . 184
4.2 Background and Theoretical Framing . . . 187
4.2.1 Environmental inequality formation: power resources and the path of least resistance . . . 187

4.2.2 Siting, sorting, and exit: the formation and redrawing of disparity 189
4.2.3 Material profiles and the limits of least resistance: why the transition's geography is fuel-specific 190
4.2.4 From a static cross-section to a trajectory: positioning and scope 193
4.2.5 Research questions 193
4.3 Data and Methods 194
4.3.1 Data sources 194
4.3.2 Sample and unit of analysis 194
4.3.3 Community demographic measurement 195
4.3.4 Energy-type classification 198
4.3.5 Analytic strategy 198
4.3.6 Identification and interpretation 201
4.4 Expectations 203
4.4.1 E1: Trajectory (RQ1) 203
4.4.2 E2: Decomposition (RQ2) 203
4.4.3 E3: Entrant composition (RQ3) 204
4.5 Results 206
4.5.1 The changing geography of energy-type communities (RQ1) 207
4.5.2 Turnover versus neighborhood change (RQ2) 210
4.5.3 Where new capacity lands (RQ3) 212
4.5.4 The Indigenous dimension 216
4.6 Discussion 217
4.7 Limitations 222
4.8 Conclusion 223
4.9 Appendix 224
4.9.1 Variable inventory 224
4.9.2 Plant-location measurement 228
4.9.3 All energy types: full trajectory reference 230
4.9.4 Decomposition: all fuels and both windows 231

4.9.5 Entrant composition: all fuels and outcomes . . . 232
4.9.6 Distributional context: solar host communities over time . . . 234
4.9.7 Robustness: buffer radius . . . 235
4.9.8 Full results tables . . . 236
4.9.9 Robustness: functional form . . . 238
4.9.10 Robustness: Indigenous proximity threshold . . . 239
4.9.11 Robustness: natural-gas prime mover . . . 240

5 Conclusion . . . 242
5.1 The distributive question, answered . . . 242
5.2 Findings in review . . . 243
5.2.1 The Organization of Pollution Control . . . 243
5.2.2 Environmental Justice at the Operating Margin . . . 244
5.2.3 The Unequal Geography of the Green Transition . . . 244
5.3 What the three chapters say together . . . 245
5.3.1 From adoption to operation . . . 245
5.3.2 The peaker at the junction . . . 246
5.3.3 Turnover as redistribution . . . 247
5.4 Policy implications . . . 248
5.5 Limitations and the road forward . . . 249
5.6 Coda . . . 249

References Cited . . . 251

List of Figures

1.1 Annual U.S. electricity net generation by source, electric power sector, 2001 to 2025. Source: Table 7.2a ("Monthly Energy Review" 2026). 20

1.2 Where nitrogen oxides form in combustion, and where the two control margins act. 27

1.3 Schematic of the selective non-catalytic reduction temperature window: NO_x conversion peaks within a narrow band of flue-gas temperature, while unreacted ammonia escapes as slip below it. Vertical scale is relative, not a percentage. Window bounds after Richards (2000); form after Pronobis (2020). 29

2.1 NO_x emission intensity by fuel type (ridges) and snapshot year (panels). 67

2.2 NO_x emission intensity by the joint state of the two control functions. 68

2.3 Fleet composition across the four control configurations, by fuel and snapshot. . . 69

2.4 Sequential variance decomposition of log NO_x intensity, mirroring the primary specification. 69

2.5 Conceptual map of the primary analysis (full fossil fleet, $\geq$25 MW). 75

2.6 Organizational-structure coefficients on the two NO_x-control adoption margins. . 84

2.7 Organizational-structure coefficients on log NO_x intensity (full sample), ranked by magnitude. 88

3.1 The spatial distribution of the demographic predictor: each coal- or gas-fired plant in the 2018 analytical sample, shaded by the share nonwhite of its 3 km population-weighted buffer, with the predictor's Moran's I. 142

3.2 Distributions of the three demographic measures across the four linkage geographies. 146

3.3 The 2018 fleet by its two outcome measures, total NO_x mass (x) and emission rate (y), both log-scaled, colored by the plant's coal heat share. 148
3.4 Average marginal effects of community demographics on NO_x-control adoption, by margin. 149
3.5 The percent-nonwhite gradient in log NO_x intensity by snapshot year. 152
3.6 The 2021 percent-nonwhite gradient across the robustness family. 155
3.7 The aggregate percent-nonwhite gradient against the Black and Hispanic component gradients. 158
3.8 The top decile of plants by 2021 NO_x emission rate. 163
3.9 Community percent nonwhite: the top intensity decile versus the fleet, by year. . 164
3.10 Population-weighted apportionment around a dense urban example plant. 171
3.11 The apportionment construction for a sparse rural example plant. 173

4.1 Power density, generating capacity per unit of land, by energy type, on a log scale. 197
4.2 Host-community composition by energy type, net of the national trend, 2010–2023. 209
4.3 Transition-core fuels in the race–class plane, 2010–2023. 210
4.4 Melitz–Polanec decomposition of share nonwhite, transition-core fuels, 2010–2019. 212
4.5 Entrant-versus-incumbent gap in share nonwhite, within and between states, by fuel. 214
4.6 Utility-scale solar plant locations by build era (pinned coordinates), with plant counts by census region; state outlines are colored on the same region scale as a geographic key (Alaska, Hawaii, and DC fall outside the base map and appear only in the bars). 215
4.7 Within-state entrant-versus-incumbent gap in the probability of a high-Indigenous host community (>=5% Indigenous, 3 km buffer), by fuel, with 95% CIs (state-clustered); colored where the interval excludes zero. 217
4.8 Host-community share nonwhite in 2023, net of the CONUS national population, by fuel, counted two ways: by plant (each generator weighted equally) and by person (each plant weighted by its buffer population). 220
4.9 Number of host communities (operating plants) by energy type, 2010–2023, for the four transition-core fuels. 221

4.10 Host-community composition for all eleven energy types, net of the national trend. 231
4.11 Melitz–Polanec decomposition of the change in host-community share nonwhite (3 km), all eleven energy types, for the primary 2010–2019 window and the 2010–2023 extension; channels signed, the black diamond the total. Fuels with no surviving plants across a window carry a total but no channel bars. 232
4.12 Entrant composition for all eleven energy types and all outcomes. 233
4.13 Per-year density of solar host-community share nonwhite, 2010–2023. 235
4.14 Solar's host-community share nonwhite, net of the CONUS national population, at 3, 5, and 10 km buffers. 236
4.15 Melitz–Polanec decomposition of the change in host-community share nonwhite (3 km), 2010–2023, with natural gas split by prime mover into combined-cycle baseload (gas_cc) and simple-cycle peaking (gas_peaker); channels signed (percentage points), the black diamond the total. 241

List of Tables

1.1 The three empirical chapters compared. 35

2.1 Organizational-type typology and construction logic: each plant is sorted by NAICS economic sector, then rate regulation, then ownership (plant-years pooled over 2009, 2018, 2021). 46

2.2 Predicted signs by hypothesis across the two adoption margins (Eq 1a, 1b), their differential, and the total and conditional outcome equations (Eq 3a, 3b). 53

2.3 Data sources for the analysis. 63

2.4 Plant characteristics by presence class (survivors, exiters, entrants) across the three snapshots; blip/gap classes excluded. 72

2.5 Descriptive statistics for primary analytical variables (coal + gas, $\geq$25 MW control-measured fleet; 2009/2018/2021 snapshots). 78

2.6 Organizational-structure coefficients on the two adoption margins, full fossil fleet (logistic regression, log-odds): combustion modification (Equation 1a) versus post-combustion treatment (Equation 1b). 82

2.7 Organizational-structure coefficients on log NO_x intensity, full sample, before and after conditioning on the two adopted-control indicators. 87

2.8 Hypothesis scorecard: predicted versus observed signs for each hypothesis across the two adoption margins (Eq 1a, 1b), their differential, and the total and conditional outcome (Eq 3a, 3b). 89

2.9 Codebook: dependent variables. 98

2.10 Codebook: ownership and regulatory status. 99

2.11 Codebook: firm scale and market position. 100

2.12 Codebook: capacity to emit. 101
2.13 Codebook: regulatory controls. 102
2.14 Codebook: political-economic context. 103
2.15 Codebook: NO_x-control functional mediators. 104
2.16 Organizational-structure coefficients in the coal-boiler within-adopter deep-dive (post-combustion adopter subsample). 106
2.17 Full adoption-equation estimates (Equations 1a and 1b), logistic regression in log-odds. 107
2.18 Full outcome-equation estimates (Equations 3a and 3b), OLS on log NO_x intensity; Equation 3b adds the two adopted-control indicators. 108
2.19 Distribution diagnostics: dependent variables (native and log forms). 109
2.20 Distribution diagnostics: transformed and composite predictors in the primary specification. 110
2.21 Distribution diagnostics: proportions, counts, rates, and ages in the primary specification. 110
2.22 Variance inflation factors, conditional outcome model (Equation 3b), all predictors (year and region fixed effects omitted); conventional flags are VIF > 5 for inspection, VIF > 10 as problematic. 110
2.23 Pearson correlations among continuous predictors. 111
2.24 R packages used in the analysis pipeline (versions as rendered). 112

3.1 Predicted signs by demographic measure across the two adoption margins (Eq 1a, 1b), their differential, and the total and conditional intensity equations (Eq 3a, 3b); signs were fixed before estimation. 125
3.2 Data sources for the analysis. 128
3.3 Community context by organizational type: median neighborhood density (census tracts within 3 km) and median 3 km buffer composition, pooled over the three snapshots. 139
3.4 Plant co-location by snapshot: distance to the nearest same-year neighbor and the share of plants with at least one neighbor within 5, 10, 25, and 50 km. 142

3.5 Descriptive statistics for the analytical sample (pooled over 2009/2018/2021), by variable role. . . . 144
3.6 Average marginal effects of community demographics on adoption of each NO_x control and the post-minus-combustion differential, utility-clustered; scaling as in Figure 3.4, on the probability scale. . . . 150
3.7 Community demographics and plant-level NO_x emission intensity (log lb/MWh): total (Eq 3a) and adoption-conditional (Eq 3b) associations, pooled with demographic × year interactions; main effects are the 2009 reference. . . . 152
3.8 Robustness of the 2021 percent-nonwhite gradient: each cell is the nonwhite × 2021 coefficient (standard error) from a re-estimation changing one element of the primary specification, for Eq 3a and Eq 3b. . . . 155
3.9 R packages used in the analysis pipeline (versions as rendered). . . . 167
3.10 Codebook: dependent variables. . . . 168
3.11 Codebook: community demographics, buffer measures. . . . 169
3.12 Codebook: community demographics, nearest-tract measures. . . . 169
3.13 Demographic estimates across the buffer radii (3 / 5 / 10 km) and the nearest-centroid tract for the two example plants (2018 snapshot). . . . 171
3.14 Codebook: functional NO_x-control adoption. . . . 174
3.15 Codebook: capacity-to-emit controls. . . . 174
3.16 Codebook: state political-context controls. . . . 175
3.17 Codebook: organizational-structure controls. . . . 176
3.18 Full adoption-equation estimates (Equations 1a and 1b), logistic regression in log-odds. 178
3.19 Full outcome-equation estimates (Equations 3a and 3b), OLS on log NO_x intensity; Equation 3b adds the two adopted-control indicators. . . . 179
3.20 Exploratory race disaggregation of the adoption equations (1a, 1b), logistic log-odds, with separate Black/Hispanic/Asian shares replacing pooled percent nonwhite. . . . 180
3.21 Exploratory race disaggregation of the outcome equations (3a, 3b), OLS on log NO_x intensity, with separate Black/Hispanic/Asian shares replacing pooled percent nonwhite. . . . 181

4.1 Data sources for the plant-year panel. 194
4.2 Pre-registered expectations by decomposition channel: directional priors committed before estimation for entry, exit, and within-survivor drift. 205
4.3 Host-community characteristics by energy type, pooled over operating plant-years, 2010–2023. 206
4.4 Codebook: identifiers and model terms. 224
4.5 Codebook: predictor and fleet dynamics. 225
4.6 Codebook: community-demographic outcomes. 226
4.7 EIA-860 single-point coordinate drift across years, by energy type (raw, pre-pinning). 228
4.8 Wind host-community demographics: turbine-array buffers minus point buffers, by radius. 229
4.9 U.S. large-scale solar facility footprint (USPVDB); area-equivalent radius. 229
4.10 Entrant composition: total entrant-versus-incumbent gap (year fixed effects only), by fuel and outcome. 237
4.11 Entrant composition: within-state entrant-versus-incumbent gap (year and state fixed effects), by fuel and outcome; cells as in the previous table. 237
4.12 Melitz–Polanec decomposition of the change in host-community share nonwhite (3 km) by fuel and channel, for the 2010–2019 window and the 2010–2023 extension. 237
4.13 Entrant-versus-incumbent gap in host-community race (full specification, percentage points), for the four transition-core fuels, under three codings of the outcome: linear share nonwhite (primary), a majority-minority indicator (share > 0.5), and a top-tertile indicator; linear-probability models for the binaries. 239
4.14 Entrant-versus-incumbent gap in high-Indigenous proximity (full specification, percentage points), for the four transition-core fuels, at the primary >=5% cutoff, a >=10% cutoff, and the continuous non-Hispanic AIAN share. 239

1 Introduction

Organizational Structure, Environmental Justice, and the Geography of the Green Transition in the U.S. Power Sector

1.1 The distributive question

Why has the response to climate change failed to match the urgency the science describes? The familiar answers run from the ideational through the material to the political: a deficit of knowledge; a conflict of values; the immaturity of the technology; the costs of abatement; and the organization of political power (Aklin and Mildenberger 2020; Meckling 2015; Klein 2014; Malm 2016; Christophers 2024). However those answers are weighed, the response they all presuppose is an energy transition: the replacement of fossil energy combustion with ways of generating energy that neither burn fossil fuel nor release greenhouse gases. Where the politics of that transition is fought has meanwhile shifted. Bosworth, surveying the American climate movement after the exhaustion of its explicitly named phase, finds climate politics persisting implicitly, carried in sectoral struggles over utilities and infrastructure rather than under its own banner (Bosworth 2026). The chapters that follow work on that terrain: their subject is the material operation of electricity generation, and climate change enters them as the transition remaking that operation rather than as an explicitly named topic. For whatever the shortfall in American federal ambition on climate change, government and private sector actions worldwide over the past fifty years have nonetheless set off a period of intense compositional change in the American electricity supply (Figure 1.1), driven by the development and commercialization of new fossil and renewable energy sources

(Yergin 2012, 2021; Lachapelle, Macneil, and Paterson 2017). Coal generation has fallen steeply from its mid-2000s peak. Natural gas has moved to the center of the mix. And, relevant to the green transition, wind and solar have climbed from the margins of the system to a substantial share of what gets built. Considered as a matter of the energy mix, the past two decades are not a period of stasis but a period of rapid change—only a fraction of which is 'green'—whose consequences remain poorly described. Who bears the burden of the fossil capacity still in service, and who receives the benefits of what is replacing it, is the question this dissertation takes up.

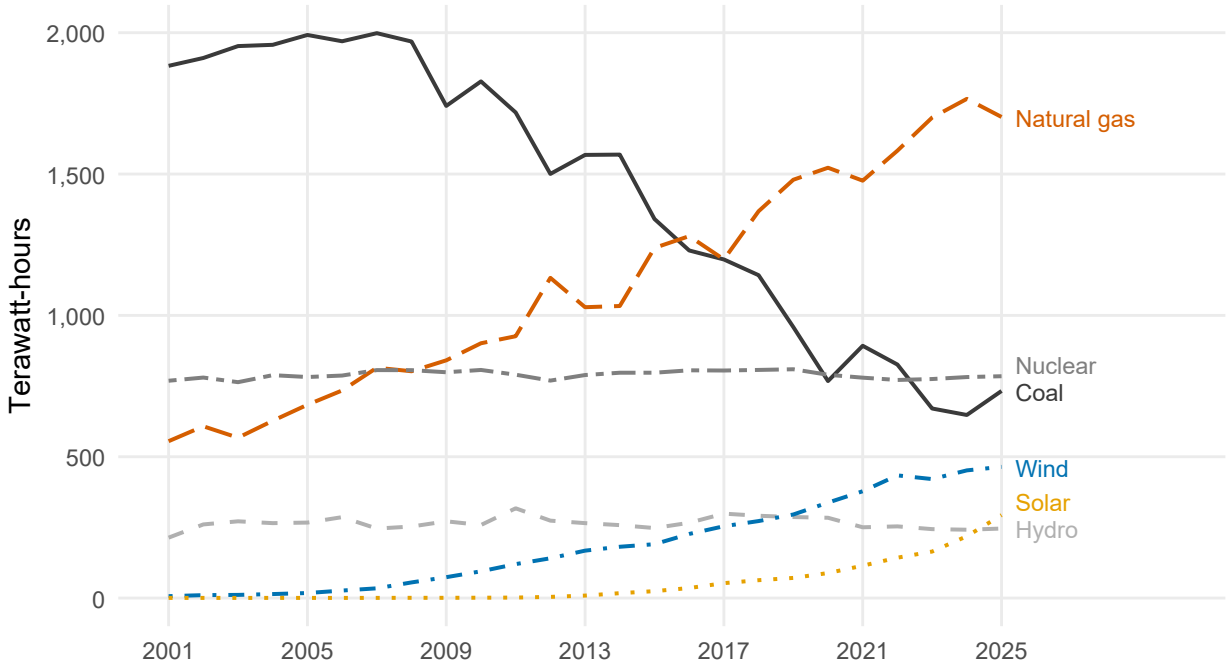

Figure 1.1: Annual U.S. electricity net generation by source, electric power sector, 2001 to 2025. Source: Table 7.2a ("Monthly Energy Review" 2026).

Why electricity generation? It is the largest stationary source of American greenhouse gas emissions. Within it, the generating facility—the site where fuel is burned to make electricity, and where the decisions that govern how much pollution accompanies that burning are made—is the transition's fundamental unit, and therefore the logical place to ask this question. Two research traditions converge there. Organizational political economy explains pollution as a product of the firms that own and operate industrial assets, their structures, and their incentives (Grant et al. 2010; Prechel and Istvan 2016). (Organizations, in Perrow's formulation, are "the most intensive and effective environmental destroyers" (Perrow 1997).) Environmental justice explains its distribution

as a product of the political and economic processes that place hazards near some communities rather than others (Pellow 2000; Mohai and Saha 2015). Both treat environmental harm as organizationally produced rather than technically given. Much work in these two veins has concentrated on what a plant is and where it sits: the fuel it burns, the technology it carries, the community that hosts it. But the plant as an asset that is operated, day after day, by an organization with its own structure, obligations, and slack has received less attention. Yet two plants burning the same fuel under the same federal rules can emit at markedly different rates. Part of that spread is by regulatory design: allowance trading lets a plant buy the right to emit rather than abate, on the premise that abatement should happen wherever it is cheapest (EPA 2025), and standards written for new sources leave the older units still in service under looser obligations. Further discretion exists between combustion and smokestack, with multiple sites at which decisions are made by owners with different structures and incentives, sited in communities with different capacities to object (Prechel and Istvan 2016; Grant, Jorgenson, and Longhofer 2016): a margin of operational discretion between a plant's technical capacity to emit and the pollution it actually produces.

If pollution is an organizational outcome, its distribution across communities may be an organizational outcome as well. The question animating the chapters that follow is accordingly distributive: how do electricity generating facilities' organizational and community contexts shape both the burden of fossil-fuel pollution and the distribution of clean-energy benefits, and do those burdens and benefits fall equitably across lines of race, income, and geography? Posed at the facility level, this question resolves into three sub-questions: 1) who installs emissions control equipment?; 2) who operates the equipment they have installed, and how cleanly do they operate it?; and 3) who hosts the generating capacity a decarbonizing fleet builds and retires? Each of these questions is a station along a single pathway rather than a separate topic, an argument made clear later in Section 1.6. The analysis will show that the influence of both organizational structure and community composition enters later along that pathway than the literature has generally looked for it. Adoption, the decision to install emissions control equipment, proves the least discriminating of the three: what a plant carries is largely settled by regulation and engineering. The variation that matters lies downstream, in operation. The analytic move is akin to disaggregating the influence of car and driver on race times: emission intensity is produced jointly by equipment and operator, and separating their contributions

means comparing plants that carry the same controls but are run by different organizations.

To this point the generating facility has figured as a generic technology: fuel in, electricity out, an organization in between. But the burden a facility places on its neighbors depends on what kind of facility it is. A coal or gas plant emits into the air its hosts breathe; a wind farm or solar array does not, whatever else it may impose; so the composition of the fleet, at any moment, is also a distribution of burden across the communities that host it. The fleet is meanwhile turning over, and turnover redistributes pollution across communities in ways no single plant-level snapshot can capture. Turnover runs on two channels: what enters, and what retires. By and large, coal retires; gas and renewables enter. Rising electricity demand, driven by electrification and now by the data centers built to train and serve artificial intelligence (Shehabi et al. 2024), is slowing retirement, and stands to hold older capacity in service, still emitting, longer than announced schedules alone would suggest. The same demand is also at work on the entry side, reviving commercial interest in nuclear generation.

1.2 How electricity gets made, and by whom

Almost all electricity in the United States is produced by spinning a magnet inside a coil of wire. Generating technologies differ mainly in what does the spinning. In a coal plant, burning fuel boils water, and the resulting steam turns a turbine. A gas plant may do the same; or it may burn fuel to drive a turbine directly, in the manner of a jet engine; or it may combine the two, capturing the exhaust heat of the first stage to raise steam for a second. Nuclear plants also raise steam, differing from coal chiefly in the heat source. Wind turbines dispense with fuel and let moving air do the work. Solar photovoltaics dispense with the turbine altogether, converting light into current with no moving parts and no combustion at all.

The second feature to hold constant is the duty cycle: facilities supplying the grid run on very different schedules. Electricity cannot be stored at scale (yet), so generation must match demand instant by instant, and demand varies by hour, by day, and by season. Grid operators meet it by dispatching plants in order of operating cost, cheapest first, which sorts the fleet into rough classes of use. Baseload plants run more or less continuously. Intermediate or load-following plants follow

the daily rise and fall of demand. Peaking plants, or ‘peakers,’ run only when demand is highest, sometimes for only a few hundred hours a year, starting and stopping as they are called. The conventional summary measure is the capacity factor: actual output over a period divided by what the plant would have produced running flat out. A well-utilized combined-cycle plant may post a capacity factor above one half; a peaker may spend well under a tenth of the year generating. How hard and how steadily a plant runs matters for its emissions as much as what it burns, for reasons taken up in the next section.

Who owns these plants matters for emissions as well, at least where fuel is burned. American generating capacity is held under several distinct economic arrangements, and the analysis in Chapter 2 sorts owners into five roles. A rate-regulated investor-owned utility recovers its costs through rates set by a state public utility commission, which allows the firm to earn a regulated return on its invested capital, the rate base; prudently incurred operating costs are passed through to customers. The arrangement has long been understood to shape investment incentives, since a firm earning a return on capital has reason to prefer capital solutions (Averch and Johnson 1962). A cooperative is owned by the customers it serves and operates without shareholders or a rate base, recovering its costs from its members. Government or public power owners, chiefly municipal utilities and federal or state authorities, finance construction with tax-exempt debt. An on-site industrial generator produces mainly for the host facility’s own process rather than for sale. A merchant generator, finally, owns plants but no captive customers: it sells its output at market prices and bears the resulting price risk itself. The Energy Policy Act of 1992 created the Exempt Wholesale Generator status that made this business model possible at scale, and merchant ownership grew from it.

That last category is inseparable from the wholesale market, which is where merchant generators sell. Beginning in the 1990s, a series of federal and state reforms separated the generation of electricity from its transmission and distribution in much of the country, opened transmission lines to third-party generators, and established independent system operators to run centralized auctions in which generators bid to supply power (Borenstein and Bushnell 2015). In those regions a merchant generator lives on the price its output clears at, hour by hour. In the rest of the country, notably much of the Southeast and West, utilities remain vertically integrated: the same firm generates, transmits, and delivers, and recovers the cost of all three through regulated retail rates. The distinction between

a merchant and a regulated owner is therefore not merely a matter of corporate form. It is a difference in what the owner is exposed to, and in whether an additional dollar of operating cost is borne by the firm or passed to a ratepayer.

These arrangements also govern the pace and shape of the clean-energy buildout, though not in the way a cost comparison alone would suggest. Christophers argues that the constraint on renewable deployment is not the cost of building wind and solar, which has fallen dramatically, but their profitability, which depends on revenues that wholesale markets deliver unreliably to generators whose marginal cost is near zero (Christophers 2024). On that account, capital moves toward clean generation when the expectation of profit warrants it, and cheapness alone does not warrant it. The questions taken up here sit downstream of that one. They concern what the capacity that does get built does to the communities near it, and how the fossil capacity still in service is operated while the transition proceeds.

1.3 Nitrogen oxides: formation, control, and measurement

Nitrogen oxides, written NO_x, are a family of reactive gases formed when fuels burn in air. Why build the analysis around this pollutant rather than another? Three properties recommend it. First, NO_x is a product of combustion itself rather than of any particular fuel: coal, gas, and oil all produce it when burned, so one emission measure can compare plants across the fuels that make up the fossil fleet. Sulfur dioxide, by contrast, is largely a coal problem, since natural gas carries almost no sulfur, and an analysis built on it would forfeit the gas fleet. Second, NO_x preserves a live operating margin. Sulfur dioxide is dominated by its control equipment: a scrubber, once installed, removes most of what would otherwise be emitted, leaving relatively little for operating practice to determine. Carbon dioxide has no control equipment at all, and its rate is fixed almost entirely by fuel chemistry and efficiency. NO_x sits between them as the one major pollutant whose emission rate is set jointly by what equipment a plant carries and by how that plant is run. Third, it is measured hourly at the stack by continuous emissions monitoring systems (EPA 2025).

Two members of the family matter for air quality: nitric oxide (NO), which makes up the great majority of what leaves a stack, and nitrogen dioxide (NO_2), into which much of it converts in

the atmosphere. A third nitrogen oxide, nitrous oxide (N_2O), is excluded from the category: it is comparatively unreactive at ground level, and its significance is as a greenhouse gas and a destroyer of stratospheric ozone rather than as a local pollutant. NO_x matters locally because it is a precursor. In sunlight it drives the formation of ground-level ozone, and it contributes to the secondary formation of fine particulate matter. EPA's most recent integrated science assessment concludes that the evidence supports a causal relationship between both short-term and long-term exposure to oxides of nitrogen and respiratory effects, with people who already have asthma at elevated risk from short-term exposure and children at elevated risk from exposure of either duration (Krajewski et al. 2026).

How much NO_x a plant produces is not a fixed property of its fuel. It is a property of how that fuel is burned. Two formation routes account for nearly all of it (Richards 2000). *Thermal* NO_x forms when the nitrogen and oxygen already present in combustion air combine at high temperature, by a chain of reactions first described by Zeldovich in 1946. The route is steeply temperature dependent: it becomes significant only above roughly 1,500°C (2,700°F) and peaks between about 1,900°C and 2,000°C, which is precisely the range in which most stationary combustion operates. Its magnitude is governed by three things: the peak flame temperature, the oxygen available in the hottest part of the flame, and how long the gases linger there. *Fuel* NO_x forms instead from nitrogen chemically bound in the fuel itself. It is consequential for coal, whose nitrogen content runs from roughly 0.2 to 3.5 percent by weight, and for heavy oil, and essentially absent for natural gas, which carries no bound nitrogen to speak of.

That two-route chemistry produces two places to intervene, and the distinction between them organizes the analysis in Chapters 2 and 3. The first margin, combustion modification, attacks formation. Low-NO_x burners, overfire air, flue gas recirculation, low excess air, and fuel reburning all work by lowering the peak flame temperature, reducing the oxygen available where the flame is hottest, or shortening the time the gases spend in that zone. Which of those levers does the work depends on the fuel: lowering the peak temperature and shortening residence time suppress the thermal route, while controlling the oxygen available where the fuel's nitrogen is released suppresses both routes and is the principal lever on the fuel one. Because natural gas carries no bound nitrogen, combustion modification at a gas plant works on the thermal route alone, where at a coal plant it

works on both. These measures are comparatively inexpensive and can often be retrofitted into an existing boiler, and they deliver moderate reductions. The second margin, post-combustion treatment, concedes the formation and removes NO_x from the flue gas afterward, injecting a nitrogen-based reagent that converts it back to molecular nitrogen and water. Selective non-catalytic reduction does this within a narrow temperature window and without a catalyst; selective catalytic reduction passes the gas over a catalyst bed and achieves substantially deeper removal (Richards 2000; Pronobis 2020). Both are capital-intensive, and selective catalytic reduction especially so.

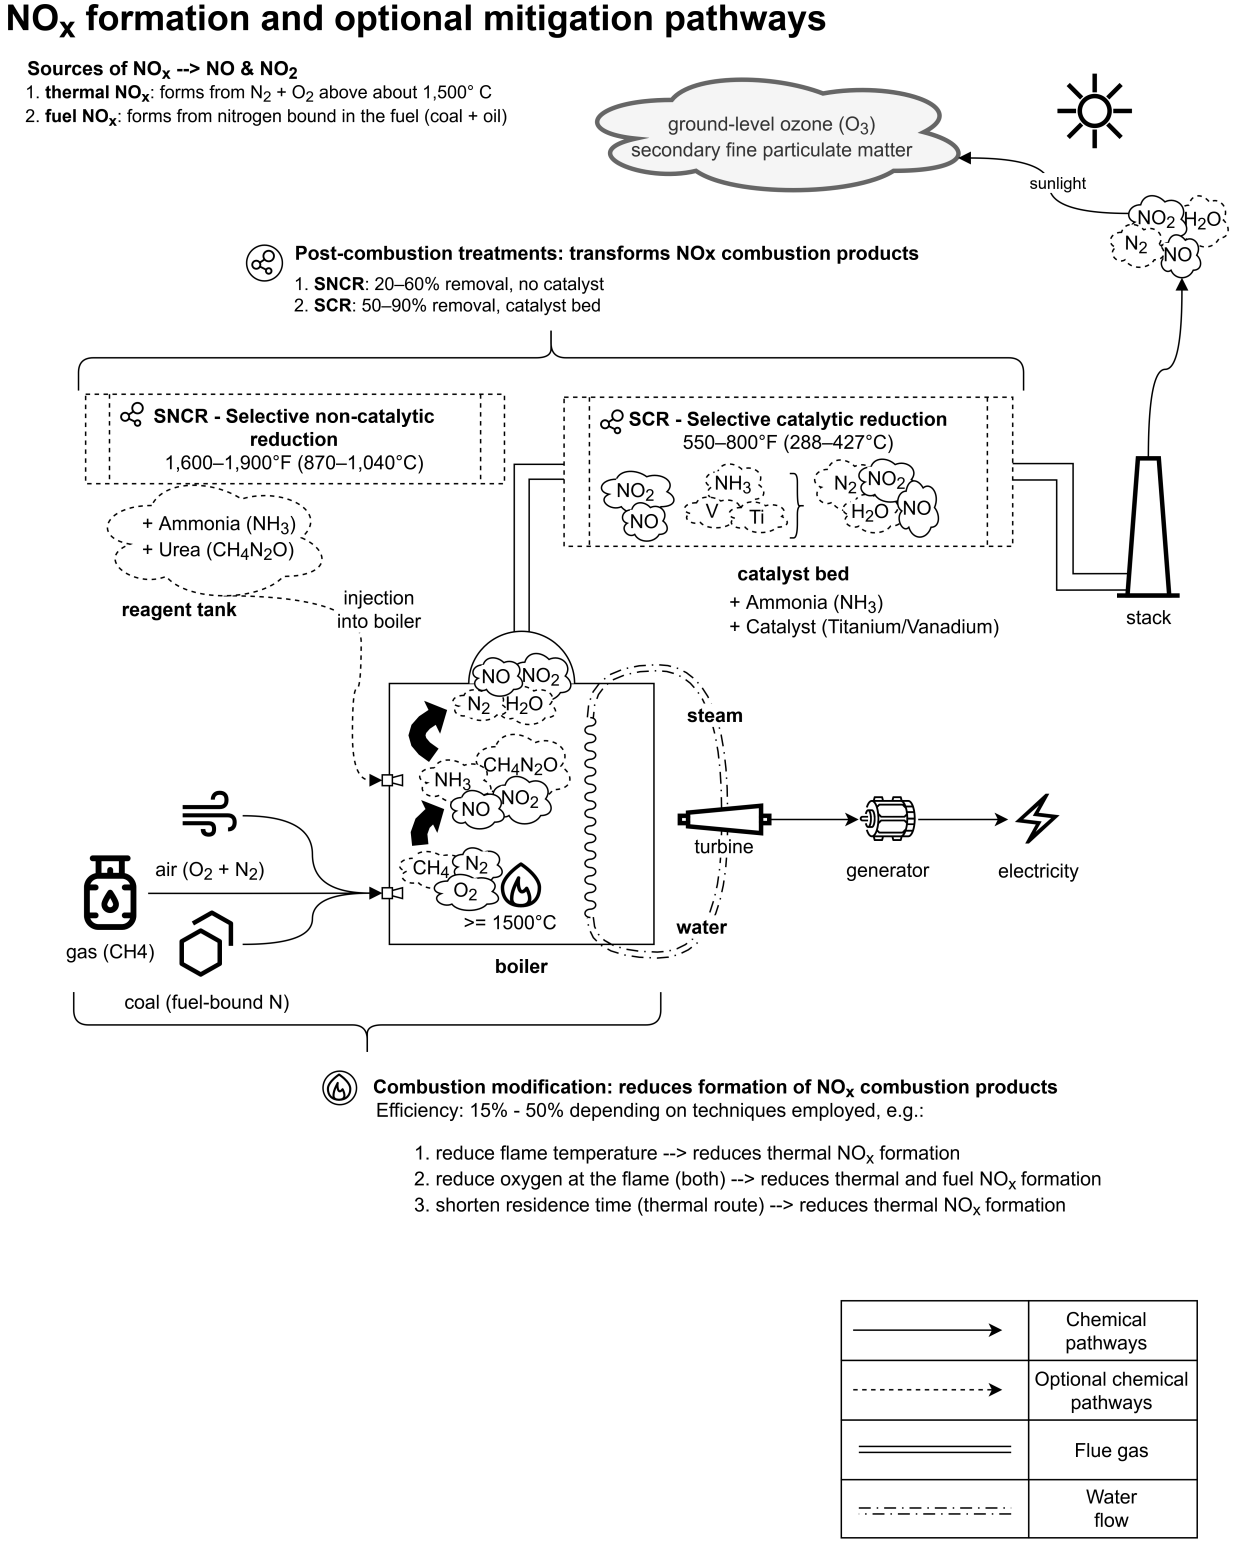

Figure 1.2: Where nitrogen oxides form in combustion, and where the two control margins act.

A plant's NO_x performance can be measured in two ways, depending on whether the object is the efficiency of its combustion or the efficiency of its conversion of fuel to electricity overall. The first is captured in pounds of NO_x per million British thermal units (MMBtu) of heat *input*, the unit in which regulatory limits are written: how much pollution each unit of burned fuel produces, whatever becomes of the energy. The second is captured in pounds per megawatt-hour (MWh) of electricity *generated*, which charges the plant for its conversion efficiency as well as its combustion; it is the rate studied throughout this dissertation, because operating behavior shows up in the efficiency term. The two are linked by the plant's heat rate, the quantity of fuel energy consumed for each unit of electricity produced, which is simply the inverse of thermal efficiency. The emission rate per megawatt-hour is the combustion rate scaled by the heat rate:

$$\underbrace{\frac{\text{lb NO}_x}{\text{MWh}}}_{\text{emission intensity}} = \underbrace{\frac{\text{lb NO}_x}{\text{MMBtu}}}_{\text{combustion rate}} \times \underbrace{\frac{\text{MMBtu}}{\text{MWh}}}_{\text{heat rate}}$$

so two plants burning equally cleanly can differ in what they emit per unit of useful output if one converts fuel to electricity less efficiently than the other. To make the scaling concrete: a perfectly efficient plant would consume 3.412 MMBtu for each megawatt-hour. A steam unit posting a heat rate of 10 MMBtu per MWh therefore converts about 34 percent of its fuel energy to electricity, and one burning at 0.15 lb of NO_x per MMBtu emits 1.5 lb per MWh. The same combustion at a combined-cycle heat rate of 7 (about 49 percent efficient) yields roughly 1.0 lb per MWh; and at the heat rate of 12 that the same steam unit might post at half load, it emits 1.8 lb per MWh.

This scaling has a consequence that matters for what follows. A plant running well below its design output is less efficient, so its heat rate worsens and every megawatt-hour carries more fuel with it. It also requires proportionally more excess air at low load, which raises the oxygen concentration in the hottest part of the flame and so increases thermal NO_x formation per unit of fuel burned. Where post-combustion control is installed, reduced load drives efficiency decreases in multiple ways. Selective catalytic reduction depends on the flue gas reaching its catalyst within a working temperature range, and low load can drive the gas beneath that threshold. Selective non-catalytic reduction is more exposed still: lacking a catalyst, it must inject its reagent wherever the gas inside the boiler happens to sit between roughly 870 and 1,040°C (1,600 to 1,900°F), and as load varies

that zone migrates relative to the fixed injectors. Too cool and the reaction fails to complete, leaving reagent to escape unreacted; too hot and the reagent is itself oxidized to nitric oxide, at which point removal efficiency turns negative (Figure 1.3). Starting and stopping compounds all three. The same physical plant, in other words, emits more per megawatt-hour when it runs briefly and intermittently than when it runs hard and steadily, which is why the duty cycle described in the previous section is not a detail of engineering but a determinant of harm.

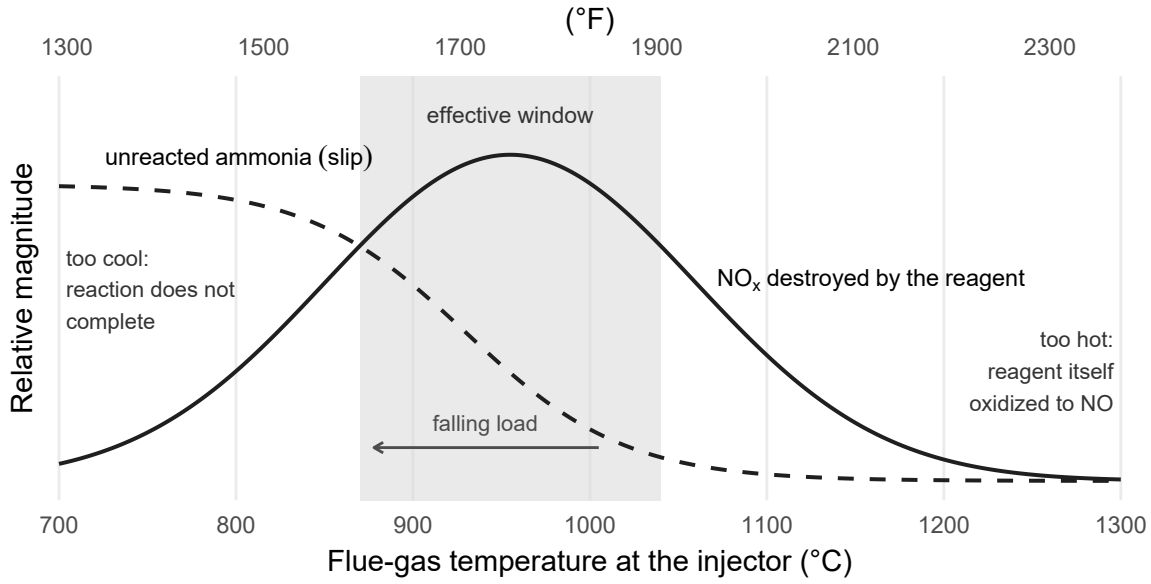

Figure 1.3: Schematic of the selective non-catalytic reduction temperature window: NO_x conversion peaks within a narrow band of flue-gas temperature, while unreacted ammonia escapes as slip below it. Vertical scale is relative, not a percentage. Window bounds after Richards (2000); form after Pronobis (2020).

1.4 Literature review

Environmental harm is not spread evenly across the facilities that produce it. Freudenburg's account of disproportionality holds that a small minority of sites accounts for a large share of the damage, so that industry-wide averages describe a distribution with a long tail rather than a representative middle (Freudenburg 2006). Greenberg develops the same observation for resource-based industries, where the concentration of harm tracks the concentration of production itself (Greenberg 2017). If harm concentrates in that way, then explaining it means explaining which facilities occupy the tail,

and that is a question about the organizations that own and operate them. Grant and colleagues have pursued it directly, relating plant-level emissions to the structure of the owning firm, its position within a corporate hierarchy, and its exposure to regulatory and market pressure (Grant et al. 2010; Grant, Jorgenson, and Longhofer 2018). Prechel and colleagues locate the same disproportionality in corporate form, finding that how a firm is organized bears on what it emits (Prechel and Istvan 2016).

Taken together, these studies compose what Grant and colleagues term an eco-structural account of pollution. Its premises are worth stating up front, since the analysis in Chapter 2 builds on them. There are three. First, the proximate agents of environmental harm are production organizations rather than individuals or technologies, because the same equipment can be operated at markedly different intensities by different owners; organizations, in Perrow's phrase, are the effective destroyers (Perrow 1997). Second, the structures through which those organizations are owned, financed, and regulated shape the decisions that determine harm, so that ownership and market position carry environmental consequences of their own (Grant et al. 2010; Prechel and Istvan 2016). Third, the harm so produced is concentrated rather than dispersed, which is what makes the identity of the heaviest emitters an explanatory problem rather than a matter of accounting (Freudenburg 2006). Disproportionality is the operative frame because it isolates the most consequential cases: where harm concentrates, the tails of the distribution, rather than its center, are the appropriate objects of explanatory analysis.

A second literature asks a different question of the same facilities: not which of them pollute most, but who lives beside them. Bullard's study of the American South established the pattern that hazardous facilities cluster near Black and low-income communities (Bullard 1990), and Mohai and Bryant found the same association in national data (Mohai and Bryant 1992). The result has been reproduced widely enough that Ringquist's meta-analysis of the accumulated studies could report a consistent, if modest, association between race and environmental hazard across the literature taken as a whole (Ringquist 2005). For electricity generation in particular, McCoy documents the coincidence of coal-fired plants with communities of color (McCoy 2017), Cushing and colleagues connect the location of fossil-fuel plants to the historical geography of redlining (Cushing et al. 2023), and Liévanos and colleagues follow the burden past the stack to the accumulation of coal

waste (R. S. Liévanos, Greenberg, and Wishart 2018).

What that coincidence means has been contested for three decades. The question, usually posed as which came first, is whether facilities were placed in communities that were already poor and disproportionately nonwhite, or whether the demographic composition of those communities changed after the facilities arrived. Been and Gupta asked whether neighborhood composition shifts in the years following a siting (Been and Gupta 1997); Pastor and colleagues put the same question to Los Angeles (Pastor, Sadd, and Hipp 2001); Mohai and Saha brought national longitudinal data to bear (Mohai and Saha 2015); and Banzhaf and colleagues set out what identifying either mechanism causally would actually require (Banzhaf, Ma, and Timmins 2019). The disagreement is not about whether the coincidence exists. It is about the process that produced it. The analyses that follow do not adjudicate it, and are not designed to: they take the distribution of plants and people as given at each point of observation and ask what happens next, in operation and in turnover.

A third body of work explains why the transition proceeds as it does. On this account it is distributive conflict, rather than any deficit of knowledge or technology, that governs the pace of decarbonization: incumbent owners of carbon-intensive assets face concentrated losses and organize accordingly, while the beneficiaries of transition are diffuse and organize less (Aklin and Mildenberger 2020; Meckling 2015). Mildenberger traces how that conflict is carried inside the institutions that make climate policy, so that opposition operates through the policymaking process rather than only against it (Mildenberger 2020). Colgan and colleagues frame the same politics as a contest over the revaluation of assets, in which what is at stake is the future worth of capital already in the ground (Colgan, Green, and Hale 2021), and Breetz and colleagues show that the political logic differs by technology and by stage of deployment (Breetz, Mildenberger, and Stokes 2018). These accounts explain the rate and direction of fleet turnover. They do not describe what turnover does to the communities living alongside it, which is the subject of Chapter 4.

Set against each other, the two principal literatures share a blind spot. The eco-structural tradition measures what a plant emits and relates it to who owns the plant, but it does not generally distinguish the equipment a plant carries from the way that equipment is run. The environmental justice tradition measures who lives near a plant and relates it to whether the plant is there at all, but it has largely treated the plant, once sited, as a fixed quantity of harm; a newer line of work follows the disparity

into the emissions themselves (Clark, Millet, and Marshall 2014; Cushing et al. 2023), and Chapter 3 continues in that direction. Both, in other words, are organized around what a facility *is* and where it *sits*. Neither has typically devoted much analytical attention to how the facilities themselves are *operated*, though the operating margin described in the previous section is where a substantial part of the variation in emissions actually lives. The analyses that follow occupy that opening. The argument here is middle-range (Lewis-Beck, Futing Liao, and Bryman 2026): it does not propose a general theory of environment and society, nor confine itself to a single case, but advances propositions at the level of the organization and the community that are specific enough to be tested against the observed behavior of an entire national fleet, and modest enough to be revised by it.

1.5 Questions and expectations by chapter

The three chapters analyze, in turn, the organizations that operate generating plants, the communities that host them, and the redrawing of the spatial demography of generation as the fleet turns over. Each builds on the one before it, and each is rooted in its own theoretical tradition.

1.5.1 The Organization of Pollution Control

Working within the eco-structural framework set out above, Chapter 2 asks three questions:

- **RQ1.** Does organizational structure predict plant-level NO_x emission intensity (lb/MWh), net of a plant's regulatory obligations, fuel mix, and capacity to emit?
- **RQ2.** Does organizational structure predict which NO_x controls a plant adopts, and does any gradient concentrate on the capital-intensive post-combustion margin rather than the inexpensive combustion one?
- **RQ3.** Does the organizational association with intensity run through adoption, or does it persist among plants carrying the same controls?

Organizational structure is theorized along five dimensions, each a hypothesis family: 1) organizational type, where government owners, rate-regulated investor-owned utilities, and cooperatives are predicted less NO_x-intensive than merchant generators, on the logic that cheap public capital

and rate-base recovery lower the cost of capital-intensive abatement (Prechel and Istvan 2016; Grant, Jorgenson, and Longhofer 2016); 2) ownership concentration, where dispersed ownership is predicted dirtier, reflecting weaker principal–agent monitoring of operators (Grant, Jorgenson, and Longhofer 2016); 3) cross-cutting status flags (leased operation and cogeneration designation), carried as exploratory; 4) market position, where dominant in-state firms are predicted dirtier; and 5) firm scale, the 'super-polluter' expectation that plants owned by larger firms pollute more intensively (Freudenburg 2006; Prechel and Istvan 2016).

1.5.2 Environmental Justice at the Operating Margin

Chapter 3 supplements that framework with an environmental justice perspective. It asks three questions, mirroring Chapter 2's with the primary variable swapped:

- **RQ1.** Does neighboring-community demographic composition predict plant-level NO_x intensity, net of organizational structure?
- **RQ2.** Does community composition predict which NO_x controls a plant adopts, and does any gradient concentrate on the capital-intensive post-combustion margin rather than the inexpensive combustion one?
- **RQ3.** Does the demographic association with intensity run through adoption, or does it persist among plants carrying the same controls?

Composition enters as percent nonwhite, poverty, and income, measured by population-weighted apportionment within a 3 km buffer of each plant. The dissertation derives the EJ literature's theoretical expectation for this context: that plants near nonwhite and low-income communities underinvest in abatement, reflecting differential regulatory enforcement and reduced community capacity to monitor and object (Bullard 1990; Been 1994), with the shortfall concentrated on the capital-intensive post-combustion margin. The two-margin design should, ideally, adjudicate the mechanism. Where a demographic association with intensity attenuates after conditioning on adoption, abatement decisions are shaped by demographics (demographics pattern who abates, and who abates sets emission intensity). Where the demographic association survives, demographic context would operate, instead, *downstream* of adoption. To make the alternatives concrete: were

plants in more nonwhite communities found to carry post-combustion control less often, and were the demographic coefficient on emission intensity to fall toward zero once that control is accounted for, the disparity would be one of equipment, and it would be addressed by what gets installed. Were those plants instead found to carry control at the same rate, and were the demographic coefficient to survive that accounting undiminished, the disparity would lie in how identically equipped plants are run, and it would be addressed in operation and enforcement. Throughout, race and class enter separately, so the racial association is read net of income and poverty (Mohai, Pellow, and Roberts 2009).

1.5.3 The Unequal Geography of the Green Transition

Chapter 4 extends the environmental justice argument beyond the fossil fleet, engaging the literatures of energy geography and of the political economy of just transition. It asks where the transition is unfolding, in three questions:

- **RQ1 (trajectory).** How do the demographic compositions of energy-type host communities evolve over 2010–2023, net of the national trend?
- **RQ2 (decomposition).** How much of each fuel's demographic change is *fleet turnover* (entry/exit) versus *change around persisting plants*?
- **RQ3 (entrant composition).** How does the community composition surrounding newly built capacity differ from that surrounding incumbent capacity, by fuel, and does that difference arise within states or between them?

Because its design is descriptive, Chapter 4 states directional expectations (fixed before estimation) rather than hypotheses under test. The expectations are organized by a trichotomy: under 1) replication, the fossil path of least resistance operates on clean capacity as it did on dirty; under 2) transformation, technoeconomic and land-use logics dominate and the demographic association weakens or shifts its axis from race toward rurality and class; under 3) the 'null,' resource potential swamps demography altogether (Cranmer et al. 2023; O'Shaughnessy et al. 2023; Mueller and Brooks 2020). Each limb implies a different pattern in the results. Replication would appear as new wind and solar capacity coming to rest in communities whose demographic profile matches that of

the fossil fleet it displaces. Transformation would appear as clean entrants landing in communities that are whiter and more rural but no less poor, so that the gradient persists on class while weakening on race. The null would appear as entrant composition tracking wind speed, solar irradiance, and transmission access, leaving no demographic pattern once those are accounted for. Different fuels are expected to fall under different limbs, with new gas replicating the fossil gradient, wind transforming it, solar sitting between, and coal's change running through exit as the decade's retirements thin its fleet (Richmond-Bryant et al. 2020).

1.6 Design and data

Table 1.1 compares the three empirical designs.

Table 1.1: The three empirical chapters compared.

	Chapter 2	**Chapter 3**	**Chapter 4**
Core contrast	Does organizational structure predict control adoption, or operation?	Do host community demographics predict adoption, or operation?	How do entry, exit, and neighborhood change redraw the spatial demography of generation?
Key explanatory variable	Organizational structure	Community demographics, 3 km buffer	Energy type, by entry and exit
Outcome	NO_x emission intensity (lb/MWh)	NO_x emission intensity (lb/MWh)	Community demographics, 3 km buffer
Sample	Coal and gas $\geq$ 25 MW, CEMS-matched	Same fleet, with ACS buffers	Full fleet, renewables included
Frame	2009 / 2018 / 2021 snapshots	2009 / 2018 / 2021 snapshots	Annual panel, 2010–2023

Chapters 2 and 3 share a functional two-margin abatement design (Table 1.1). The design relates a primary independent variable to plant-level NO_x emission intensity and examines that relationship across the two functional margins of pollution control a plant may adopt: combustion modification, which suppresses NO_x formation at the source and is relatively cheap, and post-combustion treatment, which scrubs the flue gas after the fact and is capital-intensive. It then asks whether the variable's association with emission intensity survives conditioning on those adoption decisions. In other

words, each of the two chapters separates two stories that are usually told as one: a story about who buys control equipment, and a story about who runs dirtier given the equipment they have. Because the two margins differ sharply in capital intensity, the contrast between them also does diagnostic work. It distinguishes a cost-of-capital mechanism from one operating at the combustion source regardless of capital cost. The primary independent variable in Chapter 2 is organizational structure; in Chapter 3 it is community demographics, with Chapter 2's organizational dimensions carried forward as controls.

Both chapters observe the coal- and gas-fired fleet at three cross-sectional snapshots (2009, 2018, and 2021) assembled from the Environmental Protection Agency's (EPA) Emissions & Generation Resource Integrated Database (eGRID), EPA Clean Air Markets Program facility attributes, Energy Information Administration Forms 860 and 923, and the American Community Survey (ACS). The snapshot years are a design choice: they bracket three distinct policy regimes (pre-Clean Power Plan, mature post-CPP, and post-pandemic). The chapters pool them with year fixed effects and utility-clustered standard errors, rather than estimating each year separately. The fleet's composition shifts across the three snapshots as plants enter and retire; the pooled sample retains that shift, because the object of each analysis is the fleet as it stood at each snapshot, and restricting to plants observed in all three years would condition the sample on survival. The pooled sample covers coal- and natural gas-fired plants of at least 25 MW matched to EPA continuous emissions monitoring in the continental United States: 1,666, 1,622, and 1,530 plants in the three snapshots respectively, 4,818 plant-years in all.

Chapter 4 sets the estimating-equation apparatus aside for a descriptive design applied to the full electricity generating fleet, renewables included. The panel is annual, 112,782 plant-years across 13,063 plants, 2010 through 2023, built from EIA-860 and the ACS. The variable roles reverse relative to Chapters 2 and 3: community demographics become the outcome and energy type the predictor. Three descriptive analyses carry the design. The first, a trajectory analysis, traces each energy type's mean host-community composition year by year, read as divergence from the national trend. The second, a decomposition, asks what moves those trajectories, partitioning each fuel's change into three channels: the entry of new plants, the exit of retiring ones, and demographic drift in the neighborhoods around plants that persist. The third, an entrant-composition regression, asks

where the change happens, splitting each fuel's entrant-versus-incumbent gap into its within-state and between-state parts.

How the community around a plant is drawn is itself a substantive decision: Mohai and Saha showed that conclusions about environmental inequality vary with how proximity to a hazard is measured, so that the same facilities and the same population can yield different findings under different definitions of the surrounding community (Mohai and Saha 2006). Chapters 3 and 4 accordingly state their spatial rule explicitly, report it at several radii, and treat the sensitivity of the findings to that choice as a result in its own right.

A caveat attends the demographic measures. The ACS is a sample survey, and its tract-level estimates carry sampling error. The established measure of that uncertainty is the coefficient of variation, the standard error of an estimate divided by the estimate itself, calculable directly from the ACS's published margins of error; the reliability literature recommends screening out tract estimates whose coefficient of variation exceeds 0.50 (Folch et al. 2016; Raoul S. Liévanos 2018, 2019). The analyses here apply no such screen, and the estimates should be read with that limitation in view. The resulting uncertainty is not evenly distributed: documented reliability problems concentrate in the southern United States and in low-income communities and communities of color, so the demographic estimates for precisely those communities may be the most affected.

The decennial censuses of 2010 and 2020 measure race and ethnicity with full-count precision, but they do not measure income or poverty (the long form having been retired after 2000), and their two enumeration dates align with neither the 2009/2018/2021 snapshots nor the annual panel; the ACS is therefore used throughout, with a decennial-based sensitivity analysis reserved for future work.

1.7 Plan of the dissertation

The three empirical chapters follow, each self-contained, with its own data, methods, and discussion, and each readable on its own. They are built from a shared backbone of federal data, described in Section 1.6, that each chapter extends in a different direction: toward the organizations that run the fleet, the communities that host it, and the geography of what replaces it. A concluding

chapter synthesizes the findings across the three, draws out their policy implications, and gathers the limitations and extensions that the chapters treat separately.

2 The Organization of Pollution Control

Organizational Structure, Environmental Control Technology Adoption, and NOx Emissions

Abstract

Do the organizational structures of power-plant owners shape how much pollution a plant emits per unit of electricity? And if so, does that influence work through which control equipment a plant installs, or through how cleanly it runs the equipment it already has? This chapter examines NO_x emission intensity across the U.S. coal- and gas-fired fleet of at least 25 MW, observed at three snapshots (2009, 2018, and 2021) and assembled from EIA Forms 860 and 923, EPA eGRID, and the EPA CAMD facility-attributes file. It models NO_x control as two fuel-agnostic functional margins, combustion modification and capital-intensive post-combustion treatment, and estimates a four-equation system with standard errors clustered on the utility; the gap between a plant's adoption on the two margins serves as a test for whether organizational influence runs through the cost of capital. Organizational structure barely predicts whether a plant adopts control on either margin, the lone exception being government owners, who favor the capital-intensive option. Where structure leaves its mark is the emission outcome, downstream of adoption: regulated investor-owned utilities and cooperatives run markedly dirtier than merchant generators, a penalty that survives conditioning on adoption and points to the weak operating-cost discipline of cost-of-service regulation rather than to the financing of equipment. Firm scale, the super-polluter expectation, carries no such intensity penalty. The chapter's contribution is to relocate where organizational structure leaves its mark, from the decision to install control to how plants are run.

2.1 Introduction

Fossil-fueled electricity generation is one of the largest sources of greenhouse gas and air pollutant emissions in the world (Ge, Friedrich, and Vigna 2026), and scholarly and policy attention to climate change has long concentrated on it accordingly. A useful organizing concept in this work is that of 'disproportionality' (Freudenburg 2006; Greenberg 2017). The frame of disproportionality directs our attention to the fact that a small set of sectors and technologies produces most of the emissions, and scholarly, policy, and activism effort concentrates accordingly. But this work often buries the lede by stopping at the sector level. Within the sector, a small fraction of fossil-fueled facilities produces most of the emissions; the rest contribute modestly. These are the 'super-polluters' (Grant, Longhofer, and Jorgenson 2024; InfluenceMap 2024). But super-polluters do not pollute as abstract entities. They pollute through organizations—"the most intensive and effective environmental destroyers," in Perrow's phrase (Perrow 1997, 66)—and at specific plants, with equipment, ownership structures, and regulatory contexts that shape what gets emitted and how much (Grant, Jorgenson, and Longhofer 2018). Firm scale is in fact only one dimension along which organizational structure varies. Ownership form, regulatory status, market position, and ownership concentration are also expected to shape how much a given plant pollutes, and they need not move together with size.

This chapter asks which features of organizational structure shape NO_x emission intensity in the U.S. electricity sector, and through what mechanism: whether the association runs through a plant's adoption of control equipment or persists among plants already carrying it.

To answer this question, the chapter draws on multiple relevant traditions developed largely in parallel.

Ecostructuralism explains environmental pollution as a product of the relationship between industrial facilities and the communities that host them, arguing that macro-level theories of capitalism and micro-level theories of consumption alone cannot account for emissions at the point of production (examples of this work include Grant et al. (2010); Prechel and Touche (2014); Grant, Jorgenson, and Longhofer (2016); Grant, Longhofer, and Jorgenson (2024)).

Organizational political economy draws on resource dependence theory, which explains organizational and inter-organizational behavior in terms of the critical resources an organization must

secure in order to survive and function: organizations become dependent on the entities that control those resources, the asymmetric exchange creates power relations through which external actors can constrain organizational behavior, and organizational leaders deploy strategies to manage those dependencies and preserve autonomy (B. L. Johnson (1995); cf. Wallerstein (2004) for an application to states and their political leaders in the global economy). In this tradition, organizational political economy examines how corporate structures, political embeddedness, and accumulation strategies generate differential sets of incentives and constraints, leading to strategic decisions to produce or control pollution (Meckling (2015); Paterson and P-Laberge (2018); Green et al. (2021); Colgan, Green, and Hale (2021); Prechel (2023)).

Comparative political economy (and specifically, sub-literature(s) of comparative environmental policy and comparative climate and energy policy) situates those organizational decisions within their institutional environment, showing how state-level institutions shape the regulatory, market, and political circumstances under which firms operate (cf. Hall and Soskice (2001); Oates and Portney (2003); Newell, Pattberg, and Schroeder (2012); Aklin and Urpelainen (2013); Seto et al. (2016); S. E. Kim and Urpelainen (2017)). The state-firm(-state) relationship is reciprocal: many of these are distributive accounts in which firms and the interest groups they form deploy their power resources to shape regulation and policy in their own favor, in the spirit of open-economy-politics models of preference aggregation, so the institutional environment is as much a product of organized firm power as a constraint upon it (Brulle (2018); Breetz, Mildenberger, and Stokes (2018); Aklin and Mildenberger (2020); Mildenberger (2020); Christophers (2022)). These institutional differences are what the regulatory-status and market-position measures capture.

Ecostructuralism and organizational political economy (specifically, the environmental inequality research that applies an organizational lens) both take environmental pollution as their dependent variable, yet they have only recently been brought into integrated frameworks. Following Prechel, Blazek, and Amaral (2025), this chapter draws on both at once:

> … Organizational behavior cannot be understood without examining the relationships among the component parts of the social structure and corporate polluting behavior because these relationships create dependencies, incentives, and opportunities that either constrain or advance organizational goals. (Prechel, Blazek, and Amaral 2025)

Taken together, these traditions define the framework this chapter builds and tests: an eco-structural organizational political economy of plant-level pollution, specified with the detail and precision afforded by integration of the comparative political-economic and institutionalist traditions. Ecostructuralism locates pollution at the point of production, in the relationship between a facility and its host community. Organizational political economy supplies the firm-level engine: corporate structure, resource dependence, and accumulation strategy generate incentives and constraints driving plants' pollution decisions. The comparative traditions afford an account of institutional form, function, design, and transformation, which gives greater analytical content to the measurements of institutional dimensions gathered here. These institutional distinctions enter the model directly, as the regulatory-status and market-position dimensions of organizational structure and as the political-economic context in which each plant operates, while the movement across the three regulatory regimes the analysis observes (2009, 2018, 2021) is where institutional transformation enters the design. Environmental justice and environmental inequality scholarship (cf. Pellow (2000); Downey (2005); Freudenburg (2006)) supplies the disproportionality frame that motivates the inquiry; its distributional question, who bears the resulting burden, is taken up in a later chapter.

The analysis operationalizes this consolidated framework at the plant level. The analysis uses plant-level NO_x intensity as the outcome and retains the organizational variables (ownership form, regulatory status, market position, firm scale, ownership concentration) that the organizational political economy tradition treats as primary explanatory dimensions.

The central argument of this chapter follows Grant et al. (2010) and Grant, Jorgenson, and Longhofer (2018) in its contention that organizational structure shapes plant-level emissions. This chapter aims to expand on these studies by asking through which part of the NO_x abatement pathway that influence runs. It proceeds in two moves. First, it separates abatement into two functions that differ in cost: combustion modification, which is comparatively inexpensive, and post-combustion treatment, which is capital-intensive. Comparing how organizational structure relates to adoption on each margin distinguishes a cost-of-capital mechanism from a broader ecostructural one. Second, it asks whether the association between organizational structure and emission intensity survives conditioning on those adoption decisions: an association that attenuates once adoption is held constant operates largely through the adoption choice, while one that persists points to channels

beyond it.[1]

The analysis is restricted to coal-fired and natural gas-fired plants in the continental United States. This restriction focuses the analysis on plants where both the dependent variable and the technology adoption pathway are analytically meaningful: fossil fuel generators with real emissions variation and real NO_x abatement decisions. Pure renewable plants are excluded; without fossil generation they have no abatement technology pathway to analyze and no NO_x outcomes to explain. Nuclear merits its own sentence rather than a place in a residual list: fission raises steam without combustion, so a nuclear plant produces no stack NO_x and has no abatement pathway to model, and its regulatory and organizational context, from licensing to duty cycle, differs from the fossil fleet's in ways the organizational measures here are not built to capture. The remaining types (hydro, biomass, geothermal, oil) are excluded because hydro and geothermal combust no fuel, while the biomass and oil fleets are small and in most cases carry no NO_x emissions variable in the underlying data, so there is no outcome to model in any event. Chapter 4, whose question concerns the full fleet's geography rather than combustion abatement, includes nuclear and all of these types.

Three cross-sectional snapshots anchor the analysis: 2009, 2018, and 2021. These years are chosen deliberately. The 2009 snapshot falls in the first full year of EIA Schedule 8 environmental operations data, sits under the Clean Air Interstate Rule (CAIR), and precedes the shale-driven coal-to-gas switching cycle. The 2018 snapshot captures the matured Cross-State Air Pollution Rule (CSAPR) regime, after the post-combustion retrofit wave driven by that rule had largely run its course. The 2021 snapshot marks the post-pandemic (and pre-Bipartisan Infrastructure Law) regulatory and investment landscape. The three-panel structure also aligns with future American Community Survey 5-year windows,[2] enabling a subsequent demographics chapter to be layered onto the same plant-year base.

[1] A further channel is how intensively installed equipment is then operated: the number of hours it runs and how close it performs to its tested design efficiency. A plant that has installed an SCR unit but runs it infrequently, or well below its design removal efficiency, emits substantially more than one running identical equipment at full capacity. Because that utilization measure is defined only for boiler-equipped plants, it is examined in a coal-boiler deep-dive rather than across the full fleet.

[2] ACS 5-year estimates centered on each snapshot year: 2007–2011 for 2009; 2016–2020 for 2018; 2019–2023 for 2021.

2.1.1 Research Questions

- **RQ1.** Does organizational structure predict plant-level NO_x emission intensity (lb/MWh), net of a plant's regulatory obligations, fuel mix, and capacity to emit?
- **RQ2.** Does organizational structure predict which NO_x controls a plant adopts, and does any gradient concentrate on the capital-intensive post-combustion margin rather than the inexpensive combustion one?
- **RQ3.** Does the organizational association with intensity run through adoption, or does it persist among plants carrying the same controls?

Each question maps onto the four-equation system specified below: RQ1 onto Equation 3a, RQ2 onto Equations 1a and 1b, and RQ3 onto Equation 3a against Equation 3b.

2.1.2 Research Design

2.1.2.1 Unit of Analysis

This is a plant-level panel data analysis. The unit of analysis is the electricity generating facility (hence, 'plant'-level analysis).

2.1.2.2 Scope of Analysis

The geographical scope covers the continental United States (excluding Hawaii, Alaska, and Puerto Rico). The analytical sample is restricted to coal-fired and natural gas-fired plants.

The temporal scope consists of three cross-sectional snapshots: 2009, 2018, and 2021.

Two considerations inform the design of the snapshot. First, it deliberately brackets change over time as a sequence of distinct regulatory and market regimes, rather than a data-availability fallback: 2009 observes the fleet under CAIR, before CSAPR and the post-combustion retrofit wave it drove; 2018 observes the matured CSAPR regime (effective 2015), with the bulk of that retrofit installation complete; and 2021 observes the post-pandemic, pre–Inflation Reduction Act phase. Second, three snapshots give a clear before, middle, and after trajectory while staying parsimonious and aligning with ACS five-year windows for the demographic analyses in later chapters; additional years would add regulatory noise without sharpening the regime contrast. The related choice to

pool these snapshots with year fixed effects rather than estimate an annual panel is defended in Section 2.2.2.7, and is not relitigated here.

2.1.2.3 Dependent Variables

The primary outcome is NO_x emission intensity (`nox_rate_lb_per_mwh`, lb/MWh), derived from US Environmental Protection Agency (EPA) Clean Air Markets Program Data (CAMPD) reported fields in their Emissions and Generation Resource Integrated Database (eGRID).[3] An intensity outcome is used rather than total emitted mass so that plants of very different size remain comparable: lb/MWh captures how cleanly a plant generates electricity, net of how much it generates, whereas a count of total NO_x tons conflates emission performance with generation scale and would rise mechanically with output even when a plant controls its emissions well. This keeps the dependent variable aligned with the abatement-adoption questions at the center of this analysis, which concern whether and how a plant controls its emissions, not how much electricity it produces. Total NO_x mass (ln-transformed) is retained as a robustness outcome.

2.1.2.4 Primary Independent Variable: Organizational Structure

Organizational structure is measured along five dimensions, each corresponding to a hypothesis family in Section 2.1.3. The central dimension is *organizational type*, a single mutually exclusive classification (`org_structure_type`) that assigns each plant to one of five economic-role categories: market-disciplined merchant generator (the reference), rate-regulated investor-owned utility, cooperative, government or public power, and industrial or commercial host-serving generator. It is built from EIA-860 entity type, generation sector, and rate-regulation status (Table 2.1), and it carries the ownership-and-regulatory regime on which the cost-of-capital, operating-cost-discipline, and eco-structural mechanisms turn. *Ownership concentration* (`n_owners`) is whether control rests with a single entity or is dispersed across multiple partial owners, including joint plants held by

[3] "The Emissions & Generation Resource Integrated Database (eGRID) is a comprehensive source of data from EPA's Clean Air Power Sector Programs on the environmental characteristics of almost all electric power generated in the United States. The data includes emissions, emission rates, generation, heat input, resource mix, and many other attributes. eGRID is typically used for greenhouse gas registries and inventories, carbon footprints for electricity purchases, consumer information disclosure, emission inventories and standards, power market changes, and avoided emission estimates" (eGRID website).

two or more utilities. *Market position* is the parent firm's share of its state's electrical capacity and output, following the dominant-utility framework of Grant, Jorgenson, and Longhofer (2016). *Firm scale* is the absolute size of the parent firm's portfolio in plants and nameplate capacity, distinct from its state-level share. Because the two market-position measures and the two firm-scale measures are each strongly correlated within their dimension, they enter as single standardized composite indices, `state_market_share_index` and `firm_scale_index` (detailed in Section 2.1.3), with the underlying measures retained for descriptive and robustness use. Finally, two *cross-cutting status flags* mark arrangements that overlay the org-type taxonomy rather than partitioning it, attaching to plants of any type: *third-party (leased) operation*, where a lessor owns the asset and a separate lessee runs it, and *PURPA cogeneration*, a Qualifying-Facility certification carried by combined-heat-and-power plants.[4] Neither carries a directional prediction. Together these five dimensions operationalize the incentives and abilities the framework predicts will shape differences in how plants control their emissions.

Table 2.1: Organizational-type typology and construction logic: each plant is sorted by NAICS economic sector, then rate regulation, then ownership (plant-years pooled over 2009, 2018, 2021).

1. Economic sector	2. Rate regulation	3. Ownership	→ Organizational type	Plant-years	Share
Industrial / commercial (on-site)	—	—	Host-serving	757	15.7%
Wholesale / grid-connected	Non-regulated (NR)	—	Merchant (reference)	1,848	38.4%
	Rate-regulated (RE)	Investor-owned / IPP	Regulated IOU	1,229	25.5%
		Cooperative	Cooperative	315	6.5%

[4]The Federal Energy Regulatory Commission (FERC) certifies generators under two federal statutes that define categories of nonutility generation. The Public Utility Regulatory Policies Act of 1978 (PURPA) created Qualifying Facility status for cogenerators (plants producing electricity and useful thermal energy together) and for small power producers (generally 80 MW or less, primarily renewable), obligating utilities to purchase their output at avoided cost; these enter the analysis as the cross-cutting binaries `ferc_cogen_binary` and `ferc_small_power_binary` (EIA-860 Schedule 2). The Energy Policy Act of 1992 (EPACT) created Exempt Wholesale Generator status, the legal vehicle of the modern merchant-IPP industry; rather than a separate indicator, it is subsumed in the organizational-type factor, where non-regulated wholesale generators constitute the Merchant reference category. In this coal- and gas-fired sample the PURPA flags fall almost entirely on cogeneration plants, since small power producers are predominantly renewable and out of scope by fuel.

	Government	Government	669	13.9%
All plant-years			**4,818**	**100.0%**

Note. First matching rule wins. **Definitions** — Merchant: non-regulated wholesale generator (market-disciplined; the reference). Regulated IOU: rate-regulated investor-owned or independent utility. Cooperative: rate-regulated member-owned utility. Government: rate-regulated public power (federal, state, municipal). Host-serving: industrial or commercial on-site generation. *Regulated* means economic rate regulation (cost-of-service ratemaking), not environmental regulation, which applies to all plants. Built from EIA-860 sector, regulatory_status, and entity type.

2.1.2.5 Technology Variable: NO_x Abatement Technology

The primary technology variable classifies NO_x control by *function* rather than by specific technology, which makes it fuel-agnostic. This fuel-agnosticism is what permits cross-fuel comparison of how organizational structure relates to control adoption and emission intensity. Following the standard taxonomy, NO_x controls divide into two functional families. Combustion modification technologies seek to prevent NO_x from forming in the combustion zone and include, in the case of boilers, low-NO_x burners, overfire air, lean combustion, fuel reburning, and flue gas recirculation; in the case of gas turbines, technologies include dry-low-NO_x combustors and water/steam injection. Separately, post-combustion treatment, the two chief technologies being selective catalytic reduction (SCR) and selective non-catalytic reduction (SNCR), removes NO_x from the flue gas after it forms. The two families differ in the capital intensity of their adoption (post-combustion treatment, SCR especially, is far more capital-intensive) and in their removal efficiency (SCR 80–90%, SNCR 25–50%, combustion modification 15–25%; Richards (2000)).

The NO_x abatement technology variable is operationalized as two binaries, `has_combustion_control` and `has_postcombustion_control`, each equal to 1 if the plant operates any control in that family (from EIA-860 boiler strategy records and the EPA CAMD facility-attributes file, the latter supplying the gas-turbine controls EIA does not record). They are modeled as two separate adoption decisions (Equations 1a and 1b) for several reasons. The two are not mutually exclusive and are functionally distinct: combustion modification prevents NO_x from forming, whereas post-combustion treatment removes it after the fact, and a plant may hold any combination of the two, so a single pooled (any-control) binary or an ordered tier would impose false structure on them. They also differ in how the decision is made: combustion modification is often an incremental retrofit

or operational adjustment, whereas post-combustion is a discrete, lumpy capital investment, so the organizational covariates may load differently on each margin. Collapsing the two would net out associations that can run in opposite directions and obscure the mechanism, while modeling them separately isolates the cost-of-capital contrast (Equation 1b versus 1a): post-combustion is far more capital-intensive, so the gap between organizational structure's effect on the two margins is what carries the cost-of-capital signal. Finally, because the families differ so greatly in removal efficiency, entering both indicators separately is necessary for the conditioning step (Equation 3b) to be correctly specified. The functional classification applies identically to coal boilers and gas turbines, so the primary analysis spans the full fossil-fired fleet.[5]

2.1.2.6 Control Variables

Control variables are organized under three conceptual banners.

Capacity to emit captures the baseline physical conditions that determine a plant's emissions potential independently of abatement choices: plant scale (`log_nameplate_cap_mw`), utilization (`capacity_factor`), generator age (`avg_generator_age_yrs`), prime-mover configuration (`has_steam_turbine`, `has_gas_turbine`), and fuel-mix composition (`coal_heat_share`). The fuel-mix share holds combustion composition constant across coal-, gas-, and hybrid plants, so that variation in NO_x intensity that merely reflects burning a different fuel is not absorbed into the organizational-structure coefficients; one share suffices, because heat shares sum to ≤ 1 with a negligible oil-and-other residual and, with roughly 96% of the sample burning only coal or only gas, `coal_heat_share` and `gas_heat_share` are near-perfectly collinear (coal + gas ≈ 1), so `gas_heat_share` is omitted. Plant throughput and efficiency are deliberately left out: total heat input (`log_total_heat_mmbtu_923`) is redundant with scale and utilization, and heat rate (`heat_rate_mmbtu_per_mwh`) is undefined or unstable for the many low-utilization gas peaking units in the sample and is excluded on data-availability grounds. Boiler firing type (cyclone, wall, tangential, fluidized-bed) is a further baseline determinant of the unabated NO_x formation rate, distinct from the combustion-*modification* controls of Equation 1a; it is defined only for boiler-equipped units,

[5]The within-adopter, boiler-specific questions (for example, the SCR-versus-SNCR choice and how intensively post-combustion equipment is operated, measured as average annual hours in service per control unit, `nox_control_in_service_hours`) are not defined for simple-cycle gas turbines and are analyzed separately in the coal-boiler deep-dive.

however, and is structurally absent for the simple-cycle and combined-cycle gas turbines that make up much of the fleet, so it enters the coal-boiler deep-dive rather than the full-fleet primary, with the consequences of omitting it assessed in Section 2.2.2.7.

Regulatory controls in the primary are the state-level policy environment: emissions caps, GHG reduction targets, renewable portfolio standards, carbon pricing, public benefit funds, and electric decoupling from the state policy panel, all of which apply to every plant. A plant-level compliance-vintage measure, `years_under_nox_regulation` (how long each plant has been subject to a NO_x standard, from EIA-860 boiler records), is available, but like firing type it is boiler-keyed and structurally undefined for the simple-cycle gas turbines that make up much of the fleet, so it is held out of the full-fleet primary and used only in a sensitivity check. These state- and plant-level controls form a natural hierarchy: plants sit within utilities, which sit within states. The analysis does not fold this hierarchy into a multilevel (random-effects) model, for three reasons. First, such models assume the state and utility differences are unrelated to the predictors, yet organizational structure, the main variable of interest, is likely to be related to those differences, which biases the results. Second, the usual reason to use one, that plants owned by the same utility resemble each other, is handled here by clustering the standard errors at the utility level. Finally, because the state and utility levels enter as ordinary variables rather than as random effects, each coefficient is a single estimate from the whole sample. A random-effects model would instead report the state and utility pieces as partially pooled estimates, each of which would be pulled toward the overall average (grand mean) by an amount that depends on how few plants that state or utility has, which is harder to interpret cleanly. In sum, and for the sake of interpretability, the research design treats these hierarchical elements as optional rather than necessary: the within-utility resemblance they would absorb is handled parsimoniously by clustering the standard errors at the utility level, the coefficients stay straightforward to interpret, and, most importantly, the bias that random effects would invite when the predictors are tied to state and utility differences is sidestepped entirely.

Political-economic context controls capture the state-level political and market environment within which a plant operates. The political controls come first. A state partisan-lean index (`dem_lean_index`, the standardized mean of the governor's party and the Democratic seat share in each legislative chamber) absorbs cross-state heterogeneity in environmental rule-setting and, in

particular, enforcement priorities. This is where the institutional reciprocity (both "firms-as-product-of-government" and "government-as-a-product-of-firms") noted above becomes concrete: energy interests (across the fuel-type and energy type spectrum, including renewables) use politically favorable environments to soften the regulation and enforcement they face (Brulle (2018); Aklin and Mildenberger (2020); Mildenberger (2020)), so a plant in a more Republican-leaning state may operate under weaker pressure to control its emissions than an identical plant elsewhere. Holding partisan environment constant strips differential enforcement pressure out of organizational-structure coefficients. Because firm power can also help shape that environment, however (the largest firms most of all), the partisan-lean index conceivably absorbs part of any organizational effect that runs through political influence. This actually makes the organizational estimates conservative, though; moreover, no causal interpretation is attached to the partisan-lean coefficient itself. The economic control follows: wholesale market structure enters via ISO/RTO membership[6] (`in_iso_rto`), capturing dispatch-discipline and capacity-market[7] effects that operate independently of the organizational-structure hypotheses. ISO/RTO membership is treated as a control rather than a hypothesis for two reasons: empirically, its identification is heavily confounded with geographic and state-level political-economic differences; and theoretically, its mechanism is mixed: dispatch discipline plausibly predicts cleaner operation, whereas capacity-market inertia, by keeping older peaking units in service, predicts dirtier. The variable is included only so that its variation is held constant. There is no causal interpretation attached to its coefficient.

Several of these controls overlap with the organizational measures and with one another (regulated utilities cluster outside ISO/RTO regions (Fowlie (2010)); large firms tend to hold large state shares),

[6] An Independent System Operator (ISO) or Regional Transmission Organization (RTO) is a FERC-regulated entity that operates the high-voltage transmission grid and runs centralized wholesale electricity markets across a multi-state region, independently of any single utility. The two are functionally similar; RTOs, established under FERC Order No. 2000 (1999), carry somewhat broader regional responsibilities. In ISO/RTO territory, generators are dispatched through competitive energy markets (and often capacity markets) rather than by a vertically integrated utility, which is the dispatch-discipline channel referenced here. Major U.S. ISOs and RTOs include PJM Interconnection (PJM), the Midcontinent Independent System Operator (MISO), the California Independent System Operator (CAISO), the Electric Reliability Council of Texas (ERCOT), ISO New England (ISO-NE), the New York Independent System Operator (NYISO), and the Southwest Power Pool (SPP); `in_iso_rto` equals 1 if a plant lies within one of these footprints (a non-empty eGRID `ISORTO` field).

[7] A capacity market pays a generator for its commitment to be available during periods of peak demand, distinct from the energy market, which pays per megawatt-hour actually generated. A plant can therefore earn revenue for standing ready without producing electricity. Thus, relevant to the analytical concern here, capacity-market participation could in theory sustain a low-utilization plant whose emissions profile reflects its business model rather than its abatement.

which raises an important question to address here: with so many related variables in one model, can the regression tell their effects apart? The fact is, it can. The variance inflation factor (VIF) measures how much each coefficient's uncertainty is inflated by its correlation with the other predictors; a value near 1 means a predictor carries largely independent information, while values above about 10 flag one too redundant to estimate cleanly. The largest VIF here is 5.8 (on State climate-policy count), with ISO/RTO membership at 2.2, the organizational-type factor at 5.4, and the combined state-market-share index at 1.9 (the two state shares it replaces had previously been collinear at $r = 0.87$). Despite State climate-policy count at 5.8, all variables still sit below the threshold, so the model estimates a separate, stable coefficient for each (Table 2.22 lists the VIF for every predictor).

2.1.3 Hypotheses

Building on the consolidated framework set out in Section 2.1, this chapter sketches out the eco-structural organizational political economy framework (Prechel, Blazek, and Amaral (2025)) that theoretically explains how organizational structure shapes plant-level NO_x emissions. The framework's central claim is that organizational structure conditions the opportunity-and-constraint surface within which firms pursue capital accumulation: what behaviors are profitable, what risks can be externalized, what financing is available, and what regulatory rules a firm must follow and which it has the capacity to influence or ignore. Pollution outcomes follow from where a firm sits on that surface.

The framework rests on three theoretical pillars, each linking organizational structure to pollution behavior through a distinct mechanism, and each translating into the empirical specification that follows. First, Prechel's eco-structural argument: organizational structure conditions firms' resource dependence, their embeddedness in markets and political-legal arrangements, and their capacity to mobilize politically to redefine corporate–state relations in their favor (Prechel, Blazek, and Amaral 2025; Polanyi 2010). Second, principal–agent dynamics (Miller 2005): organizational structure determines the monitoring environment between owners and operators (concentrated vs. dispersed principals, single vs. multiple, public vs. private), and therefore the discretion agents hold over operational choices, including pollution-relevant ones. Third, cost-of-capital: organizational structure conditions financing access, ease, magnitude, and hurdle rates (regulated rate-base

recovery, merchant market exposure, tax-exempt municipal bonds), and therefore the feasibility of capital-intensive abatement (Averch and Johnson 1962; L. L. Johnson 1973; Fowlie 2010; Donovan and Corbishley 2016).

However different their mechanisms, the three pillars carry the same load. Each connects organizational structure to emissions behavior by way of the opportunity-and-constraint surface introduced above: structure sets the surface a firm faces, and where the firm sits on it shapes what it emits. A corporate group that holds its plants through layered subsidiaries can shield itself from the legal and financial fallout of pollution in ways a single unified company cannot, lowering the firm's private cost of polluting (Prechel and Istvan (2016)). A power plant operated by a traditionally regulated utility can recover the cost of expensive pollution controls through customer rates set by a state regulator, while a similar plant selling electricity into a competitive wholesale market must fund those same controls from its own revenues at the returns investors demand, making the equipment far more costly to install in private terms (Fowlie (2010)). A government-owned plant whose ownership exists to provide public service may treat pollution avoidance as an end in itself rather than as a cost to be minimized; conversely, government-ownership in the case of the Tennessee Valley Authority is more likely to treat cheap electricity as the public good to be maximized, with environmental control merely another cost to be minimized.

Because the three pillars differ in *what kind* of constraint they relax, the analysis tests each organizational dimension across two adoption margins rather than one. NO_x control divides functionally into combustion modification (low-NO_x burners, overfire air, dry-low-NO_x combustors, water injection), which is relatively low-capital, and post-combustion treatment (selective catalytic and non-catalytic reduction), which is capital-intensive, especially SCR. This split is fuel-agnostic, applying identically to coal boilers and gas turbines, and that is where the analytical leverage lies. The cost-of-capital pillar acts specifically on the capital-intensive margin, so it should move post-combustion adoption (β_1^p, Equation 1b) while leaving combustion modification (β_1^c, Equation 1a) comparatively untouched; the eco-structural pillar (political clout, compliance evasion, asset inertia) and the principal-agent pillar are not capital-specific and should move both margins together. The difference $\beta_1^p - \beta_1^c$ is therefore a diagnostic for where an organizational dimension's effect lands: a large gap signals a mechanism running through capital cost, isolating the cost-of-capital

pillar from the other two. The predicted signs cover both adoption margins, their differential, and the conditional outcome (β_3, Equation 3b).[8]

How each prediction is judged depends on whether the three pillars agree. Where they point the same way (as for regulated status and government ownership), the prediction is directional and is supported when the coefficients line up across the two adoption margins and the outcome. Where they conflict (a horse race, as for ownership concentration and firm scale), no single coefficient settles the question; support is judged against a weaker criterion, that the total association points in the predicted direction and a majority of the channels point with it. Where the pillars cannot resolve even a net direction (lessee ownership, PURPA cogeneration), the hypothesis is treated as exploratory and carries no predicted sign.

The subsections below apply this framework to each dimension of organizational structure that the data measures directly. Unfortunately, the specific illustration of liability-shielding is a conceptual case only, here: identifying parent–subsidiary chains and the legal entity forms that carry shielding (LLCs, special-purpose entities, holding companies) requires corporate-affiliates data not available in EIA records. The empirical hypotheses therefore concentrate on dimensions EIA does capture (ownership form, ownership concentration, regulatory status, market position, and firm scale), treating the liability-shielding channel as one mechanism the framework names but the analysis only weakly proxies through related contrasts: notably traditionally regulated vs. competitive-market plants, and joint ownership.

Table 2.2 collects the predicted signs for every hypothesis; the subsections that follow develop the mechanism behind each.

Table 2.2: Predicted signs by hypothesis across the two adoption margins (Eq 1a, 1b), their differential, and the total and conditional outcome equations (Eq 3a, 3b).

Organizational trait	Pillar	(1a) Comb.	(1b) Post	Diff.	(3a) Total	(3b) Cond.
H1 - Organizational type						
Government	CoC + ES	+	+ (strong)	>= 0	-	-
Regulated IOU	CoC vs AJ	+ (weak)	+ (strong)	» 0	-	-
Cooperative	CoC vs OpDisc	+	+ (strong)	>= 0	-	-

[8]The within-adopter SCR-vs-SNCR choice (θ) and utilization (β_2) predictions are in the coal-boiler deep-dive.

Host-serving	Expl.	n.p.	n.p.	n.p.	+	+
H2 - Ownership concentration						
Number of owners	Horse-race	-	-	> 0	+	+
H3 - Cross-cutting status flags						
Third-party ownership	Expl.	n.p.	n.p.	n.p.	n.p.	n.p.
FERC cogeneration	Expl.	n.p.	n.p.	n.p.	n.p.	n.p.
H4 - Market position						
State market-share index	ES	-	-	~ 0	+	+
H5 - Firm scale (super-polluter)						
Firm scale index	Horse-race	-	- (offset)	> 0	+	+

Note. Signs are theoretical predictions on the log-odds of adoption (Eq 1a, 1b) and on log NOx intensity (Eq 3b). The differential is the post-combustion minus the combustion adoption prediction: a large positive gap is the cost-of-capital signature (the trait moves the capital-intensive margin far more than the cheap one), while a near-zero gap indicates the absence of that signature rather than positive evidence of another channel. `n.p.` = not predicted (exploratory). Pillar: CoC = cost-of-capital, ES = eco-structural, PA = principal–agent; horse-race = PA+ES vs CoC, judged on the weaker total-plus-majority-of-channels criterion. Predicted and observed signs are compared in 2.8.

2.1.3.1 H1 — Organizational type

Organizational type enters as a single mutually exclusive factor, `org_structure_type`, with five levels built from EIA-860 entity type crossed with regulatory status (Table 2.1): merchant independent power producer (IPP), regulated investor-owned utility (IOU), cooperative, government, and host-serving on-site generator. The merchant IPP is the reference category, a market-disciplined wholesale generator that earns a market price and funds capital at market rates, against which the other four types are contrasted in the H1 family below. Joint ownership, a separate ownership-structure dimension, is examined under H2 (ownership concentration).

Two forces are hypothesized to recur across the contrasts and act on different margins, so they are stated once here. Cost-of-capital is expected to govern the *adoption* margin: organizational types with cheaper or cost-recovered financing can more readily afford capital-intensive control equipment. On the other hand, operating-cost discipline is expected to govern the *efficiency* margin:

types insulated from market cost pressure (cost-of-service recovery, fuel pass-through, non-profit operation) run less efficiently, and so dirtier per MWh, holding adoption constant (Fowlie (2010); Borenstein and Bushnell (2015)). The merchant reference sits at the disciplined end of both, at the mercy of the market for access to capital and behold to its profits, which are a direct function of the efficiency of their operation.

2.1.3.1.1 H1a — Government

What this captures. The government level flags plants whose operating entity is federal, municipal, state, or a political subdivision, contrasted against the merchant reference.

Theory. Government-owned plants face the cheapest cost-of-capital environment in the sample: tax-exempt municipal bonds, federal financing arrangements, and lower effective hurdle rates make capital-intensive abatement equipment more affordable to install than the market-rate financing merchants face (Khanchel, Lassoued, and Khiari 2025; Fowlie 2010). They also operate under a different organizational mission, public service rather than capital accumulation, so pollution avoidance can be an end in itself rather than a cost to be minimized. The cost-of-capital and eco-structural pillars therefore predict government plants are more likely to adopt control, the capital-intensive post-combustion option in particular, and run cleaner than merchants. The expected direction is theoretical and the magnitude likely modest, because the level pools heterogeneous subtypes: federal entities like the Tennessee Valley Authority (TVA) may face a counter-pull, since soft budget constraints and political resistance to retirement can keep older uncontrolled units running, whereas large municipal utilities are often closer to the cleaner end of the public-power spectrum. Ultimately, this level is municipal-dominated (59% of government plant-years are municipal).

H1a. Government-owned plants are less NO_x-intensive than merchant plants.

2.1.3.1.2 H1b — Regulated IOU

What this captures. The regulated-IOU level flags investor-owned utilities whose plants are rate-regulated (cost-of-service) and not government-owned, contrasted against the merchant reference.

Theory. Regulated IOUs recover approved environmental capital spending through ratemaking, with state public utility commissions (PUCs) allowing "prudent" investment into the rate base,

while merchant generators face market-rate financing with no comparable recovery mechanism (Borenstein and Bushnell 2015). On the *adoption* margin, where abatement is a capital purchase, this lowers the cost of installing control equipment, so cost-of-capital predicts IOUs adopt more readily (Fowlie (2010) finds this among coal-fired plants). But because the utility earns its return on the capital in its rate base rather than on how efficiently it runs, and passes most fuel and operating cost through to ratepayers, its operational discipline is dulled. The Averch-Johnson theory (Averch and Johnson 1962; L. L. Johnson 1973) suggests rate-regulated firms tend to "employ more capital relative to variable inputs (including labor and fuel) than is consistent with cost minimization" (as characterized in Fowlie (2010)). That bias toward capital over variable inputs makes regulated firms readier to install abatement capital, the cleaner-adoption side; but the same organizational and budgetary slack over variable inputs like fuel (externalized to ratepayers) implies, on the *efficiency* margin, less efficient operation, and so dirtier plants. It is the slack the principal-agent pillar identifies, with diffuse ratepayer principals imposing looser cost-minimization pressure than the institutional investors and industrial parents behind merchant plants (cf., Wamsted (2023)). The eco-structural pillar adds PUC accountability pressure, leaning cleaner (its political-mobilization channel is ambiguous). On balance, then, IOUs face two competing forces, and this chapter predicts the adoption advantage dominates: their cheaper access to control capital makes them readier to install both combustion and post-combustion controls, and that greater adoption outweighs whatever operational slack the cost-recovery bargain creates, leaving IOUs cleaner than merchants overall. The stagewise adoption design tests whether that bet holds.

H1b. Rate-regulated investor-owned utilities run less NO_x-intensive plants than merchant generators.

2.1.3.1.3 H1c — Cooperative

What this captures. The cooperative level flags member-owned, not-for-profit utilities (predominantly generation-and-transmission cooperatives), contrasted against the merchant reference.

Theory. Cooperatives combine features of the government and regulated-IOU types but match neither exactly. Like government plants, they finance through subsidized, low-cost capital (federal Rural Utilities Service loans and cooperative lenders), so the cost-of-capital pillar predicts ready adoption of capital-intensive control, and cleaner plants, on the adoption margin. Like regulated

IOUs, they operate on a cost-recovery basis with weak market discipline, so the operating-discipline pillar predicts less efficient operation, and dirtier plants, on the efficiency margin. Unlike IOUs, however, cooperatives have no shareholders earning a return on rate base, so the Averch-Johnson over-capitalization motive does not apply. Instead, the operating-discipline prediction here rests on cost pass-through and the absence of market pressure, not on a capital-base incentive. As with the regulated IOU, we predict the cost-of-capital/adoption channel dominates on net, so cooperatives are cleaner than merchants, with operating discipline the countervailing force.

H1c. Cooperative plants are less NO_x-intensive than merchant plants.

2.1.3.1.4 H1d — Host-serving on-site

What this captures. The host-serving level flags industrial and commercial on-site generators, contrasted against the merchant reference. These plants serve a host facility's process load, often producing process steam, rather than selling primarily into the wholesale market.

Theory. For host-serving generators, electricity and emissions are byproducts of serving a host process rather than the core business, so the eco-structural and principal-agent pillars predict abatement loses out to process priorities: capital and managerial attention go to the host's production, not to control equipment, and many such plants are small, where post-combustion economies of scale are weak. The cost-of-capital pillar is ambiguous, since financing depends on the host firm, so the net leans dirtier. The prediction is exploratory: the level is heterogeneous, and for the combined-heat-and-power plants within it the per-MWh accounting bias discussed under H3b can inflate measured intensity independent of behavior.

H1d. Host-serving on-site generators are more NO_x-intensive than merchant plants (exploratory).

2.1.3.2 H2 — Ownership concentration

What this captures. Ownership concentration is measured by `n_owners`, the count of distinct ownership entities at the plant. It distinguishes sole-owned plants (a single owner) from joint arrangements and, among joint plants, indexes how dispersed control is. Joint arrangements are predominantly coal plants (49% coal-fired), where several utilities split capital and operating risk and

a regulated investor-owned utility (IOU) is the operator in 44% of cases; the sole-owned reference is heterogeneous.

Theory. The three pillars push in opposite directions, making H2 a horse race. The eco-structural pillar predicts dispersion is dirtier: asset inertia worsens when either abatement or closure requires coordination across owners, and a single holdout can keep an aging uncontrolled plant running past its economic life. The principal-agent pillar agrees, since more principals create coordination friction and free-riding on monitoring, weakening accountability over the operator. The cost-of-capital pillar pulls the other way: capital pooling lowers each partner's individual capital burden and can enable capital-intensive abatement a single owner could not afford. The horse-race prediction is that the eco-structural and principal-agent pillars together overwhelm the cost-of-capital advantage, so dispersion is net dirtier; the weaker criterion above applies.

Because most plant-years sit at the sole-owned baseline, the variation that identifies `n_owners` lives almost entirely in the joint subset, so a large-magnitude estimate with a wide confidence interval is expected.

H2. More dispersed plant ownership is associated with greater NO_x intensity.

2.1.3.3 H3 — Cross-cutting status flags

2.1.3.3.1 H3a — Third-party / leased operation

What this captures. Lessee ownership flags plants where ownership and operation are split: a *lessor* owns the asset and a separate *lessee* operates it under a long-term lease. The economic substance (who funds capital spending, who runs operations, who carries environmental liability) is split across the lease contract and varies case by case.

Theory. The three pillars give no clean directional prediction for lessee plants, because the lease contract itself determines which party faces which incentive. Under the eco-structural pillar, the lessee operates the plant day-to-day and is closest to the local political-legal environment while the lessor holds the asset and bears long-term capital risk, so rule-shaping capacity sits with whichever party is more politically embedded: ambiguous. The principal-agent pillar is probably the most relevant here, since the lease is a literal principal-agent split: the lessor wants to preserve the asset's value, the lessee to maximize operating revenue over the lease term. That split tends to

misalign incentives over capital spending (usually the lessor's to fund) and operations (the lessee's to run), which on balance leans toward less abatement, though how strongly is ambiguous. The cost-of-capital pillar is itself the rationale for these structures, since sale-leaseback and leveraged leases[9] exist to lower the lessor's effective cost of capital through tax and risk-sharing benefits, but whether that advantage reaches the plant's abatement decisions again depends on the contract. H3a is therefore exploratory.

H3a. Lessee ownership has no predicted association with NO_x intensity (exploratory).

2.1.3.3.2 H3b — PURPA / CHP designation

What this captures. `ferc_cogen_binary` is 1 when the plant holds Federal Energy Regulatory Commission (FERC) certification as a cogeneration Qualifying Facility (QF) under the Public Utility Regulatory Policies Act (PURPA). Almost all PURPA cogenerators are nonutility, so H3b lives within the nonutility group and identifies a within-nonutility contrast.

PURPA cogeneration certification is a legal status a plant can hold only if it is physically a combined heat and power (CHP) plant: one that produces both electricity and useful heat, such as process steam, from the same fuel. The dataset has no separate indicator for that physical configuration, so `ferc_cogen_binary` can only proxy it with the certification. Some plants are physically CHP but never seek certification, typically because they serve on-site industrial load and have no need for the mandatory-purchase, avoided-cost power purchase agreement (PPA) that Qualifying Facility status confers; these plants are coded 0 and fall into the comparison group. The coefficient therefore identifies the effect of PURPA certification cleanly, but captures the broader physical-CHP effect only with attenuation toward null, since the comparison group is contaminated with uncertified CHP plants. A second, distinct measurement issue runs the other way and sits in the outcome rather than the regressor: CHP plants typically book all fuel input against electricity output in EIA reporting, so the thermal output does not enter the denominator of `nox_rate_lb_per_mwh` (EPA eGRID 2021), inflating their measured NO_x rate by construction and making them

[9] Two financing maneuvers that split legal ownership from operation. In a sale-leaseback, the operator sells the plant to an investor and immediately leases it back, freeing up capital and handing the asset's tax benefits (notably depreciation) to the buyer while continuing to run it. A leveraged lease goes further: an equity investor funds part of the purchase and lenders supply the rest as non-recourse debt, with the investor-lessor claiming the tax benefits and the operator-lessee paying rent. In both, the point is financial, not operational.

appear dirtier regardless of operational performance. The H3b coefficient thus mixes behavioral and accounting components, and because the accounting bias points the same way as the eco-structural and principal-agent dirtier predictions, artifact and behavior are hard to separate.

Theory. The three pillars push in different directions. Cost-of-capital leans weakly cleaner: QFs hold mandatory utility PPAs at avoided-cost rates, giving predictable revenue and a lower hurdle rate than merchant IPPs face, though QFs tend to be small plants, where selective catalytic reduction (SCR) economies of scale are weak, and avoided-cost rates fell over the sample period (Borenstein and Bushnell 2015). The eco-structural and principal-agent pillars both point dirtier: QFs serve industrial or commercial hosts whose priority is process steam rather than emissions, ownership is split between host and cogen developer, and many QFs date to the original PURPA buildout and may carry substantial asset inertia, so abatement loses out to host process priorities.

H3b. PURPA cogeneration status has no predicted association with NO_x intensity; observed differences partly reflect measurement (exploratory).

2.1.3.4 H4 — Market position

What this captures. Market position is the parent firm's share of state-level electricity activity, from two underlying quantities: `state_capacity_share` (share of total nameplate capacity in the plant's state) and `state_output_share` (share of state net generation), both computed at the EIA-reporting-entity level. Because the two shares carry the same signal and are each right-skewed, they enter the primary specification as a single standardized composite, `state_market_share_index`, mirroring the `firm_scale_index` construction; both raw shares remain in the parquet for descriptive use.

There are two caveats worth noting here. First, Grant's "dominant utility" framework (Grant, Jorgenson, and Longhofer 2016) is conceptually invoked here, but the sample does not contain dominant utilities in the strict sense. Instead the variable identifies a continuous gradient of state-level presence rather than a binary dominance contrast. Second, market shares are computed at the reporting-entity level rather than the ultimate parent company; when a parent owns multiple subsidiaries each reporting separately, its true state-level concentration is understated, so the H4 estimates are likely attenuated.

Theory. The dominant mechanism is eco-structural, with the other pillars weak or partly absorbed by H5 firm scale. Eco-structural theory suggests firms with larger state-level presence wield political clout at the state PUC, shaping environmental rule-making and enforcement in their favor (Grant, Jorgenson, and Longhofer 2016), so larger state share predicts more permissive enforcement and dirtier outcomes. Cost-of-capital offers only a weak counter-pull toward cleaner, since larger state presence may ease state-PUC financing of capital spending in the regulated subset. The principal-agent mechanism is largely absorbed by H5 firm scale: conditional on H5, H4 captures the residual regional-concentration effect, with direction unclear net of H5.

The eco-structural dirtier prediction aligns with the SuperPolluter sign pattern at H5, but the underlying mechanism is different: H5 is national-scale plus asset inertia plus ownership dispersion, while H4 is state-level political clout. Both hypotheses can find empirical support through their respective mechanisms even with overlapping signs.

H4. Plants owned by firms with a larger share of their state's electricity market are more NO_x-intensive.

2.1.3.5 H5 — Firm scale (the SuperPolluter hypothesis)

What this captures. Firm scale is the absolute size of the plant's parent firm's portfolio, measured by the number of plants owned (`owner_n_plants`) and total nameplate capacity (`owner_total_nameplate_capacity_mw`), combined into a single standardized composite, `firm_scale_index`; both raw variables remain in the parquet for descriptive use.

Theory. This is the SuperPolluter hypothesis: that large, gain-driven firms with substantial asset inertia make systematically more polluting decisions than smaller or more constrained firms (Grant, Longhofer, and Jorgenson (2024)). The three pillars of the unified framework push the prediction in distinct directions. The eco-structural pillar predicts that large firms have greater political mobilization capacity to influence environmental rule-setting and enforcement in their favor, and that they carry asset inertia in the form of older, more diverse fleets that include uneconomic-to-retire coal units protected by political and labor coalitions. The principal-agent pillar predicts that at scale, ownership is more dispersed (public companies, institutional investors), the chain between principals and operators is longer, and cost-minimization pressure on managers translates into less abatement

investment, cheaper controls when adopted, and less utilization once installed. The cost-of-capital pillar runs in the opposite direction: large firms have cheaper access to capital, deeper financing options, and lower effective hurdle rates, all of which should make capital-intensive abatement easier to fund (Donovan and Corbishley 2016; Fowlie 2010). The SuperPolluter prediction therefore runs across the two adoption margins and the outcome. On adoption, the eco-structural and principal-agent pillars predict that large firms are *less* likely to control on both the combustion and post-combustion margins, while the cost-of-capital pillar pushes back on the capital-intensive post-combustion margin specifically, where cheaper financing should matter most; the net is under-adoption on both margins with post-combustion partly offset, leaving a positive post-minus-combustion differential (the same cost-of-capital signature read for the other dimensions). On the outcome, the eco-structural and principal-agent pillars dominate: larger firms are predicted to run dirtier per MWh both in total (Equation 3a) and conditional on adoption (Equation 3b), the mechanism being asset inertia and operational slack rather than the binary decision to install control.

The composite folds in two facets of scale that point the same way: portfolio breadth (`owner_n_plants`: political reach across many state commissions and bureaucratic inertia at the holding-company level) and sheer size (`owner_total_nameplate_capacity_mw`: economies of scale in compliance and legal capacity, lobbying-budget magnitude, and the size of the asset base to be defended). Both predict more polluting, so `firm_scale_index` carries their shared signal.

Because the pillars conflict on the adoption margins, the weaker criterion above applies.

H5. Plants owned by larger firms (the super-polluters) are more NO_x-intensive.

2.2 Data & Methods

This section sets out the sample, the data, and the estimation strategy in the order the analysis uses them. The sample is the full fossil-fired fleet (every coal- and natural gas-fired plant of at least 25 MW in the continental United States), observed at three snapshots (2009, 2018, and 2021) and pooled into a single cross-section with year fixed effects. Plant generation, capacity, ownership, and control-equipment records come from EIA Forms 860 and 923; installed NO_x controls are supplemented with the EPA CAMD facility-attributes file, which captures the simple-cycle gas-

turbine controls EIA does not record; plant emissions come from EPA eGRID; and the regulatory and political-economic context is assembled from a state policy panel. The analytical strategy follows directly from the technology mechanism: NO_x control is measured on two fuel-agnostic functional margins: combustion modification and post-combustion treatment. These map onto four estimating equations run on the full fossil fleet (two adoption logits and a nested outcome pair). The analysis starts with a descriptive variance decomposition: it asks how much of the variation in NO_x intensity is already accounted for by observable technical factors, specifically, the fuel a plant burns and the two control functions it operates. These measurements fix the size of the residual gap within which organizational structure might plausibly operate, which ultimately motivates the organizational question: with 51.2% of the variation in NO_x emission intensity explained by technical factors, to what degree do organizational structural characteristics explain the remaining variation in NO_x emission intensity? Following this descriptive exploration, the section sets up the estimating equations. Each estimating equation is stated explicitly alongside the year fixed effects and utility-clustered standard errors that structure it and the identifying assumptions that bound what the adoption differential and the 3a-to-3b ('total'-to-'conditional-on-adoption') attenuation can and cannot claim. The section then closes with a complete inventory of every analytical variable, organized by conceptual role: the dependent variable, the organizational-structure predictors, the functional NO_x-control mediator, and the control blocks, each with its data source and construction recipe.

2.2.1 Data sources

Table 2.3 lists the public datasets the analysis draws on, the version or coverage of each, and the role it plays (EIA 2025; EPA 2026, 2025; EPA eGRID 2009, 2018, 2021; Gonzalez and Carlotto 2026; NCSL 2026; Grossmann, Jordan, and Mccrain 2021); every variable's source is given in full in the variable inventory below.

Table 2.3: Data sources for the analysis.

Dataset	Provider	Version / coverage	Role in the analysis

Form EIA-860 (Annual Electric Generator Report)	U.S. Energy Information Administration	2009, 2018, 2021	Generators, boilers, ownership, and control-equipment records; the plant spine
Form EIA-923 (Power Plant Operations Report)	U.S. Energy Information Administration	2009, 2018, 2021	Heat input, fuel consumption, and Schedule 8 environmental-operations data
eGRID (Emissions & Generation Resource Integrated Database)	U.S. Environmental Protection Agency	8th, 13th, 16th editions (2009, 2018, 2021 data)	Plant NOx emissions and emission rate; ISO/RTO membership
Clean Air Markets Program Data (CAMD/CEMS facility attributes)	U.S. EPA Clean Air Markets Program	2009, 2018, 2021	Simple-cycle gas-turbine NOx-control flags EIA does not record
State climate-policy dashboard	ClimateXChange	2009, 2018, 2021	State climate-policy count, RPS targets, carbon pricing
State partisan composition	NCSL; Correlates of State Policy Project	2008–2024	Governor and legislative party control (political-economic context)

2.2.2 Analytical Strategy

The three snapshot years (2009, 2018, 2021) are pooled into a single cross-sectional sample with year fixed effects. The year fixed effects absorb residual macro-level mean shifts: gas price movements, aggregate demand changes, and anything not already captured by the control variables. The substantive content of what makes one year different from another (the regulatory pressure that compelled NO_x control adoption) is carried by the control variables themselves: the state emissions cap indicators, the carbon price variable, the RPS targets, and the political-economic context variables. Year fixed effects are therefore a residual mop-up, not a substitute for that content. The specification likewise absorbs NERC-region fixed effects, so the organizational coefficients are identified within grid region, net of regional differences in fleet composition and air-quality regulation. What the organizational structure coefficients then estimate is variation in adoption and emissions intensity *conditional on* facing the same regulatory environment, which is precisely the theoretical claim of this chapter. Standard errors are clustered at the utility level to account for

within-utility correlation across plants.

The system is estimated as four equations on the full fossil-fired fleet: two adoption logits (Equations 1a and 1b) and two outcome specifications (Equations 3a and 3b).

2.2.2.1 Descriptive Context

Before the formal estimating equations, it is useful to contextualize the dependent variable: how much of the variance in NO_x intensity is explained by the technical and contextual factors the primary specification controls for, and how much room remains for organizational structure? A sequential variance decomposition that mirrors the primary specification addresses this, entering the same blocks in order: snapshot year, a plant's capacity to emit (scale, utilization, generator age, prime-mover configuration, and fuel mix), the abatement controls it operates, and its regulatory and political-economic context, with organizational structure added last.

Together these factors, namely snapshot year, capacity to emit, the two functional control margins (combustion modification and post-combustion treatment), and regulatory and political-economic context, explain approximately 51.2% of the total variance in log NO_x emission intensity. The remaining 48.8% is unexplained by the measurable characteristics the primary specification conditions on. It is within that gap that organizational-structure heterogeneity is expected to operate: plants facing identical regulatory obligations, burning identical fuels, with identical capacity to emit, and technically capable of adopting identical controls may nonetheless make systematically different adoption and operational decisions depending on how they are organized.

Figure 2.1 shows the distribution of NO_x intensity by fuel type and year: coal-only and mixed-fuel plants cluster at higher intensity than gas-only plants, and all three distributions shift leftward from 2009 to 2021 (the secular regulatory trend accounts for 1.6 pp of explained variance, and the capacity-to-emit block, which folds fuel mix in with scale, utilization, age, and prime mover, a further 39.7 pp). Figure 2.2 shows why control function matters: operating either function is associated with markedly lower median NO_x intensity (the marginal median moves from 0.47 to -0.27 in log lb/MWh for combustion modification, and from 0.32 to -1.35 for post-combustion treatment), and partitioning the fleet by the joint state of both controls shows intensity falling from the uncontrolled group through the single-control groups to the both-controls group. Figure 2.3 shows how the

adoption landscape has shifted by fuel type and year: combustion control is widespread across the fleet (77% of plant-years), and post-combustion treatment (46%) is substantial across fuels (common even among gas plants, not only coal), rising across the snapshots under CAIR/CSAPR pressure. Figure 2.4 integrates the sequential decomposition (year, then capacity to emit, then the adopted controls, then regulatory and political-economic context), accumulating to 51.2%, leaving the 48.8% gap within which organizational structure can produce heterogeneous effects. (Per-block increments are order-dependent; the total non-organizational R^2 and the organizational-structure increment are order-invariant.)

Organizational structure occupies a modest but statistically reliable part of that gap. Adding the organizational-structure block to the full set of non-organizational controls, net of capacity to emit, the adoption margins, and regulatory and political-economic context, raises the explained share of log-NO_x variance by 1.0 percentage points; equivalently, organizational structure accounts for 2.0% of the variance those technical and contextual factors leave unexplained. The increment is statistically significant (nested $F(10, 4761) = 9.8$, $p < .001$), and the adjusted R^2 rises from 51.0% to 51.8%, so the gain is not an artifact of adding regressors. This is a descriptive partition rather than a causal decomposition; the conditional associations in the sections that follow show which organizational features drive it and whether they persist once adoption is conditioned on.

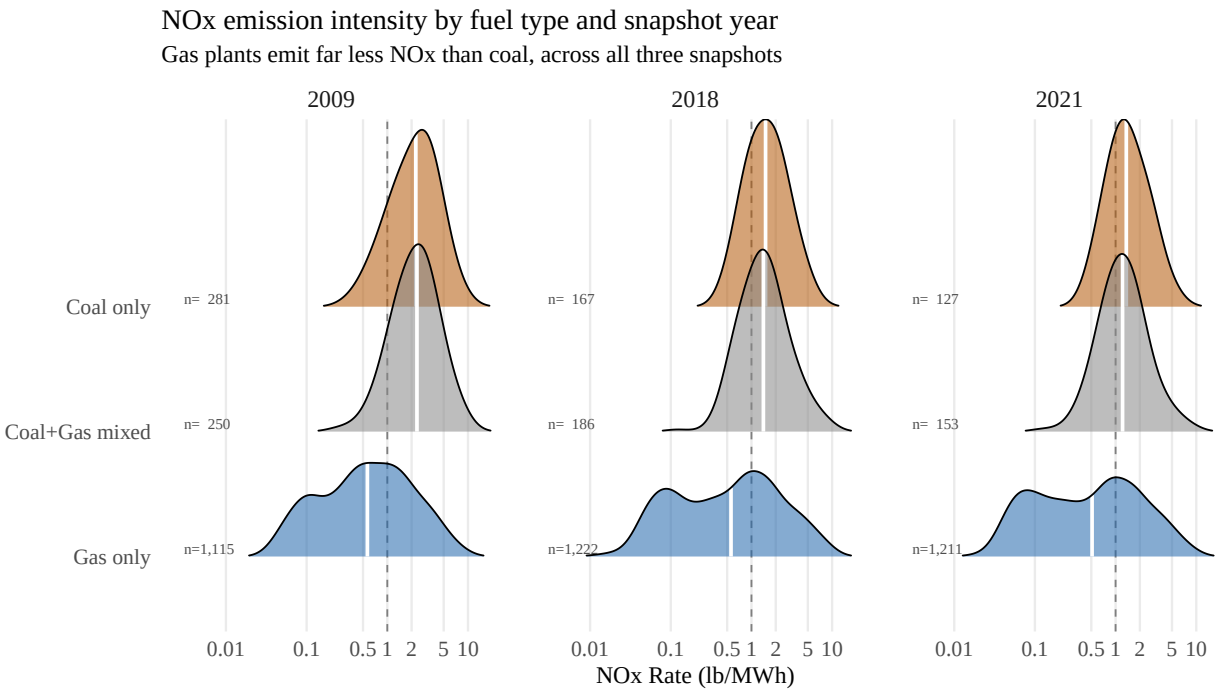

Figure 2.1: NO_x emission intensity by fuel type (ridges) and snapshot year (panels). The x-axis is log-spaced, in lb/MWh; the dashed line marks 1 lb/MWh and white lines the group medians.

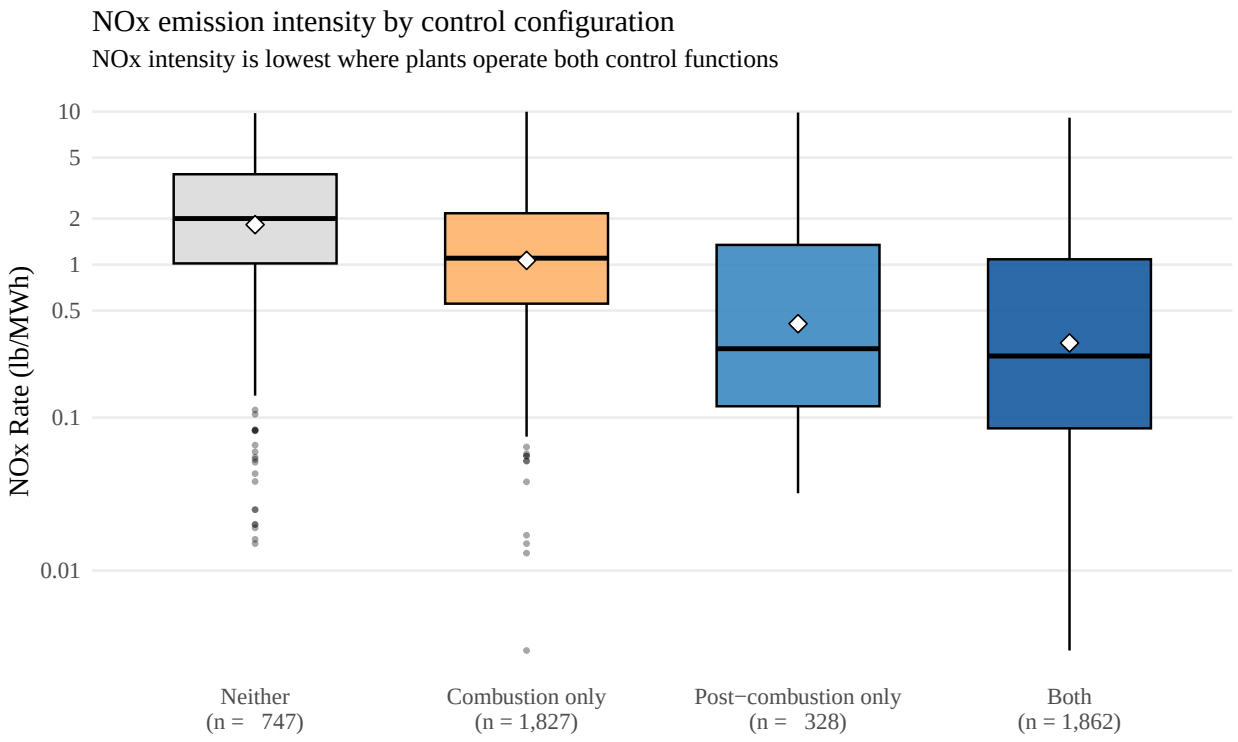

Figure 2.2: NO_x emission intensity by the joint state of the two control functions. Each plant-year falls in one of four mutually exclusive groups, ordered left to right by decreasing mean intensity; boxes show the median and IQR, the white diamond the group mean, with group sizes under each label. The y-axis is log-spaced, in lb/MWh.

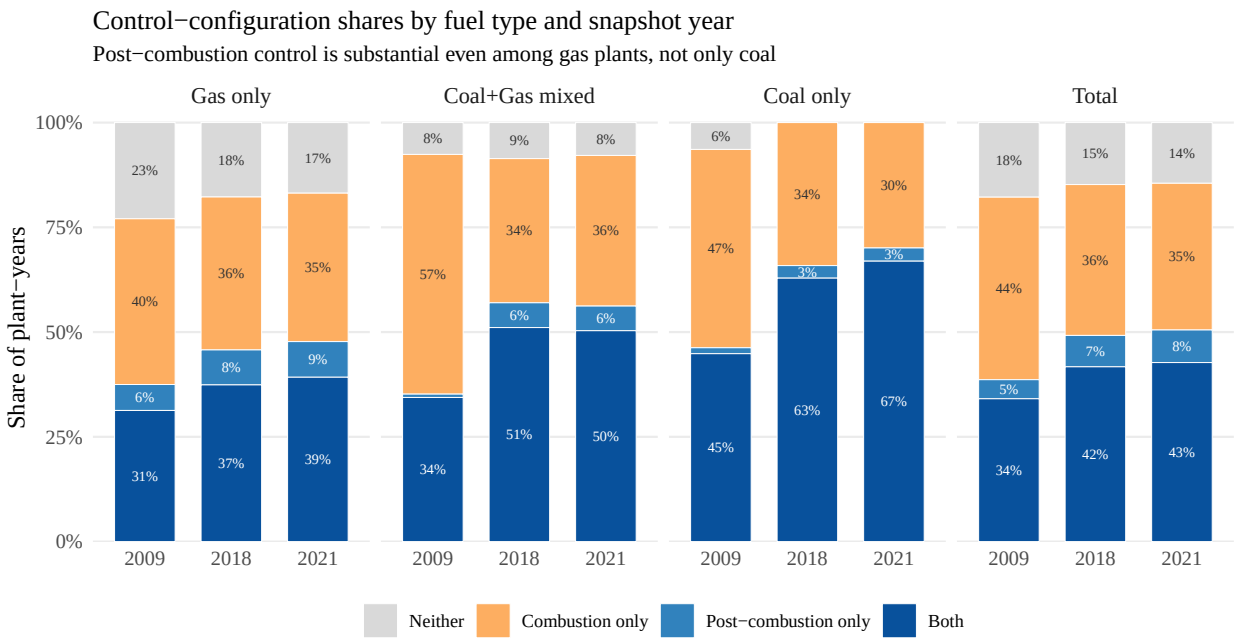

Figure 2.3: Fleet composition across the four mutually exclusive control configurations, by fuel type and for the fleet total, across the three snapshots. Each bar sums to 100% of its group's plant-years; colors are shared with Figure 2.2.

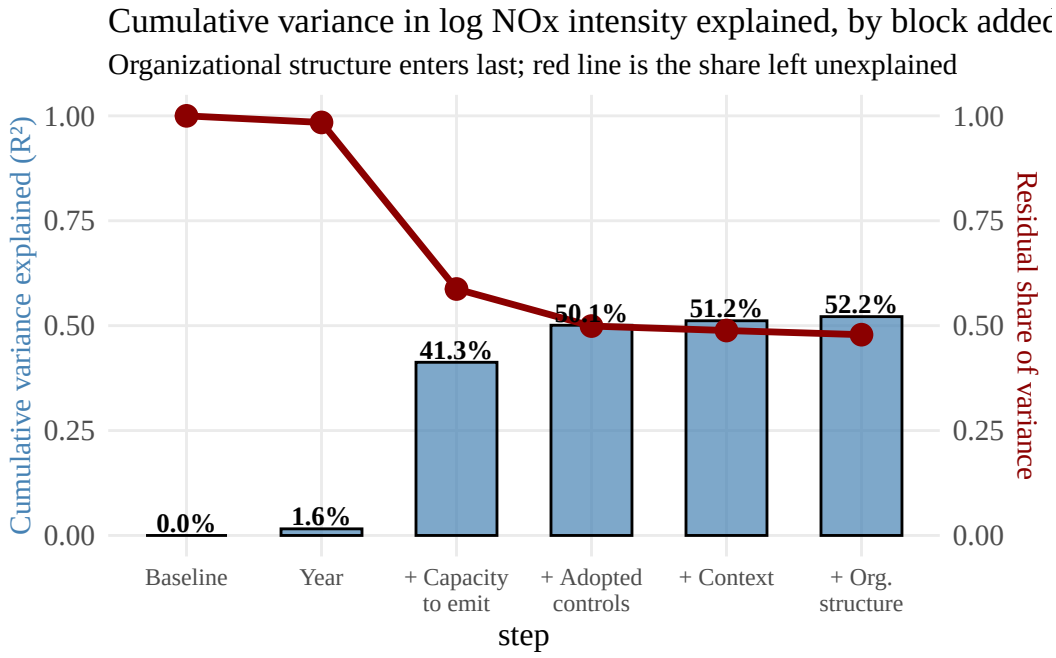

Figure 2.4: Sequential variance decomposition of log NO_x intensity, mirroring the primary specification. Blue bars show cumulative R^2 as the blocks enter in order, organizational structure last; the red line is the share of variance left unexplained.

2.2.2.2 Equation 1 — Adoption (two functional margins)

Two logistic regressions on the full sample (n = 4,818 plant-years), one per control function: Equation 1a for combustion-modification control and Equation 1b for post-combustion control.

$$P(\text{Control}_i^f = 1) = \text{logit}^{-1}(\alpha_1^f + \beta_1^f \cdot \text{OrgStructure}_i + \gamma_1^f \cdot \text{Controls}_i + \delta_t \cdot \text{Year}_t), \quad f \in \{c, p\}$$

where Control^c is `has_combustion_control` and Control^p is `has_postcombustion_control`, and δ_t are year fixed effects (2009 omitted). A positive β_1^f indicates that plants with that organizational characteristic are more likely to adopt control function f. Estimating the two margins as separate binaries rather than a single ordered ladder is deliberate: the margins differ in capital intensity (post-combustion treatment, especially SCR, is capital-intensive; combustion modification comparatively cheap), so the contrast $\beta_1^p - \beta_1^c$ distinguishes the cost-of-capital mechanism (which should load on the post-combustion margin) from the eco-structural one (which should move both about equally). The functional split, combustion modification versus post-combustion treatment, following the standard EPA APTI 418 taxonomy, is fuel-agnostic, applying identically to coal boilers and gas turbines, so both equations are estimated on the full fossil fleet.[10]

2.2.2.3 Equation 3 — Outcome: NO_x Emission Intensity

The outcome is `ln_nox_rate_lb_per_mwh` (log NO_x emission intensity, lb/MWh), estimated on the full fossil fleet (n = 4,818) as a nested pair. Log is the primary functional form — the rate is right-skewed, effects are multiplicative, and the coefficients read as semi-elasticities; a levels specification (`nox_rate_lb_per_mwh`) is retained as a robustness check. Equation 3a omits the adopted-control indicators and recovers the *total* association between organizational structure and NO_x intensity:

[10] Equation 2, analyzing utilization, or how intensively an adopted control is actually run, is reserved for the coal-boiler deep-dive. The control-in-service hours it requires are collected only for boiler-associated controls (EIA-923 Schedule 8), so the measure is structurally undefined for the simple-cycle and combined-cycle gas units that make up much of the fleet. The primary analysis therefore presents Equations 1 and 3 on the full fossil fleet; Equation 2 is developed within the coal-boiler subsample, where utilization is observed.

$$\text{(3a)} \quad \ln(\text{NO x Rate}_i) = \alpha_3^{\text{tot}} + \beta_3^{\text{tot}} \cdot \text{OrgStructure}_i + \gamma_3^{\text{tot}} \cdot \text{Controls}_i + \delta_t \cdot \text{Year}_t + \varepsilon_i$$

Equation 3b adds the two adopted-control indicators and recovers the association *conditional on* adoption:

$$\text{(3b)} \quad \ln(\text{NO x Rate}_i) = \alpha_3' + \beta_3' \cdot \text{OrgStructure}_i + \lambda^c \cdot \text{Combustion}_i + \lambda^p \cdot \text{PostComb}_i + \gamma_3' \cdot \text{Controls}_i + \delta_t \cdot \text{Year}_t + \varepsilon_i$$

where Combustion_i is `has_combustion_control` and PostComb_i is `has_postcombustion_control`. The movement of the organizational coefficients from β_3^{tot} (3a) to β_3' (3b) characterizes how much of each organizational association with NO_x intensity is accounted for by which control functions the plant has adopted, while coefficients that persist in 3b indicate organizational variation in emissions intensity beyond the adoption decision. Because both equations are estimated on the full coal-and-gas population, the comparison is not restricted to plants that have selected into post-combustion controls.

2.2.2.4 How the four equations map to the design

The four primary equations make explicit the structural logic: organizational structure is allowed to act at two points in the abatement process, and the design separates them rather than collapsing them into one number.

The first point is adoption. Equations 1a and 1b ask whether organizational structure predicts adoption of combustion-modification control and of post-combustion control, and estimating the two margins separately is what makes the cost-of-capital test sharp. Because post-combustion treatment is the capital-intensive option, a coefficient that loads on the post-combustion margin (Equation 1b) while sparing the combustion margin (Equation 1a) is the cost-of-capital signature, whereas one that moves both reflects a non-capital, eco-structural channel; the difference $\beta_1^p - \beta_1^c$ is read directly across the two equations.

The second point is the outcome. Equations 3a and 3b estimate NO_x intensity without and with the two adopted-control indicators, so the movement from the total association (3a) to the conditional one (3b) isolates how much of the organizational signal is bound up with what the plant has adopted. What that movement does and does not license is taken up in Section 2.2.2.6.

Both stages are estimated on the full fossil fleet, because the combustion/post-combustion split is functional rather than technological and so applies identically to coal boilers and gas turbines.[11]

2.2.2.5 Sample composition across snapshots

Table 2.4 cross-references plant characteristics by their 'presence' across the three snapshots. Plants that appear in all three snapshots are in the 'Survivor' class (in binary terms, 111); plants that leave at any point (either 110 or 100) are 'Exiters'; 'Entrants' join late (either 011 or 001); Finally, plants that enter then leave are 'Blips' (i.e., 010) and plants that appear in the first, leave, then re-enter (i.e., 101) are 'Gaps' – both 'Blips' and 'Gaps' are excluded from Table 2.4. Relevant to note here is that the sample is *not* a balanced panel observed repeatedly; rather, the sample is three independent cross-sections, each a near-complete census of the operating coal- and gas-fired fleet in its snapshot year. A plant absent from a later snapshot is therefore not a dropped observation but a unit that was not operating that year, so the plants that a critic might identify as "closing" are actually retirements that should be absent rather than missing data.

Table 2.4: Plant characteristics by presence class (survivors, exiters, entrants) across the three snapshots; blip/gap classes excluded.

	Presence class (first-seen)		
Variable	Survivor	Exiter	Entrant
Sample size			
Plants (n)	1,266	375	239
Outcome			
NOx intensity (lb/MWh)	1.35	2.51	1.65
Organizational structure			

[11] The finer, boiler-only distinctions (the SCR-versus-SNCR choice, capacity-weighted via `is_scr_dominant`; the chemical-efficiency gap between SCR, $≈$80–90% removal, and SNCR, $≈$25–50%; and operating-hours utilization) are not defined for simple-cycle gas turbines and are taken up in the coal-boiler deep-dive (Section 2.7).

Merchant / IPP (reference)	38.9%	45.3%	43.9%
Regulated IOU	25.9%	25.1%	18.4%
Cooperative	6.2%	4.5%	6.7%
Government	13.4%	14.9%	12.1%
Host-serving	15.6%	10.1%	18.8%
Owner portfolio (plants)	48.5	44.1	56.2
Owner total capacity (MW)	20,596	18,887	17,048
Technology			
Any NOx control	85.6%	72.0%	74.1%
SCR	24.5%	11.2%	29.7%
Plant characteristics			
Nameplate capacity (MW)	557	345	402
Gas heat share	0.66	0.47	0.74
Coal primary	24.5%	45.1%	8.4%

Note. First-seen timing: each plant summarized at its earliest in-sample snapshot. Survivor = 2009+2018+2021; exiter = present early, absent later; entrant = absent early, present later. Blip/gap excluded.

The three classes of Survivors, Entrants, and Exiters differ systematically on the variables of relevance to the project. Exiters are predominantly the retiring coal fleet: mean NO_x intensity 2.51 lb/MWh against 1.35 for survivors, 45% coal-primary, and only 11% operating selective catalytic reduction. Entrants are the natural-gas buildout: 74% gas by heat input, 8% coal-primary, and, like every presence class, predominantly merchant (merchant share 44%, against 39% among survivors); their one distinctive organizational feature is a tilt away from rate-regulated utilities (Regulated IOU 18% versus 26%), with no rate base over which to socialize abatement costs. Survivors are the larger, better-emission-controlled core (SCR 24%, against 11% among exiters). In sum, the dominant compositional change is fuel, from (uncontrolled) coal toward gas across the window; organizational composition is comparatively stable, merchant-plurality throughout, with the one tilt being entrants away from rate-regulated utilities. The fuel shift part of this compositional change is conditioned out by the fuel-mix controls, and its temporal pattern by the year fixed effects. The organizational variation, by contrast, is what the design actively seeks to exploit. The organizational structure coefficients are identified from within-year, within-region, within-fuel contrasts between regulated and merchant plants, so the cross-year drift in merchant share does not bias them.

One might object that organizational structure could shift *within* a plant across snapshots. It

almost never does, though: only 6% of survivors (plants present in all three snapshots) change organizational type, and 94% hold the same type throughout. Because organizational structure is effectively fixed within a plant, entry into and exit from the sample tracks fuel and period, both of which are already controlled, rather than the variable of interest. This leaves entity presence in the sample largely uncorrelated with (orthogonal to) organizational structure.

A related concern is attrition. Because sample turnover sorts plants on utility status (the exiting plants are somewhat more merchant, with merchant share 45% versus 39% among survivors, as are the entrants), attrition could bias the regulated-versus-merchant contrast; the direction cannot be signed a priori, because it depends on which group is truly dirtier. The protection is that organizational type is effectively fixed within a plant (as discussed above), so sample turnover tracks fuel and period rather than the variable of interest, and the year and region fixed effects and the regime-spanning design absorb period- and geographic-level confounding. What none of these can rule out, because organizational type is essentially fixed within a firm, is unobserved firm-level heterogeneity correlated with regulated-versus-merchant status. This final issue is addressed in Section 2.2.2.7.

2.2.2.6 Stagewise Conditional Analysis

Equations 1 and 3 are estimated across the full fossil-fired population. Equation 1 is two adoption logits (one for combustion-modification control, one for post-combustion control), and Equation 3 is estimated with and without those two adopted-control indicators (3a versus 3b). This structure enables a stagewise comparison of organizational structure coefficients across models. When organizational structure coefficients attenuate (shrink) after conditioning on the two adopted control functions, it is consistent with the notion of adoption carrying part of the association between organizational structure and NO_x emissions intensity. In other words, the variation in emissions that org structure predicted unconditionally is partially accounted for by which control functions the plant has adopted (intuitively, this makes sense given the pre-determined, invariant efficacy of pre- and post-combustion NO_x reduction methods). Conversely, coefficients that persist after conditioning on adoption indicate organizational structure shapes emissions through channels other than the adoption decision itself. This may include, for example, operational choices, maintenance investment, utilization intensity, or unmeasured firm-level practices. Of course, these are firmly

bounded claims: this comparison does not decompose the total effect into attributed fractions of causal effect; it merely characterizes the conditional relationship at each stage and lets the pattern of attenuation and persistence speak to where in the abatement process organizational differences are most consequential.

Figure 2.5 maps this structure. Organizational structure is tested against the two fuel-agnostic adoption channels, combustion modification (β_1^c, Equation 1a) and post-combustion treatment (β_1^p, Equation 1b), and against NO_x intensity directly (β_3, Equation 3b, conditional on both controls); thick arrows mark the organizational-structure coefficients the hypotheses predict, thin arrows (λ^c, λ^p) the estimated control-to-outcome conditioning relationships, with controls and year fixed effects entering every equation but omitted for clarity. The contrast $\beta_1^p - \beta_1^c$, whether organizational structure moves the capital-intensive post-combustion margin more than the cheap combustion margin, is the diagnostic for a cost-of-capital channel; a near-zero contrast indicates the absence of that signature rather than positive evidence of another.

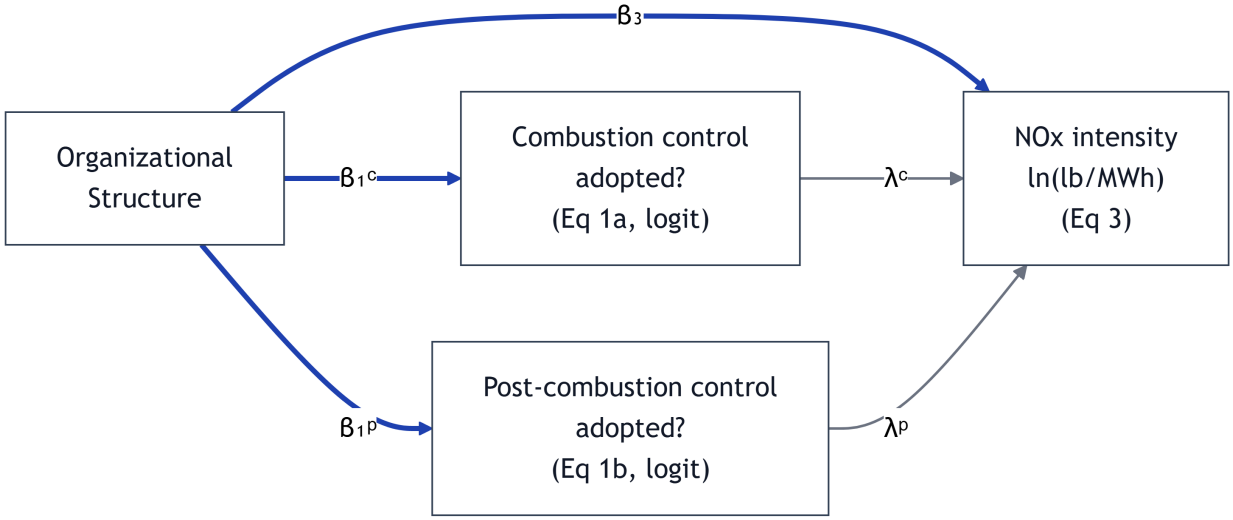

Figure 2.5: Conceptual map of the primary analysis (full fossil fleet, ≥25 MW).

The organizational-structure coefficients in the diagram are: β_1^c, the effect of organizational structure on the probability of adopting combustion-modification control (Equation 1a, logit); β_1^p, the effect on the probability of adopting post-combustion control (Equation 1b, logit); and β_3, the effect on log NO_x intensity conditional on both control functions (Equation 3b). The conditioning parameters λ^c and λ^p capture the average difference in log NO_x intensity associated with each

control function and appear in the regression tables for transparency. The central quantity of interest is the differential $\beta_1^p - \beta_1^c$: because post-combustion treatment is capital-intensive and combustion modification is not, an organizational dimension whose effect concentrates on the post-combustion margin (β_1^p) bears the signature of a cost-of-capital channel. A differential near zero indicates the absence of that signature: the dimension shifts both margins alike, as a broad abatement propensity would, without the capital-intensity asymmetry the cost-of-capital mechanism predicts. A positive coefficient on log NO_x intensity means dirtier than reference; a negative coefficient means cleaner.

2.2.2.7 Identifying assumptions

This section walks through the assumptions underlying the conditional associations. It begins with the core caveat (organizational structure is not randomly assigned, so the coefficients read as conditional associations, not causal effects), then bounds what the stagewise conditional analysis can claim. It next defends pooling three cross-sections with year fixed effects against the alternative of an annual plant-level panel with plant fixed effects, and closes with the residual threats to the conditional-association interpretation and the robustness checks that address them.

Organizational structure is not randomly assigned. A plant's ownership form and regulatory status are choices made by firms, shaped by history, strategy, regulation, and competition. They are not 'treatments' handed out by an experiment. That matters here because whatever unobserved firm characteristics drive those choices may also drive abatement: a firm disposed toward a particular ownership form may, for reasons not measured here, be equally disposed toward (or against) cleaning up its plants. The analysis therefore makes no claim to have isolated the causal effect of organizational structure on emissions. Every coefficient reported below is a conditional association. A conditional association is interpreted as the difference across organizational forms that survives the controls; it is read that way throughout.

The conditioning structure shapes what each equation may plausibly claim. The two adoption equations (Equations 1a and 1b) and the total outcome equation (Equation 3a) are estimated on the full fossil-fired fleet and condition only on exogenous controls. The conditional outcome equation (Equation 3b) adds the two adopted-control indicators, which are themselves outcomes of organizational structure rather than characteristics that precede it. Therefore, conditioning on

such post-treatment variables means the Equation 3b organizational coefficients, and the 3a-to-3b attenuation, are descriptive rather than a causal decomposition. The attenuation comparison, specifically Equation 3a without the adopted-control indicators versus Equation 3b with them (see Section 2.2.2.6), asks how much of the organizational association with NO_x intensity remains once adoption is conditioned on: whether the organizational coefficient shrinks, and by how much. This is not a causal decomposition; the analysis is only able to claim that organizational structure and adoption are correlated (or not), and conditioning on adoption either reduces the residual organizational coefficient (or not).[12]

Why pool three regime cross-sections with year fixed effects rather than estimate an annual plant-level panel with plant fixed effects? Because the treatment (organizational structure) is very nearly fixed within a plant over the window: only 6% of survivors change organizational type between 2009 and 2021 (94% remain stable), so plant fixed effects would absorb almost all of the organizational variation the contrast relies on. eGRID publishes most years across the window, so a roughly biennial panel is constructible; pooling regime snapshots is therefore a design choice, not a data necessity.

Two classes of firm-level confounding remain, and neither can be partitioned away via fixed effects. Because organizational structure is essentially fixed within a firm, and barely moves even within a plant (the 94% of survivors hold their organizational type across the window), parent-firm or plant fixed effects would absorb the very organizational variation the analysis is identified from; the treatment and the fixed effect are collinear, so that estimator is unavailable. That leaves two residual threats of unequal severity. Time-invariant firm-level unobservables, an unmeasured firm characteristic correlated with both organizational form and abatement, cannot be ruled out and are acknowledged as a limitation; the organizational coefficients are read as between-firm associations accordingly: differences across firms of different organizational types, not the effect of a given firm changing its type. Time-varying confounding, an unobserved firm-level shift that moves organizational structure and abatement together (for instance a strategic pivot that changes ownership and lifts abatement investment at once), is more constrained: because organizational

[12]The within-adopter analysis (which post-combustion technology is operated and how intensively) is estimated separately on the coal-boiler subsample; because it conditions on having adopted, its coefficients describe that adopter subpopulation and should not be read as effects across the full fleet.

structure barely moves within a plant, there is little within-plant covariation for such a confounder to exploit, so that channel is both rare and minimal when present.[13]

2.2.2.8 Estimation procedure

All four primary equations (the two adoption logits 1a/1b and the two outcome specifications 3a/3b) are estimated with cluster-robust standard errors (CR1) clustered at `utility_id`, computed with `vcovCL` from the `sandwich` package and reported via `coeftest`. The outcome variable in Equation 3 is `ln(NO~x~ rate)`. The within-adopter SCR-vs-SNCR and utilization analyses are reported in the separate coal-boiler deep-dive.

2.3 Results

2.3.1 Descriptive Statistics

Table 2.5 presents descriptive statistics for key variables in the analytical dataset, on the full coal-plus-gas fleet at or above the 25 MW control-measurability floor (4,818 plant-year observations across the three snapshot years). The mediator is the pair of functional control binaries (combustion and post-combustion).

Table 2.5: Descriptive statistics for primary analytical variables (coal + gas, ≥25 MW control-measured fleet; 2009/2018/2021 snapshots).

	Mean by year

[13] A further omitted control is boiler firing type, which strongly shapes the unabated NO_x rate (cyclone units high, fluidized-bed low, wall- and tangentially-fired in between) but is defined only for boiler-equipped plants and so cannot enter the full-fleet primary without discarding the gas turbines that make up much of the sample. Omitting it biases the organizational coefficients only insofar as firing type is correlated with organizational structure net of the controls already in the model. That correlation is likely small: firing type is a legacy engineering choice fixed at construction and driven by coal rank, plant vintage, and region rather than by ownership form, and the specification already conditions on coal heat share, plant scale, generator age, and prime-mover type, which absorb most of what firing type co-varies with. The residual scope for bias is thus confined to the coal subsample and is of indeterminate sign; the more likely consequence is a modest loss of precision, since a strong predictor of the outcome is left in the error term. Two things bound the concern. Most sample variation in NO_x intensity is across fuel, control adoption, and scale, all held constant, so the unexplained firing-type variance is a small share of the total; and the coal-boiler deep-dive, estimated where firing type is observed, does control for it, so if the organizational coefficients there track those in the primary, the omission is not distorting them.

Variable	N	2009	2018	2021	SD	Cover %
Outcome	—	—	—	—	—	—
NOx intensity (lb/MWh)	4,793	1.62	1.368	1.233	1.641	99.5
log NOx intensity	4,793	-0.193	-0.515	-0.63	1.482	99.5
Organizational Structure	—	—	—	—	—	—
Merchant / IPP, reference (share)	4,818	0.405	0.375	0.369	0.486	100.0
Regulated IOU (share)	4,818	0.256	0.252	0.258	0.436	100.0
Cooperative (share)	4,818	0.057	0.07	0.07	0.247	100.0
Government (share)	4,818	0.136	0.142	0.138	0.346	100.0
Host-serving (share)	4,818	0.146	0.16	0.166	0.364	100.0
Number of ownership entities	4,818	1.256	1.24	1.219	0.993	100.0
State market-share index (std.)	4,818	0.027	-0.027	0	0.967	100.0
Firm scale index (std.)	4,818	0.012	-0.009	-0.003	0.973	100.0
NOx Control Technology (functional)	—	—	—	—	—	—
Combustion control (0/1)	4,818	0.771	0.769	0.769	0.421	100.0
Post-combustion control (0/1)	4,818	0.386	0.486	0.501	0.498	100.0
Any NOx control (0/1)	4,818	0.819	0.846	0.85	0.369	100.0
Capacity to emit	—	—	—	—	—	—
log(Plant nameplate capacity, MW)	4,818	5.518	5.614	5.564	1.271	100.0
Avg generator age (yrs)	4,818	24.88	29.083	30.468	16.523	100.0
Capacity factor	4,818	0.327	0.303	0.322	0.265	100.0
Coal heat share	4,818	0.279	0.181	0.15	0.389	100.0
Steam turbine (0/1)	4,818	0.517	0.473	0.445	0.5	100.0
Gas turbine (0/1)	4,818	0.408	0.426	0.443	0.494	100.0
Regulatory controls	—	—	—	—	—	—
State climate-policy count (0–7)	4,818	2.538	2.879	3.207	2.008	100.0
Political-economic context	—	—	—	—	—	—
ISO/RTO member (0/1)	4,818	0.651	0.641	0.639	0.479	100.0
Democratic-lean index (std.)	4,817	0.129	-0.187	0.058	0.82	100.0

Note. Primary sample n = 4,818 plant-years. Bold rows are group separators. Means by year use each snapshot's plant-years. Cover % is the share of plant-years with non-missing values.

2.3.2 Primary Results

This section discusses the primary results derived from the models specified earlier, in Section 2.2.2. To ground readers in the purpose of the inquiry conducted here, the research questions are restated:

1. Does organizational structure predict plant-level NO_x emission intensity (lb/MWh), net of a plant's regulatory obligations, fuel mix, and capacity to emit?
2. Does organizational structure predict which NO_x controls a plant adopts, and does any gradient concentrate on the capital-intensive post-combustion margin rather than the inexpensive combustion one?
3. Does the organizational association with intensity run through adoption, or does it persist among plants carrying the same controls?

The results are discussed in the order the equations are set up. Equation 1a and 1b addressing adoption come first: organizational structure barely moves whether a plant installs control on either functional margin, and what little signal there is lives in the gap between the two margins, the cost-of-capital signature, cleanest for government ownership, which is less likely to carry the cheap combustion control but more likely to carry capital-intensive post-combustion treatment (Table 2.6, Figure 2.6). Equation 3, the question of direct organizational effect on NO_x emission intensity, follows: does organizational structure predict emission intensity at all? Selectively, and with one large surprise, regulated investor-owned utilities run markedly dirtier than merchants, the reverse of what the cost-of-capital logic predicted, while larger firms run cleaner, though the association does not reach significance. Finally, the section compares coefficients across Equation 1a and Equation 1b, a stagewise comparison, which asks whether those intensity associations run through what a plant adopts. What the section covers is the headline finding that, mostly, intensity associations do not run through adoption, surviving almost intact once the adopted-control indicators enter (Table 2.7, Figure 2.7). This result points to differences in the underlying fleet rather than in plant-level abatement behavior. Throughout, the coefficients are associations conditional on the controls, not causal effects.

2.3.2.1 Adoption margins: combustion vs post-combustion

Table 2.6 reports the organizational-structure coefficients on the two adoption decisions: whether a plant operates combustion-modification control (Equation 1a) and whether it operates post-combustion control (Equation 1b). Both are logistic regressions on the full fossil fleet, so the coefficients are log-odds. Whether an organizational characteristic moves the capital-intensive post-combustion margin more than the cheap combustion margin can be read by examining the contrast between the two columns; this is the $\beta_1^p - \beta_1^c$ differential that isolates the cost-of-capital mechanism from the others. All coefficients reported in Table 2.6 are log-odds on a binary adoption indicator and carry no lb/MWh reading.

On the combustion-modification margin (Equation 1a), organizational structure does little. Combustion controls (e.g., low-NO_x burners, overfire air, dry-low-NO_x combustors) are cheap and nearly ubiquitous, and the organizational coefficients reflect that: most are statistically indistinguishable from the merchant reference. The clearest signals are that government-owned plants are *less* likely to operate combustion control than merchants (-0.44 log-odds), as are PURPA cogeneration facilities (-0.39), consistent with their smaller, process-driven configurations. The cost-of-capital contrast is silent here: regulated IOUs are no more or less likely than the market-disciplined merchant reference to adopt control on the cheap margin (-0.10), and cooperatives and host-serving plants are likewise indistinguishable from it. This is what the framework predicts, since capital access should not govern a low-capital decision.

The post-combustion margin (Equation 1b) is the capital-intensive decision. Both SCR and SNCR carry large up-front capital costs, so this is where organizational differences in financing access should theoretically make themselves apparent. Mostly, they do not. Government ownership flips sign from the combustion margin: government-owned plants are now *more* likely to operate post-combustion control than merchants (0.54 log-odds), consistent with the notion of cheap public capital making the expensive option affordable. PURPA cogenerators move the other way, and they move more strongly than on the cheap margin (-0.61): their smaller, process-oriented plants are the least likely in the fleet to carry post-combustion equipment. The contrast theorized in Section 2.1.3 to most directly influence capital access, between regulated utilities and merchants, is insignificant here: regulated IOUs are statistically indistinguishable from the merchant reference (-0.10), as

are cooperatives. That the predicted cost-of-capital margin shows minimal organizational signal (government aside) is the first hint that the effects of IOU and merchant structures on plant-level decision-making, to the extent that they exist, do not run through whether a plant adopts.

The point of estimating the margins separately is to read the two coefficient estimates against each other. The gap between a trait's post-combustion and combustion coefficients (i.e., the $\beta_1^p - \beta_1^c$ differential) says *where* organizational difference is predicted to influence the adoption decision. A wide positive gap concentrated on the capital-intensive margin is the cost-of-capital signature. By contrast, a small gap that moves both margins together is the eco-structural signature. Only one trait displays a cost-of-capital signature: government ownership swings from -0.44 on the cheap combustion margin to 0.54 on the capital-intensive one. This is the lone case where both coefficients are individually significant and the gap runs in the predicted direction. The contrast built to carry the cost-of-capital test does not deliver it: regulated-IOU status, predicted to show the largest positive differential of any variable, is in fact flat and non-significant on both margins (-0.10 and -0.10), so the regulated-versus-merchant gap the design was built to read is essentially zero. Figure 2.6 ranks the traits by this gap, but outside government the ordering rests on coefficient pairs too imprecise to separate. Therefore, the differential chart is best interpreted as displaying one firm result against a field of nulls, not a clean rank order (despite the appearance). The takeaway is very specific: the cost-of-capital channel appears to impact whether a plant adopts only for government owners; for the regulated-versus-merchant contrast the framework leans on most heavily, the adoption margin is silent, and any emission difference between them must operate after the adoption decision. This question of the post-adoption decision-making process as it impacts NO_x emission intensity is the question the outcome equations take up next.

Table 2.6: Organizational-structure coefficients on the two adoption margins, full fossil fleet (logistic regression, log-odds): combustion modification (Equation 1a) versus post-combustion treatment (Equation 1b).

	Adoption (logit, full fleet)	
	(1a) Combustion	(1b) Post-combustion
Government	-0.443* (0.216)	0.535* (0.229)
Regulated IOU	-0.102 (0.242)	-0.097 (0.247)
Cooperative	-0.033 (0.347)	-0.500 (0.310)

Host-serving	-0.181 (0.207)	-0.445† (0.237)
Number of owners	0.114 (0.092)	-0.091 (0.061)
Third-party ownership	0.273 (0.213)	0.080 (0.170)
FERC cogeneration	-0.392* (0.176)	-0.609** (0.187)
FERC small power	-0.312 (0.349)	-0.324 (0.368)
State market-share index	-0.141 (0.129)	-0.024 (0.120)
Firm scale index	-0.076 (0.107)	-0.026 (0.094)

Note. CR1 standard errors clustered on utility in parentheses. Plant, fuel, state-policy, and partisan controls and year and region fixed effects are included but omitted from the table. n = 4,817. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

Figure 2.6 plots the same two adoption coefficients for each organizational trait, with the connecting segment marking the $\beta_1^p - \beta_1^c$ differential: the wider the rightward gap, the more strongly the trait loads on the capital-intensive post-combustion margin relative to the cheap combustion margin.

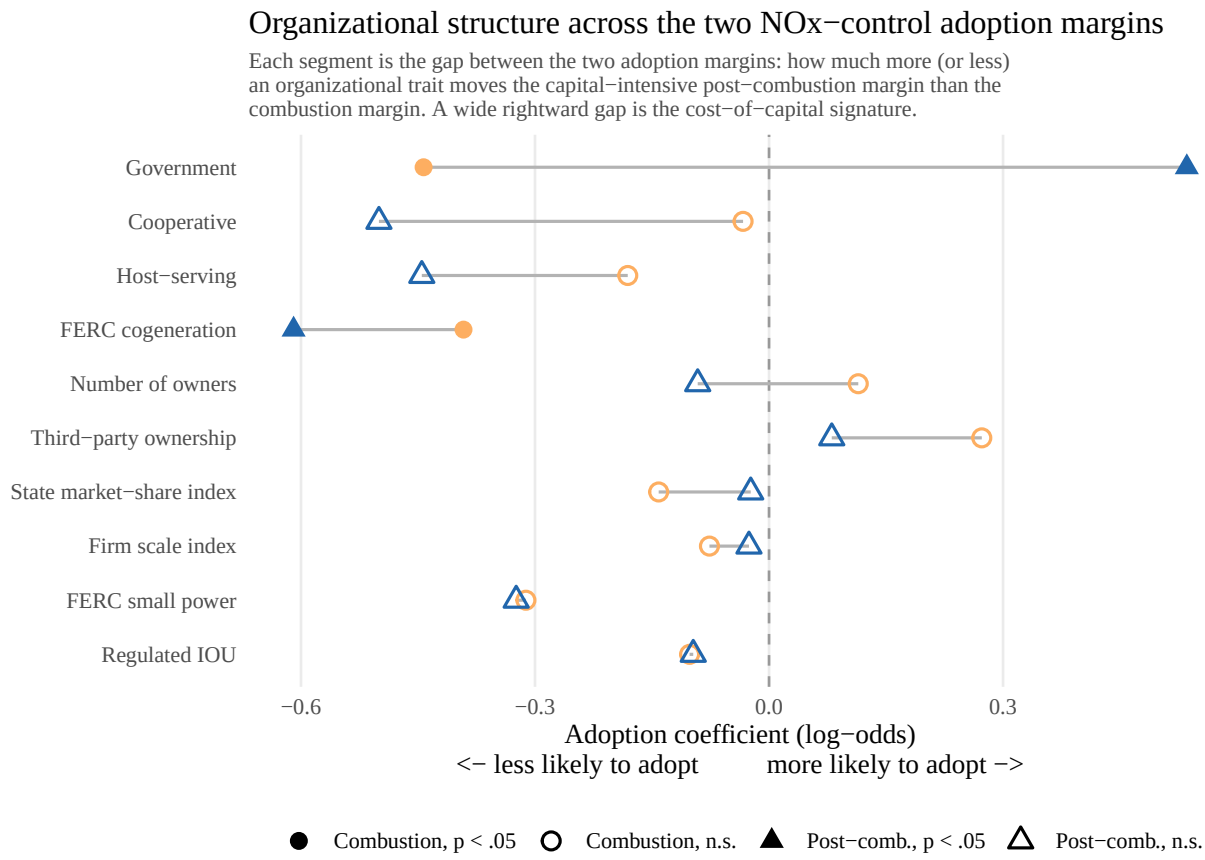

Figure 2.6: Organizational-structure coefficients on the two NO_x-control adoption margins (full fossil fleet, logistic log-odds). Orange circles: combustion modification (Eq 1a); blue triangles: post-combustion (Eq 1b); the connecting segment is the post-minus-combustion differential, by which rows are sorted. Filled points are significant at $p < .05$.

2.3.2.2 Direct association: does organizational structure predict intensity?

Before asking whether organizational structure's emission association runs *through* adoption, the original research question of this chapter bears answering: does organizational structure predict NO_x intensity? The answer is selectively. Among the organizational-type contrasts, government ownership and host-serving status return outcome coefficients indistinguishable from the merchant reference. Four traits move emissions clearly, all dirtier. Regulated IOUs run markedly dirtier than merchants (0.26 log points, roughly 30% higher intensity), opposite the cost-of-capital prediction that rate-base recovery should make them cleaner; cooperatives are dirtier by almost the same margin (0.32, roughly 37% higher); PURPA cogenerators are also dirtier (26%, though the per-MWh

accounting bias of Section 2.1.3 inflates it); and more dispersed ownership is modestly dirtier (7% per additional owner). One further association is marginal ($p < .10$): firms with a larger state-market share are suggestively dirtier. Larger firms, for their part, point cleaner (-6% per standard deviation), the cost-of-capital direction that runs opposite the SuperPolluter expectation that bigger firms are behaviorally dirtier per unit, but the association does not reach significance. The answer to the first research question is therefore selective: organizational structure shapes emission intensity for some forms and not others, and the two cost-of-service forms most associated with per-unit dirtier behavior, regulated IOUs and cooperatives, point in a direction the framework did not predict (the lower cost of capital from rate-base recovery did not improve NO_x emission intensity, as theorized).

2.3.2.3 Attenuation across the adoption margin

Table 2.7 places the full-sample outcome coefficients with and without the two adopted-control indicators side by side. A coefficient that shrinks from the total column to the conditional column operates through the adoption decision; one that persists operates through other channels.

Having seen in Section 2.3.2.2 which traits are directly associated with greater NO_x emissions intensity, the payoff to the stagewise design is whether those associations run *through* adoption. The picture is mixed. Among the traits with a meaningful direct association, most attenuate at least partially once the two adopted-control indicators enter (Figure 2.7), indicating that part of why these plants run dirtier is that they adopt control less often. More dispersed ownership shrinks 0.06 to 0.05; the cooperative penalty shrinks by about a fifth (0.32 to 0.25); PURPA cogeneration attenuates most sharply, from 0.23 to 0.11 (losing significance); and the marginal state-market-share association also shrinks (0.06 to 0.04), though from a base never distinguishable from zero. The one significant association that conditioning leaves essentially intact is the regulated-IOU penalty. (Government and third-party ownership move more in proportional terms, but off small, non-significant bases, so little rests on them.) The lesson is that adoption carries part, but not all, of how organizational structure shapes emission intensity: for most traits some of the association runs through the binary decision to install control, while the IOU penalty operates almost entirely through other channels.

The most consequential results are the two cost-of-service forms, regulated-IOU status and cooperative form, which carry the largest organizational coefficients in the outcome model. They

differ in how much of the penalty runs through adoption. The regulated-IOU penalty survives conditioning almost intact: IOUs run dirtier by 0.26 log points in the total specification and 0.24 once the two adopted-control indicators are added, roughly 27% higher NO_x intensity, so the penalty is not an adoption-selection artifact (at adoption IOUs were statistically indistinguishable from merchants). The cooperative penalty is nearly as large (0.32 total, 0.25 conditional) and remains highly significant, but about a fifth of it runs through lower adoption, so for cooperatives adoption-selection explains part, though not most, of the gap. For both forms a substantial penalty persists once adoption is held constant. This is the opposite of what the cost-of-capital pillar predicted: H1b expected rate-base recovery to make capital-intensive abatement easier for regulated IOUs and so leave them cleaner, yet they, alongside cooperatives, are the dirtiest organizational forms in the sample.

The intuitive explanation is composition, that regulated IOUs simply remain the custodians of the aging coal baseload fleet. The evidence does not support it. Splitting the sample by fuel, the IOU penalty appears among gas-fired plants (+0.293**) and is undetectable among coal-fired plants (+0.059); the two subsamples are not statistically distinguishable (interaction $p = 0.96$), but either way there is no sign the penalty is a coal phenomenon. The primary specification also absorbs NERC-region fixed effects, so it is not a matter of where IOUs operate.[14] What remains, once composition, geography, and the adoption and operation of controls are absorbed, points to a genuine organizational difference between cost-of-service utilities and merchants. The interpretation this chapter favors is that cost-of-service recovery, by passing fuel and operating costs through to ratepayers, dulls the operating-cost discipline a merchant generator faces directly, allowing the plant to run less efficiently per MWh. Additional evidence aligns with this interpretation. Cooperatives carry essentially the same penalty as IOUs despite having no shareholders earning a return on rate base. So the increased NO_x emissions intensity associated with cooperatives actually points away from a pure Averch-Johnson over-capitalization motive and toward the broader cost-recovery, weak-market-discipline mechanism that IOUs and cooperatives share. One caveat on reading these contrasts: because the type factor takes the merchant generator as its reference, the regulated-IOU and cooperative coefficients are *cost-of-service-versus-merchant* comparisons rather than "regulated

[14] And the coal-boiler deep-dive (Section 2.7) shows the penalty survives conditioning on which post-combustion technology a plant runs (SCR versus SNCR) and how intensively it operates it (annual hours in service).

versus unregulated" in general.

Larger firms run cleaner by -0.07 log points per standard deviation conditional on adoption (the index is modestly left-skewed, with a median of −0.14, reflecting that most plants belong to smaller firms, but it spans roughly four standard deviations without extreme outliers, so the per-SD coefficient is not driven by a handful of large-firm observations). This amounts to roughly -6% lower NO_x intensity. The association is not statistically significant in the full sample, and what signal there is barely moves between the total and conditional specifications.[15] The conditioning test is the same one applied to the IOU penalty, only on a far weaker coefficient: the firm-scale signal survives conditioning on the abatement behaviors the data can observe. Here, though, that survival points to fleet composition, the kinds of plants large firms own, rather than to a distinct operating regime or more aggressive abatement at a given plant. The direction is nonetheless the one the cost-of-capital pillar implied and the reverse of the SuperPolluter prediction that larger, gain-driven firms are dirtier per unit (a behavioral story). As established earlier, this is a per-unit result and says nothing about aggregate emissions, where large firms remain the largest emitters by virtue of size.

Table 2.7: Organizational-structure coefficients on log NO_x intensity, full sample, before and after conditioning on the two adopted-control indicators.

	Total	Conditional
Government	-0.136 (0.098)	-0.091 (0.091)
Regulated IOU	0.262** (0.088)	0.241** (0.076)
Cooperative	0.318*** (0.094)	0.247*** (0.075)
Host-serving	-0.084 (0.112)	-0.174 (0.111)
Number of owners	0.063** (0.022)	0.051** (0.019)
Third-party ownership	-0.054 (0.081)	-0.022 (0.081)
FERC cogeneration	0.227** (0.076)	0.108 (0.075)
FERC small power	0.444 (0.293)	0.395 (0.309)
State market-share index	0.055† (0.030)	0.042 (0.026)
Firm scale index	-0.062 (0.045)	-0.066 (0.044)

[15]The deep-dive shows it is not about abatement-behavior either: larger firms are no more likely to run SCR over SNCR and no different in how many hours they operate their controls, yet they remain directionally cleaner within the adopter subsample (-0.08 log points), though that estimate does not reach significance either.

Note. Dependent variable: log NOx intensity. CR1 standard errors clustered on utility in parentheses. Both columns include plant, fuel, state-policy, and partisan controls and year and region fixed effects; the conditional column adds the combustion and post-combustion adoption indicators. n = 4,792. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

Figure 2.7 reports the conditional outcome coefficients (Equation 3b) with their 95% confidence intervals, ranked by magnitude; navy marks significance at $p < .05$. The amber open ring on each row marks the total association (Equation 3a), so the gap between the ring and the filled dot is the part of the organizational association that runs through the adoption decision; rings that sit almost on top of their dot are coefficients that barely move under conditioning.

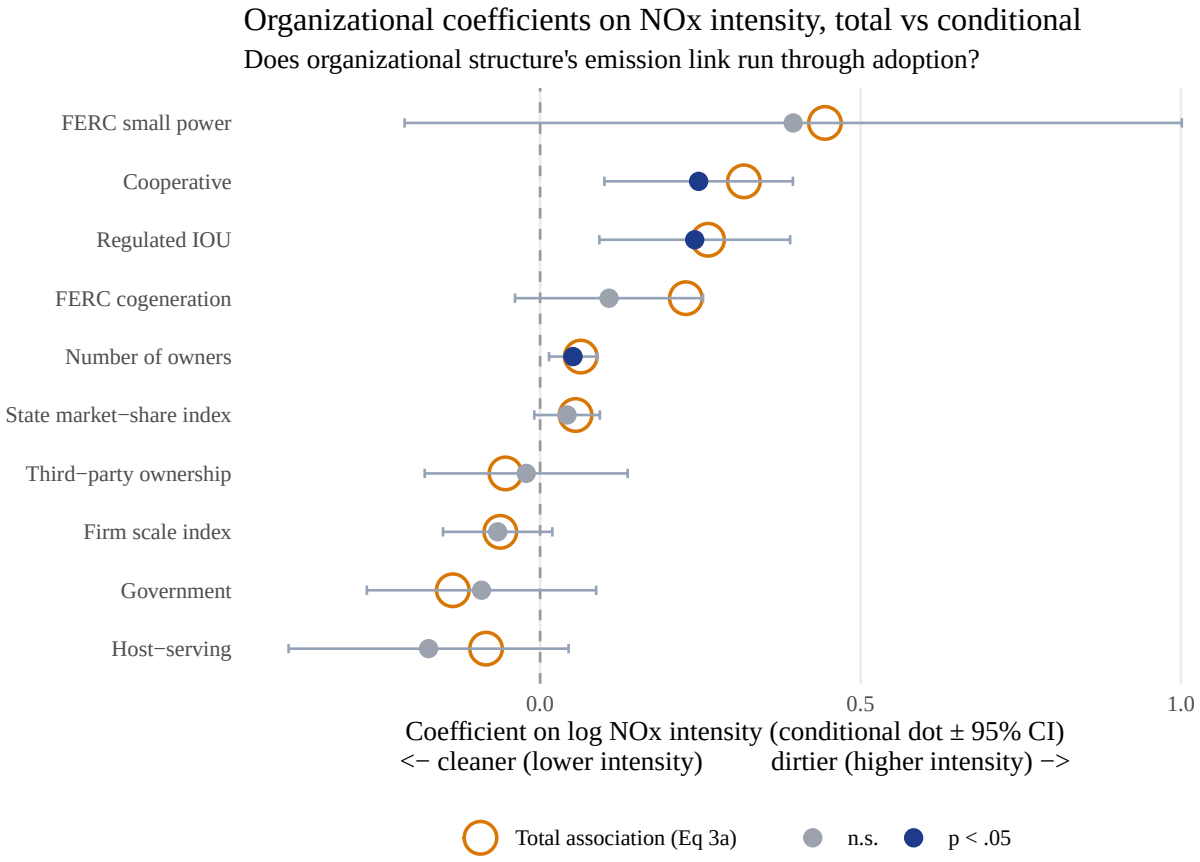

Figure 2.7: Organizational-structure coefficients on log NO_x intensity (full sample), ranked by magnitude. Filled dots: conditional estimates (Eq 3b) with 95% confidence intervals, navy where significant at $p < .05$; amber rings: total associations (Eq 3a). The ring-to-dot gap is the share of the association running through adoption; left of zero is cleaner.

2.3.2.4 Hypothesis scorecard

Table 2.8 collects the predicted and observed signs for each hypothesis across the two adoption margins, their differential, and the outcome before and after conditioning on adoption. This table synthesizes the raw predictions and results in a visually ingestible format. Section 2.4 will adjudicate the hypotheses with reference to this table.

Table 2.8: Hypothesis scorecard: predicted versus observed signs for each hypothesis across the two adoption margins (Eq 1a, 1b), their differential, and the total and conditional outcome (Eq 3a, 3b).

	(1a) Comb.		(1b) Post		Diff.		(3a) Total		(3b) Cond.	
Organizational trait	Pred.	Obs.	Pred.	Obs.	Pred.	Obs.	Pred.	Obs.	Pred.	Obs.
H1 - Organizational type										
Government	+	**-***	+	**+***	>= 0	**+0.98**	-	**-**	-	**-**
Regulated IOU	+	**-**	+	**-**	» 0	**+0.00**	-	**+****	-	**+****
Cooperative	+	**-**	+	**-**	>= 0	**-0.47**	-	**+*****	-	**+*****
Host-serving	n.p.	**-**	n.p.	**-†**	n.p.	**-0.26**	+	**-**	+	**-**
H2 - Ownership concentration										
Number of owners	-	**+**	-	**-**	> 0	**-0.21**	+	**+****	+	**+****
H3 - Cross-cutting status flags										
Third-party ownership	n.p.	**+**	n.p.	**+**	n.p.	**-0.19**	n.p.	**-**	n.p.	**-**
FERC cogeneration	n.p.	**-***	n.p.	**-****	n.p.	**-0.22**	n.p.	**+****	n.p.	**+**
H4 - Market position										
State market-share index	-	**-**	-	**-**	~ 0	**+0.12**	+	**+†**	+	**+**
H5 - Firm scale (super-polluter)										
Firm scale index	-	**-**	-	**-**	> 0	**+0.05**	+	**-**	+	**-**

Note. **Pred.** = theoretical prediction (gray); **Obs.** = observed sign read live from the fitted model (bold), starred at the levels below. The differential is the post-combustion minus combustion adoption coefficient (log-odds), shown as a signed number with no significance test (the two margins are separate logits). `n.p.` = not predicted (exploratory). Predicted signs are shown bare; their strength qualifiers and theoretical pillars appear in 2.2. Full coefficients and standard errors appear in 2.6 and 2.7. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

2.4 Discussion

This chapter set out to answer three questions: whether organizational structure predicts plant-level NO_x emission intensity, whether it predicts which controls a plant adopts, and whether any association with intensity runs through adoption or persists among plants carrying the same controls. The results deliver a two-part answer. First, organizational structure barely predicts *whether* a plant adopts NO_x emission control technology. Across both combustion and post-combustion adoption margins, most organizational coefficients are indistinguishable from zero. The only clean signal is government ownership, and it only manifests on the capital-intensive post-combustion margin (combustion coefficient -0.44 versus post-combustion coefficient 0.54). Where organizational structure appears to matter is the consequential outcome of NO_x emissions. The regulated-IOU characteristic, with cooperatives close behind, carries one of the two largest penalties (0.24, roughly 27% higher intensity) and barely moves when adoption is held constant. Firm scale, by contrast, points the other way, cleaner but not significantly so. The takeaway is hard to see, but can be compressed into this nugget: organizational differences act less on the binary decisions to install controls and more on the (admittedly, harder to measure) dimension of what happens after the control is installed: how hard the equipment is run, how the plant is operated, and, pre-control, the legacy of the broader fleet a firm carries. So, the contribution of this chapter is less about answering its titular question of 'Who cleans up?', and more about 'Who runs dirtier given the same controls?' The rest of this section will unpack this question by adjudicating the theoretical pillars underpinning the analysis (i.e., where the cost-of-capital actually influences decision-making and emissions outcomes) and narrativizing some of the analysis' stronger (and unexpected) findings. All coefficients are interpreted as associations, not causal effects. Also, Table 2.8 offers a tabular summary of the discussion that follows.

If adoption is where organizational structure leaves its fingerprints, the easiest way to lift them is to measure the distance between the two margins, the $\beta_1^p - \beta_1^c$ differential. The logic is that a positive gap, concentrated on the capital-intensive post-combustion margin, is the cost-of-capital fingerprint (financing access should matter most where the equipment is most expensive). By this test, the cost-of-capital pillar distinguishes government ownership, specifically: government-owned plants swing from -0.44 on the cheap combustion margin to 0.54 on the expensive one, both individually

significant. Adoption, however, does not necessarily reduce emissions. Government's outcome coefficient is indistinguishable from zero (-0.14 total, -0.09 conditional), so public owners adopt the expensive control more often and yet are no cleaner per MWh.

And government is the only place the cost-of-capital effect turns up at all. The two contrasts the design was actually built to test come back empty. Regulated-IOU status, predicted to carry the biggest positive differential of any variable, is flat and non-significant on both margins (-0.10 and -0.10), so the regulated-versus-merchant gap the design was built to read is essentially zero. So beyond the lone government result, the adoption margin does not really adjudicate the pillars. What it does tell us is narrow but useful: for the regulated-versus-merchant axis the framework leans on most heavily, *whether* a plant adopts is silent, and any emission difference between these forms must therefore operate after the decision to install, which is exactly what the outcome equations take up next.

The single most consequential result is also the one the framework got most wrong. The cost-of-capital pillar did not merely mispredict the regulated-IOU coefficient; it predicted the opposite sign. H1b reasoned that rate-base recovery, by letting regulated utilities pass approved capital spending through to ratepayers, should lower their effective cost of abatement and leave them cleaner than merchants financing the same equipment at market rates. Instead, regulated IOUs, alongside cooperatives, are the dirtiest organizational forms in the sample (0.24, roughly 27% higher intensity). The puzzle, then, is not whether the association is real but why the financing advantage that was supposed to make IOUs cleaner coincides with their running dirtier.

The reading this chapter settles on is that cost-of-service regulation cuts two ways: rate-base recovery lowers the effective cost of abatement capital (increasing adoption), but cost pass-through *also* blunts the operating-cost discipline a merchant generator selling into a competitive market faces directly. In other words, what separates the two organizational forms is how efficiently their plants are run. NO_x intensity ($\text{lb } NO_x/\text{MWh}$) is the combustion rate per unit of fuel ($\text{lb } NO_x/\text{MMBtu}_{\text{fuel}}$) scaled by heat rate ($\text{MMBtu}_{\text{fuel}}/\text{MWh}$), so a less efficient plant emits more per MWh at the same combustion rate, and the rate worsens further at low load and under cycling. A merchant pays for every wasted unit of fuel out of its own margin, so it has a direct reason to hold heat rate down, dispatch economically, and idle or retire its least efficient units. By contrast, a regulated utility

passes most fuel cost through to ratepayers and earns its return on the asset whether or not it runs lean, so that discipline is weaker.[16] So within regulation the two forces the framework had folded into a single cost-of-capital channel actually pull in opposite directions, and on the per-MWh margin the operating-discipline effect dominates the financing one. This pattern is the environmental analogue of a long-standing critique of rate-of-return regulation, that cost-plus recovery dulls the incentive to minimize cost (the tradition Averch-Johnson launched (Averch and Johnson 1962; L. L. Johnson 1973)), measured here in pollution rather than dollars. Two caveats, already noted, keep this an association rather than a verdict: the contrast is IOU-versus-merchant rather than regulated-versus-unregulated in general, and because regulated status barely varies within a firm, the design cannot difference out an unobserved firm-level characteristic that might travel with it.

The verdict for cooperatives is the similar to regulated IOUs. H1c predicted member-owned utilities would run cleaner than merchants on the strength of subsidized capital; instead they are among the dirtiest organizational forms in the sample (0.25, roughly 28% higher conditional on adoption), fully matching the regulated-IOU penalty, so H1c is not supported in its predicted direction. The failure is informative, though: cooperatives recover costs like IOUs but have no shareholders earning a return on a rate base, so the Averch-Johnson over-capitalization motive cannot be driving them. That a form without the capital-bias incentive carries the same penalty as one with it points past Averch-Johnson rate-recovery-driven over-capitalization to the broader mechanism the two share, cost pass-through and weak market discipline dulling the incentive to run lean. The cooperative result therefore converges with H1b, and that convergence is itself the evidence: the operating-discipline reading of the cost-of-service penalty does not rest on the rate-base return unique to IOUs.

More dispersed ownership is the other organizational feature that moves emissions in its own right, and its direction is the one H2 predicted. Plants with more owners run dirtier (0.06 per additional owner in total, 0.05 conditional on adoption), a modest but statistically reliable per-owner penalty. H2 was a horse race, eco-structural asset inertia and principal-agent coordination friction against the capital-pooling advantage of shared finance. The association attenuates only partially once adoption

[16] The deep-dive points the same way from the other side: the cost-exposed merchant tends to run its controls fewer hours and still ends up cleaner, which says the difference lives in how the plant is operated, not in how hard the scrubber is worked.

is held constant, so part of the penalty is that multi-owner plants install control somewhat less often, while most of the association persists, pointing again to operation and inertia rather than the adoption decision alone. However, two notes of caution are in order. The individual adoption-margin coefficients are not significant on their own, so the verdict rests on the outcome equations alone, and because nearly all the identifying variation sits in the minority of jointly owned plants, the estimate speaks mostly to that subset rather than the fleet at large.

The super-polluter claim (H5) predicted that larger firms would run *dirtier* per unit, on the logic that organizational scale breeds principal-agent slack and the eco-structural insulation that lets a big firm absorb the reputational and regulatory cost of a dirty plant. The data point the other way. Firm scale is associated with *lower* intensity (-0.06 total, -0.07 conditional), though the association does not reach significance; the sign is the cost-of-capital direction rather than the super-polluter one: bigger firms, if anything, run a touch cleaner per MWh. H5 is not supported, therefore. However, two qualifications are necessary here. First, this is a result about *intensity*, not mass: larger firms own more and bigger plants, so they remain the largest aggregate NO_x emitters in absolute tons, and the disproportionality intuition that motivates the super-polluter frame still holds at the level of total emissions. Second, the firm-scale coefficient is not statistically significant in the full sample, so the honest reading is a failure to find scale running dirtier, not a clean demonstration that it runs cleaner.

The remaining dimensions add little to the picture. Market position (H4) points dirtier as predicted, the eco-structural sign, but only marginally in total (0.06) and not at all once adoption is held constant. Moreover, this measure is dominated by a few near-monopoly plants holding the right tail and shares are only measured at the reporting-entity (`utility_id`) rather than the parent level (discussed in Section 2.5), the result is as much a limit of measurement as a finding. PURPA cogeneration (H3b) is dirtier in total (0.23) but the association is largely an adoption story, since these plants are the least likely in the fleet to carry control and the coefficient fades to insignificance once adoption enters (0.11). Interpretation of the results for H3b is further confounded by the combined-heat-and-power accounting convention that inflates measured intensity, so the exploratory framing stands and no clean behavioral reading is available. Third-party-operated plants (H3a, exploratory) are statistically indistinguishable from merchants throughout. Host-serving generators (H1d, also exploratory) are likewise indistinguishable on the outcome, but are marginally less likely than merchants to adopt

post-combustion control (-0.45, $p \approx .06$). None of these dimensions changes or nuances the central result that cost-of-service forms (in addition to ownership dispersion) are the primary correlates of dirtier NO_x emission outcomes.

Stepping back from the individual verdicts, a pattern emerges. Wherever organizational structure moves NO_x emission intensity, it does so downstream of the adoption decision: the cost-of-service forms, ownership dispersion, and the asset inertia behind them all register in the outcome equations and survive conditioning on what a plant has installed, while the one dimension that clearly moves adoption, government ownership, leaves emissions untouched. In other words, organizational structure mostly does not touch adoption margin; the outcome margin is where it leaves its fingerprints. Rearranging the framework's pillars in light of the evidence: the cost-of-capital pillar, around which most of the tests were structured, proved to fairly immaterial to adoption and, where it appeared to influence matters, it pointed the wrong way for the outcome; meanwhile, the eco-structural and principal-agent pillars seem to carry what little signal there is, but through operation and asset inertia rather than the financing of equipment. The cost-of-service-discipline reading developed for regulated IOUs is thus generalizable: cooperatives share that operating-discipline channel, and ownership dispersion adds a parallel operation-and-inertia channel, so the chapter's central claim is that organizational influence on emissions stems from how plants are run, not in whether they are equipped to control emissions, a claim established for the snapshot years analyzed here (2009, 2018, and 2021) and extended beyond them only as far as the regulatory and market regimes they bracket persist.

These results speak to three literatures at once. They confirm the eco-structural, organizational political economy claim that organizational structure leaves a measurable imprint on environmental outcomes (Prechel, Blazek, and Amaral 2025), but they nuance this claim by locating that imprint on a specific margin, operation rather than the adoption of control. They complicate the SuperPolluter thesis: large firms are not dirtier per unit of output, so the super-polluter signature is strictly a claim about mass, the sheer quantity of generation a large firm commands, not about the rate at which any one plant pollutes. That reading squares with prior plant-level evidence, where larger plants emit more in total while the per-MWh rate goes untested (Grant, Bergstrand, and Running 2014). That distinction reconciles the present null with the disproportionality tradition (Freudenburg 2006; Pellow

2000), in which a few actors account for an outsized share of total emissions: disproportionality can hold at the level of aggregate tonnage even where per-MWh intensity does not increase with firm scale.[17] Finally, the regulated-utility result ties the organizational-structure literature to the economics of dispatch and operating discipline (Fowlie 2010; Borenstein and Bushnell 2015), extending what those studies show about market exposure and operating efficiency from cost and dispatch to the emissions rate itself.

2.5 Limitations

Several limitations bound what these results can claim. The design is associational rather than causal: the analysis pools three cross-sections with year and NERC-region fixed effects and clusters standard errors on the utility, but organizational structure is not randomly assigned, so every coefficient is a conditional association, not a treatment effect. Furthermore, regulated and government status barely vary within a firm, so parent-firm fixed effects cannot be used to difference out unobserved firm-level heterogeneity behind the IOU and government results. Those magnitudes are inherently between-firm comparisons, and an unmeasured firm-level characteristic that travels with organizational form cannot be excluded.

A related omission sits not at the firm level but at the level of place. The specification carries no measure of local air-quality regulation, and one such measure has teeth here: plants in ozone- or $PM_{2.5}$-nonattainment areas face NO_x control obligations under Reasonably Available Control Technology (RACT) that their attainment-area counterparts do not (Richards 2000), and nonattainment designations skew toward urban airsheds. This is not the geography the NERC-region fixed effects and the state partisan-lean index already absorb. Nonattainment is a county- or area-level designation that varies *within* both NERC region and state, so the regional and state-level components of regulatory pressure are conditioned out while its sub-regional, county-level component is not. The concern is that organizational forms sort unevenly along the urban-rural gradient, with cooperatives markedly rural and merchants and government plants more urban, so the merchant-versus-cost-of-service

[17] It is the difference between a league's leading scorer and its most efficient one: a high-usage star tops the total-points table on sheer volume of shots without posting the best points-per-shot, just as a large firm tops the emissions tonnage on the scale of its generation without running any single plant dirtier per MWh.

contrast may be carrying a share of this unmeasured local regulatory difference. The direction is, at least, legible: because the cleaner group (merchants) leans urban and nonattainment-strict while the dirtier group (cooperatives) leans rural and attainment-lax, holding nonattainment status constant would compress rather than widen the gap. The organizational contrast is therefore best read as an upper bound on the purely organizational difference, not as the difference net of where these plants happen to sit.

Data availability imposes a second set of limits. The liability-shielding mechanism the framework names, parent-subsidiary structure and the legal entity forms that carry it (LLCs, special-purpose entities, holding companies), is not recorded in EIA-860 and enters only as a weak proxy through related ownership contrasts. The 25 MW capacity floor is itself a data-availability boundary: units below it generally fall outside continuous emissions monitoring, so eGRID estimates their NO_x from fuel use and emission factors rather than measuring it, which would make the outcome partly a construct of the same fuel and technology variables used as predictors, and utilization data for their NO_x controls are rarely reported. Flooring these units out removes that circularity, at the cost of narrowing the findings to the utility-scale fleet rather than the full population of emitters. Finally, EIA-860 does not record gas-turbine NO_x controls, and the EPA CAMD data that fill that gap cover mainly continuously monitored units, so control adoption too is measured most completely for the larger plants.

Third, the panel is open. Plants enter and exit across the three snapshots, and the closure diagnostics (Section 2.2.2.7) show that exiters and entrants differ systematically from survivors, with dirtier coal units leaving and cleaner gas units arriving. The pooled estimates therefore describe the fleet as observed at each snapshot rather than a fixed set of plants followed over time, and the compositional shift across snapshots is a feature of the sample to be read descriptively, not a within-plant trend.

Finally, the analysis covers NO_x from coal- and gas-fired plants only. Whether the same organizational pattern holds for other pollutants is an open question; SO_2 and CO_2 robustness analyses, which require their own datasets, are outside the scope of this analysis.

2.6 Conclusion

This chapter began from a plain question: who cleans up? Across the coal- and gas-fired fleet, which owners install the equipment that holds NO_x down, and which do not? The analysis returns an answer that redirects the question itself. Organizational structure barely predicts *whether* a plant adopts control; where it leaves its mark is on how cleanly the plant is run once equipped. The contribution, then, is to relocate where organizational structure registers, from the decision to install control to the operating margin downstream of it, and to name the channel doing the work: the weak operating-cost discipline of cost-of-service regulation. The animating question is better posed not as who cleans up, but as who runs dirtier given the same controls.

The empirical heart of that answer is a pair of cost-of-service forms. Regulated investor-owned utilities and cooperatives run dirtier than merchant generators, a regulated-IOU penalty on the order of 27% higher NO_x intensity, and the gap survives conditioning on what control a plant has installed. That it survives is the point: the gap is not a matter of which owners buy the equipment, but of how the equipment-bearing plant is run. Cooperatives carry a comparable penalty yet have no shareholders earning a return on a rate base, so the Averch-Johnson over-capitalization motive cannot be what drives it; what the two forms share is cost pass-through and the dulled operating-cost discipline it brings. Firm scale, the super-polluter expectation, shows nothing of the kind: larger firms are no dirtier per unit, and if anything run marginally cleaner. The super-polluter pattern does not dissolve so much as relocate—from the rate at which any one plant pollutes to the sheer mass a large firm commands.

The implication concerns where policy instruments seek purchase, and it is offered as a direction rather than a policy evaluation. Instruments aimed at abatement have mostly been aimed at adoption: subsidies and mandates that put equipment on site. If the operative lever is instead how the equipped plant is run, those instruments act upstream of where organization bites. The finding points toward the incentive structure of cost-of-service ratemaking itself, and toward performance-based forms of discipline that reach operation. Such obligations would be written on the emission rate a plant achieves, rather than on the equipment it carries.

Whether the dulled operating discipline documented here also tracks the race and income of the communities beside these plants, making the organizational mechanism an environmental-justice

mechanism as well, is the question the next chapter takes up.

Appendix — Variable inventory

Every variable in the analytical dataset is inventoried below, organized by conceptual role rather than by plant scope: the dependent variable, the organizational-structure predictors, the functional NO_x-control mediator, and the control blocks. Each row lists the variable name, a plain-English label, units, the primary data source, and a one-line construction recipe. Every variable describes the same coal- and gas-fired sample defined above; the few measures defined only for boiler-equipped plants (the coal-boiler deep-dive controls) are flagged where they appear.

2.6.1 Dependent Variables

Primary outcome: NO_x emission intensity (lb/MWh). Total NO_x mass (ln-transformed) is the level robustness check. SO_2 and CO_2 robustness analyses are separate documents to be added once the primary analysis is complete.

Table 2.9: Codebook: dependent variables.

Variable	Label	Units	Source	Operationalization
`nox_rate_lb_per_mwh`	NOx emission intensity (primary DV)	lb/MWh	eGRID PLNT	Direct field from eGRID PLNT; primary outcome
`ln_nox_rate_lb_per_mwh`	Log NOx emission intensity	log(lb/MV	constructed	log(nox_rate_lb_per_mwh); used in OLS specifications

2.6.2 Primary IV: Organizational Structure

2.6.2.1 Ownership and Regulatory Status

A plant's organizational type and economic-regulatory status enter as a single mutually exclusive factor, `org_structure_type` (Merchant/IPP reference), built from EIA-860 entity type crossed with regulatory status (Table 2.1). Ownership concentration (`n_owners`) and third-party operation

(has_third_party_ownership) enter separately, and the PURPA cogeneration and small-power certifications enter as cross-cutting flags.

Table 2.10: Codebook: ownership and regulatory status.

Variable	Label	Units	Source	Operationalization
org_structure_type	Organizational type (typology factor; primary regulatory-status IV)	factor (5 levels)	EIA-860 Sched 1 UTILITY	Mutually-exclusive ownership/regulatory type from entity_type × regulatory_status: Merchant (reference), Regulated IOU (entity_type I & RE), Cooperative (C), Government (F/M/S/P), Host-serving (IND/COM). Only investor-owned (I) splits by RE/NR; the ~35 pre-2013 plants lacking entity_type are recoded from sector so none are unclassified. Supersedes has_govt_ownership, is_regulated_utility, and ferc_ewg_binary in the primary spec.
n_owners	Number of distinct ownership entities	count	EIA-860 Sched 4 OWNER	Count distinct ‘ownership_id‘ values across all generators at plant
has_third_party_ownership	Third-party / leased ownership (code W)	0/1	EIA-860 Sched 4 OWNER	1 if any generator has ‘ownership_code == ’W’‘
ferc_cogen_binary	FERC cogeneration status (PURPA)	0/1	EIA-860 Sched 2 PLANT	Direct field, binarized
ferc_small_power_binary	FERC small power producer certification	0/1	EIA-860 Sched 2 PLANT	Direct field, binarized

2.6.2.2 Firm Scale and Market Position

Two composite indices enter the primary specification: a firm-scale index measuring the parent company's absolute portfolio size, and a state market-share index measuring relative power concentration normalized to the plant's state and year. Each replaces a pair of collinear raw measures with their standardized mean.[18]

Table 2.11: Codebook: firm scale and market position.

Variable	Label	Units	Source	Operationalization
`firm_scale_index`	Firm-scale index (composite; H5 primary)	z-score composite	constructed	Mean of the z-scored logs of owner_n_plants and owner_total_nameplate_capacity_mw; the H5 regressor in place of the two collinear scale measures. Mirrors state_market_share_index. Both raw measures retained for descriptives and robustness.
`state_market_share_index`	State market-share index (composite; H4 primary)	z-score composite	constructed	Mean of the z-scored log1p of state_capacity_share and state_output_share; the H4 regressor in place of the two collinear (r = 0.87) shares. Mirrors firm_scale_index. Right-skewed (a few near-monopoly plants). Both individual shares retained for descriptives.

[18] The firm-scale index is the standardized mean of the logs of `owner_n_plants` (parent firm total plants) and `owner_total_nameplate_capacity_mw` (parent firm total capacity); `log_owner_n_plants` and `log_owner_total_nameplate_capacity_mw` are the logged inputs. The state market-share index is the standardized mean of the `log1p` of `state_capacity_share` and `state_output_share`. The composites replace the collinear raw measures in the spec; all raw components are retained in the analytical dataset for descriptives and robustness. Firm-scale measures the parent company's absolute size across all utilities and plants owned, independent of state context; market-share measures relative power within the plant's state. Both are computable from Form 860 ownership records for every owner regardless of organizational type (IOU, IPP, municipal, cooperative, industrial cogenerator). A firm-revenue measure was considered and rejected: Form 861 reports only retail revenue, structurally absent for independent power producers and merchant generators, and even where present captures a regulator-allowed cost-recovery amount rather than a market-pursued revenue figure; the same numeric variable would mean materially different things across organizational types.

2.6.3 Capacity to emit

Baseline physical conditions that determine a plant's emissions potential independently of abatement choices: scale (`log_nameplate_cap_mw`), utilization (`capacity_factor`), generator age (`avg_generator_age_yrs`), prime-mover configuration (`has_steam_turbine`, `has_gas_turbine`), and fuel-mix composition (`coal_heat_share`). The fuel-mix share holds combustion composition constant across the mixed coal-plus-gas sample; one share suffices, because the heat shares sum to ≤ 1 (residual = oil and other fuels) and, for the ~96% of the sample burning only coal or only gas, coal share + gas share ≈ 1, so `gas_heat_share` is omitted.

Table 2.12: Codebook: capacity to emit.

Variable	Label	Units	Source	Operationalization
`log_nameplate_cap_mw`	Log nameplate generating capacity	log(MW)	constructed	log(total_nameplate_mw)
`capacity_factor`	Capacity factor (plant-level)	fraction 0–1	eGRID PLNT	Net generation ÷ (nameplate × 8,760 hrs).
`avg_generator_age_yrs`	Average generator age	years	EIA-860 Sched 3 GENERATOR	‘obs_year − mean(operating_year)‘ across plant generators
`has_steam_turbine`	Steam turbine (ST)	0/1	EIA-860 Sched 3 GENERATOR	1 if any generator has ‘prime_mover_code == ’ST’‘
`has_gas_turbine`	Combustion/gas turbine (GT)	0/1	EIA-860 Sched 3 GENERATOR	1 if any generator has ‘prime_mover_code == ’GT’‘
`coal_heat_share`	Share of heat input from coal	fraction 0–1	EIA-923 Sched 5	Sum coal-type heat input ÷ total heat input; 0 for plants burning no coal

2.6.4 Regulatory controls

State climate and efficiency policies enter the primary specification as a single composite covariate, `state_climate_policy_count`, a 0–7 ordinal index summing the seven individual binary indica-

tors. The composite is used in place of the seven separate binaries because none of the individual policies directly targets NO_x (NO_x-specific instruments such as the NO_x Budget Program, CAIR, and CSAPR are not in this bundle); collectively the seven binaries index *state climate-policy posture*, which affects NO_x only indirectly through fuel-switching and dispatch effects. The individual policy binaries remain in the analytical dataset for descriptive transparency.

Table 2.13: Codebook: regulatory controls.

Variable	Label	Units	Source	Operationalization
`state_climate_policy_count`	Composite state climate policy count (primary)	0–7 ordinal	State policy panel (constructed)	Sum of the seven binaries below (excluding ‘rps_target_pct‘); primary-spec covariate per §1.1 of paper1_analysis.qmd
`emissions_cap_active`	State emissions cap in effect	0/1	State policy panel	Direct field; component of ‘state_climate_policy_count‘
`ghg_target_active`	GHG reduction target in effect	0/1	State policy panel	Direct field; component of ‘state_climate_policy_count‘
`rps_active`	Renewable portfolio standard in effect	0/1	State policy panel	Direct field; component of ‘state_climate_policy_count‘
`rps_target_pct`	RPS target (% of retail sales)	%	State policy panel	Direct field; descriptive only — excluded from primary spec per §1.1 of paper1_analysis.qmd
`has_carbon_price`	Carbon pricing market active (RGGI or CA C&T)	0/1	State policy panel	‘ecp_usd_per_tco2 > 0‘; component of ‘state_climate_policy_count‘
`public_benefit_fund_active`	Public benefit fund in effect	0/1	State policy panel	Direct field; component of ‘state_climate_policy_count‘
`electric_decoupling_active`	Electric decoupling in effect	0/1	State policy panel	Direct field; component of ‘state_climate_policy_count‘
`eers_active`	Energy efficiency resource standard in effect	0/1	State policy panel	Direct field; component of ‘state_climate_policy_count‘

2.6.5 Political-economic context

State partisan control enters through a single composite, `dem_lean_index`; the raw partisan measures are retained for descriptives and a disaggregated-partisan robustness spec.[19] ISO/RTO market participation enters through `in_iso_rto` and the categorical `iso_rto`.

Table 2.14: Codebook: political-economic context.

Variable	Label	Units	Source	Operationalization
`dem_lean_index`	Democratic-lean index (composite; primary partisan control)	z-score composite	constructed	Standardized mean of dem_governor, hou_dem_share, and sen_dem_share; the partisan control used in place of the five raw measures. Spans Republican-control to Democratic-control monotonically (divided government in the middle), so trifecta_dem and divided_govt are dropped from the spec. The five raw partisan vars are retained for descriptives and the disaggregated-partisan robustness spec.
`in_iso_rto`	In ISO/RTO territory	0/1	eGRID PLNT	1 if ‘ISORTO‘ is non-null and non-empty; join on ‘orispl‘
`iso_rto`	ISO/RTO name (categorical)	factor	eGRID PLNT	Raw ‘ISORTO‘ field; expanded to dummies (‘iso_miso‘, ‘iso_pjm‘, ‘iso_ercot‘, ‘iso_caiso‘, ‘iso_isone‘, ‘iso_nyiso‘, ‘iso_spp‘) with non-ISO/RTO as reference

[19] The Democratic-lean index (`dem_lean_index`) is the standardized mean of `dem_governor`, `hou_dem_share`, and `sen_dem_share`; it spans Republican to Democratic control monotonically, with divided government in the middle. It replaces the five raw partisan measures (`trifecta_dem`, `divided_govt`, `dem_governor`, `hou_dem_share`, `sen_dem_share`) in the primary spec; all five are retained in the analytical dataset for descriptives and the disaggregated-partisan robustness spec.

2.6.6 NO_x Control — Functional Mediator (Primary)

The primary mediator is a pair of fuel-agnostic functional binaries (`has_combustion_control` and `has_postcombustion_control`), each set to 1 when the plant operates a control of that function in either EIA-860 boiler / equipment records or EPA CAMD facility attributes (the CAMD union recovers the simple-cycle gas-turbine controls the boiler-keyed EIA flags miss). Combustion modification (low-NO_x burners, overfire air, lean combustion, fuel reburning, flue gas recirculation; dry-low-NO_x combustors and water/steam injection on turbines) suppresses NO_x at the source; post-combustion treatment (SCR, SNCR) removes it from the flue gas after it forms. The two are estimated as separate adoption logits (Equations 1a and 1b). The remaining rows are the EIA-860 and EPA CAMD component flags that the two binaries union.

Table 2.15: Codebook: NO_x-control functional mediators.

Variable	Label	Units	Source	Operationalization
`has_combustion_control`	Combustion-modification NOx control present (primary mediator A)	0/1	EIA-860 boiler strategy + EPA CAMD	1 if any in-furnace control present: LNB, OFA, LEA, FGR, steam/water injection, fuel reburn, staged combustion, DLN, cell burner. OR-merge of has_combustion_nox_control (EIA) and camd_has_combustion (CAMD). APTI 418 Ch. 6–8.
`has_postcombustion_control`	Post-combustion NOx control present (primary mediator B)	0/1	EIA-860 ECE + EPA CAMD	1 if SCR or SNCR present (capital-intensive flue-gas treatment). OR-merge of has_scr|has_sncr (EIA) and camd_has_postcombustion (CAMD). APTI 418 Ch. 11–12.
`camd_has_combustion`	CAMD combustion-mod flag (merge input)	0/1	EPA CAMD Facility Attributes	Plant-year max over CEMS units (~>=25 MW) of DLN, water/steam injection, LNB, OFA, fuel reburn, combustion mod, cell burner. See R/camd/01_parse_camd_controls.R.

camd_has_postcombustion	CAMD post-combustion flag (merge input)	0/1	EPA CAMD Facility Attributes	Plant-year max over CEMS units of SCR or SNCR. Standalone 'Ammonia Injection' and 'Other' excluded. See R/camd/01_parse_camd_controls.R.
has_scr	Selective catalytic reduction present	0/1	EIA-860 emissions control equipment	equipment_type = 'SR'
has_sncr	Selective non-catalytic reduction present	0/1	EIA-860 emissions control equipment	equipment_type = 'SN'
has_combustion_nox_control	Combustion modification present (EIA component)	0/1	EIA-860 boiler strategy records	Codes LN, OV, LA, FR, CF, AA, ST/STM, H2/H2O, FU, SC, BF, BO (pre-2013: low_nox_process1/2/3; 2013+: nitrogen_oxide_control_existing_strategy_1/2/3)
has_any_nox_control	Any NOx control installed (either function)	0/1	EIA-860 + EPA CAMD	1 if has_combustion_control == 1 OR has_postcombustion_control == 1 (functional, CAMD-aware). The peaker any-control rate jumps off 4.6% here once CAMD turbine controls merge in.
n_nox_controls	Count of installed NOx control units	count	EIA-860 emissions control equipment	Count of control records per plant-year

2.7 Appendix — Coal-boiler deep-dive

The within-adopter questions the fuel-agnostic primary sets aside, namely *which* post-combustion technology a plant operates (SCR vs SNCR) and *how intensively* it runs it, are boiler-specific and defined only on the coal/boiler adopter subsample (post-combustion adopters with reported utilization

hours, n = 1,604 plant-years). Conditioning on reported hours excludes the post-combustion adopters whose in-service hours are unreported (2.8% of the adopter pool; EIA-923 Schedule 8C leaves the field blank for some control units), so these within-adopter estimates describe the reporting subsample rather than all adopters. Table 2.16 reports the organizational-structure coefficients for the three within-adopter fits; they corroborate the primary's mechanism story without entering the fuel-agnostic specification.

Table 2.16: Organizational-structure coefficients in the coal-boiler within-adopter deep-dive (post-combustion adopter subsample).

	SCR vs SNCR (logit)	Hours (OLS)	ln NOx (OLS)
Government	0.864 (0.550)	341.844 (209.430)	-0.196 (0.127)
Regulated IOU	0.731 (0.453)	243.496† (147.398)	0.069 (0.088)
Cooperative	0.024 (1.058)	142.524 (225.838)	0.109 (0.106)
Host-serving	0.530 (0.520)	854.860*** (234.503)	0.278 (0.216)
Number of owners	-0.091 (0.076)	-2.867 (34.469)	0.031† (0.016)
Third-party ownership	-0.015 (0.391)	187.215 (141.832)	-0.002 (0.082)
FERC cogeneration	-0.631 (0.455)	200.286 (215.156)	0.190 (0.131)
FERC small power	-2.893*** (0.834)	974.482** (318.086)	0.815*** (0.236)
State market-share index	-0.275* (0.133)	-21.484 (63.905)	0.089** (0.031)
Firm scale index	-0.327 (0.227)	46.027 (68.105)	-0.078 (0.049)

Note. θ: SCR-versus-SNCR choice (logit log-odds); Hours: average annual control hours (OLS); ln NOx: within-adopter emission intensity (OLS). CR1 standard errors clustered on utility in parentheses. Post-combustion adopter (Tier 2–3) subsample, n = 1,604. Plant, fuel, state-policy, and partisan controls and year and region fixed effects are included but omitted. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

2.8 Appendix — Full regression tables

Table 2.17 and Table 2.18 report every coefficient in the four primary equations, with model-fit statistics, for readers who want the complete specification behind the curated panels in the Results. Standard errors are CR1, clustered on utility; year and region fixed effects are included throughout (each summarized as a single line rather than printed).

Table 2.17: Full adoption-equation estimates (Equations 1a and 1b), logistic regression in log-odds.

	(1a) Combustion	(1b) Post-combustion
Organizational Structure		
Government	-0.443* (0.216)	0.535* (0.229)
Regulated IOU	-0.102 (0.242)	-0.097 (0.247)
Cooperative	-0.033 (0.347)	-0.500 (0.310)
Host-serving	-0.181 (0.207)	-0.445† (0.237)
Number of owners	0.114 (0.092)	-0.091 (0.061)
Third-party ownership	0.273 (0.213)	0.080 (0.170)
FERC cogeneration	-0.392* (0.176)	-0.609** (0.187)
FERC small power	-0.312 (0.349)	-0.324 (0.368)
State market-share index	-0.141 (0.129)	-0.024 (0.120)
Firm scale index	-0.076 (0.107)	-0.026 (0.094)
Capacity to emit		
log(Nameplate capacity, MW)	0.874*** (0.084)	0.702*** (0.084)
Capacity factor	-0.700* (0.289)	1.356*** (0.271)
Avg generator age (yrs)	-0.033*** (0.006)	-0.045*** (0.006)
Steam turbine	0.694*** (0.203)	0.696*** (0.200)
Gas turbine	-0.254† (0.141)	-0.907*** (0.133)
Coal heat share	1.378*** (0.289)	0.103 (0.217)
Regulatory controls		
State climate-policy count	-0.106† (0.056)	0.056 (0.061)
Political-economic context		
ISO/RTO member	-0.137 (0.165)	0.437* (0.177)
Democratic-lean index	0.208* (0.096)	0.290** (0.096)
Constant		
Constant	-2.259*** (0.622)	-3.781*** (0.626)
Model fit		
Year fixed effects	Yes	Yes
Region fixed effects	Yes	Yes

N	4,817	4,817
McFadden pseudo-R^2	0.252	0.253

Note. Logistic regression; coefficients are log-odds. CR1 standard errors clustered on utility in parentheses. Estimated on the full fossil fleet. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

Table 2.18: Full outcome-equation estimates (Equations 3a and 3b), OLS on log NO_x intensity; Equation 3b adds the two adopted-control indicators.

	(3a) Total	(3b) Conditional
Organizational Structure		
Government	-0.136 (0.098)	-0.091 (0.091)
Regulated IOU	0.262** (0.088)	0.241** (0.076)
Cooperative	0.318*** (0.094)	0.247*** (0.075)
Host-serving	-0.084 (0.112)	-0.174 (0.111)
Number of owners	0.063** (0.022)	0.051** (0.019)
Third-party ownership	-0.054 (0.081)	-0.022 (0.081)
FERC cogeneration	0.227** (0.076)	0.108 (0.075)
FERC small power	0.444 (0.293)	0.395 (0.309)
State market-share index	0.055† (0.030)	0.042 (0.026)
Firm scale index	-0.062 (0.045)	-0.066 (0.044)
Adopted Controls		
Combustion control adopted		-0.472*** (0.081)
Post-combustion control adopted		-0.818*** (0.064)
Capacity to emit		
log(Nameplate capacity, MW)	-0.417*** (0.035)	-0.264*** (0.034)
Capacity factor	-0.948*** (0.131)	-0.793*** (0.116)
Avg generator age (yrs)	0.030*** (0.002)	0.021*** (0.002)
Steam turbine	0.004 (0.094)	0.145 (0.102)
Gas turbine	0.501*** (0.058)	0.342*** (0.055)
Coal heat share	1.402*** (0.081)	1.475*** (0.078)
Regulatory controls		
State climate-policy count	-0.058** (0.021)	-0.056** (0.018)

Political-economic context		
ISO/RTO member	-0.215** (0.069)	-0.179** (0.069)
Democratic-lean index	0.014 (0.035)	0.068* (0.033)
Constant		
Constant	1.039*** (0.293)	1.040*** (0.280)
Model fit		
Year fixed effects	Yes	Yes
Region fixed effects	Yes	Yes
N	4,792	4,792
R^2	0.454	0.522
Adjusted R^2	0.451	0.518

Note. OLS; dependent variable is log NOx intensity (lb/MWh). CR1 standard errors clustered on utility in parentheses. Equation 3b adds the two adopted-control indicators, which are post-treatment. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

2.8.1 Distributional diagnostics

Table 2.19, Table 2.20, and Table 2.21 report the distributional shape (skewness and excess kurtosis) of the continuous variables entering the primary specification, in native and (where transformed) log form. Skewness measures asymmetry (positive values indicate a right tail); excess kurtosis measures tail weight relative to the normal. Both can distort OLS standard errors in small-to-moderate samples, and log transformation is the standard remedy. Conventional rule-of-thumb cut-offs for OLS robustness are |skew| < 2 and |excess kurtosis| < 7 (Curran and West 1996; H.-Y. Kim 2013).

Table 2.19: Distribution diagnostics: dependent variables (native and log forms).

Variable	N	Mean	SD	Median	Skewness	Excess kurt.
`nox_rate_lb_per_mwh`	4,793	1.41	1.64	0.87	2.00	4.62
`ln_nox_rate_lb_per_mwh`	4,793	-0.44	1.48	-0.13	-0.78	0.95
`nox_short_tons`	4,793	849.76	2.24K	117.05	5.68	51.95
`ln_nox_short_tons`	4,793	4.76	2.14	4.77	0.01	-0.40

Note. Log forms cut DV skewness substantially while leaving excess kurtosis moderate; both forms are acceptable inputs to OLS.

Table 2.20: Distribution diagnostics: transformed and composite predictors in the primary specification.

Variable	N	Mean	SD	Median	Skewness	Excess kurt.
log_nameplate_cap_mw	4,818	5.56	1.27	5.67	-0.03	-1.14
firm_scale_index	4,818	0.00	0.97	-0.14	0.32	-1.14
state_market_share_...	4,818	0.00	0.97	-0.36	4.46	33.78

Note. Plant scale enters in log form (the native right tail is well compressed by log1p); the two composites are standardized. The transform justification and the component-collinearity rationale are recorded in `paper1_analysis.qmd`.

Table 2.21: Distribution diagnostics: proportions, counts, rates, and ages in the primary specification.

Variable	N	Mean	SD	Median	Skewness	Excess kurt.
n_owners	4,818	1.24	0.99	1.00	7.42	74.65
max_ownership_share	4,818	0.97	0.12	1.00	-3.57	11.80
coal_heat_share	4,818	0.21	0.39	0.00	1.45	0.16
capacity_factor	4,818	0.32	0.27	0.28	0.39	-1.15
avg_generator_age_yrs	4,818	28.07	16.52	25.33	0.47	-0.53
state_climate_polic...	4,818	2.87	2.01	3.00	0.35	-1.00
dem_lean_index	4,817	0.00	0.82	0.08	0.09	-0.83

Note. Bounded proportions, counts, and rates carry modest skew, so no transformation is applied.

2.8.2 Collinearity diagnostics

Table 2.22 reports the variance inflation factor for every predictor in the conditional outcome model (Equation 3b), the full set behind the four values quoted in the Control Variables section.

Table 2.22: Variance inflation factors, conditional outcome model (Equation 3b), all predictors (year and region fixed effects omitted); conventional flags are VIF > 5 for inspection, VIF > 10 as problematic.

Variable	VIF

State climate-policy count	5.8
org_structure_type	5.4
Democratic-lean index	3.0
log(Nameplate capacity, MW)	2.8
Steam turbine	2.6
Firm scale index	2.5
ISO/RTO member	2.2
Avg generator age (yrs)	2.1
Coal heat share	2.0
State market-share index	1.9
Capacity factor	1.9
FERC cogeneration	1.7
Gas turbine	1.4
Post-combustion control adopted	1.4
Combustion control adopted	1.3
FERC small power	1.3
Number of owners	1.1
Third-party ownership	1.1

Note. VIF computed on the Equation 3b predictor matrix; year and region fixed effects excluded. Largest VIF = 5.8.

Table 2.23 reports pairwise Pearson correlations among the continuous predictors and ownership-concentration measures used in the analysis (log forms where available), complementing the VIF table above. Pairwise complete observations. Correlations $|r| \geq 0.50$ are bold; $|r| \geq 0.80$ additionally carry a ‡ multicollinearity flag.

Table 2.23: Pearson correlations among continuous predictors.

	1	2	3	4	5	6	7	8	9	10
(1) log_cap_mw	—	0.04	0.29	-0.03	0.16	-0.13	0.49	**0.53**	-0.24	-0.21
(2) cap_factor	0.04	—	0.29	-0.02	0.13	-0.17	-0.16	0.24	-0.13	-0.09
(3) coal_sh	0.29	0.29	—	0.32	0.22	-0.22	0.18	0.35	-0.16	-0.14
(4) gen_age	-0.03	-0.02	0.32	—	-0.00	0.02	0.04	0.02	-0.06	-0.09
(5) n_own	0.16	0.13	0.22	-0.00	—	**-0.67**	0.12	0.24	-0.03	-0.04
(6) max_own	-0.13	-0.17	-0.22	0.02	**-0.67**	—	-0.11	-0.20	0.04	0.04
(7) firm_idx	0.49	-0.16	0.18	0.04	0.12	-0.11	—	0.27	-0.21	-0.16

(8) mkt_idx	**0.53**	0.24	0.35	0.02	0.24	-0.20	0.27	—	-0.09	-0.06
(9) clim_cnt	-0.24	-0.13	-0.16	-0.06	-0.03	0.04	-0.21	-0.09	—	**0.76**
(10) dem_lean	-0.21	-0.09	-0.14	-0.09	-0.04	0.04	-0.16	-0.06	**0.76**	—

Note. Pearson correlations, pairwise complete observations. **Bold** $|r| \geq 0.50$; ‡ $|r| \geq 0.80$ (potential multicollinearity flag). Variable names abbreviated; see the Variable Inventory for full names.

Computational environment

The analysis pipeline was implemented in R, spanning data construction, estimation, diagnostics, and rendering. The packages used are listed below; full citations appear in the References Cited.

Table 2.24: R packages used in the analysis pipeline (versions as rendered).

Package	Version	Purpose
Estimation		
sandwich	3.1.1	Cluster-robust (CR1) variance estimators
lmtest	0.9.40	Coefficient testing (coeftest)
car	3.1.5	Variance inflation factors (collinearity)
Data		
arrow	23.0.1.2	Read and write Parquet data files
here	1.0.2	Project-relative file paths
duckdb	1.5.1	In-process SQL / columnar queries
DBI	1.2.3	Database interface for DuckDB
glue	1.8.0	String interpolation
dplyr	1.2.1	Data manipulation
tidyr	1.3.2	Data tidying and reshaping
tibble	3.3.1	Modern data frames
Tables & figures		
ggplot2	4.0.2	Statistical graphics
gt	1.3.0	Formatted tables
kableExtra	1.4.0	Supplementary table formatting
knitr	1.51	Dynamic report generation
moments	0.14.1	Skewness and kurtosis (summary statistics)

3 The Social Incidence of Pollution

Community Demographics, Environmental Control Technology Adoption, and NOx Emissions

Abstract

Communities living beside the United States' fossil-fueled power plants are, on average, disproportionately poor and nonwhite — among the more thoroughly documented regularities in the study of environmental inequality. The best-documented disparities concern where plants are *sited*; this chapter asks instead how they are *run*: among the plants disadvantaged communities already host, how cleanly do those plants operate? Pairing EPA continuous-emissions data with EIA plant records for the coal- and gas-fired fleet ($\geq$ 25 MW) across three snapshots (2009, 2018, and 2021), it tests whether community demographics predict NO_x-control adoption and, in turn, NO_x emission intensity, entering race and class separately so the racial gradient is read net of class. The results show no demographic gradient for emission controls adoption: plants in more-nonwhite communities are no less likely to install controls, on either the inexpensive combustion margin or the capital-intensive post-combustion one. The disparity appears instead in operation. By 2021, though not in 2009 or 2018, plants in more-nonwhite communities run measurably dirtier per megawatt-hour, conditional on the emissions control technologies they carry. The gap is one of emission *rate*, not aggregate mass. As the coal fleet retired, the dirtiest-per-megawatt-hour unit relocated from the large rural coal plant to the small, aging urban peaker, concentrating in more-nonwhite metropolitan load pockets — zones where local demand outstrips import capacity, so small units must run for reliability regardless of cost. The chapter thus identifies an emerging, operational face of environmental injustice in the

power sector — distinct from the siting disparities long studied, and deepening as demand growth and a warming climate keep these high-rate units online.

3.1 Introduction

The communities situated nearest to the United States' fossil-fueled power plants are, on average, disproportionately poor and disproportionately nonwhite, a regularity that is among the more thoroughly documented in the literature on environmental inequality (Mohai and Bryant 1992; Mohai and Saha 2015). Adjacent to this regularity is the question this chapter takes up, concerning not the siting of plants but their operation: among the plants that disadvantaged communities find themselves living beside, how cleanly do those plants actually run? A power plant, after all, does not impose a fixed quantity of harm upon its surroundings. The pollution it emits into the surrounding air is, to a considerable degree, a function of choices that its operator makes and revisits over time, as regulatory standards tighten and as equipment ages. Those choices are of two kinds: 1) whether to install control equipment; and 2) how cleanly the plant is subsequently run with whatever equipment it carries. What this chapter asks is whether those two choices, no less than the siting decisions that originally placed the plants, are patterned by the demographic composition of the surrounding community.

That question draws its analytical interest from a technical circumstance: the installation choice is itself not one decision but two, distinguished principally by their cost. Combustion modification, which operates by preventing NO_x from forming in the first place, is relatively inexpensive; post-combustion treatment (and selective catalytic reduction in particular), which removes NO_x after it has formed, is capital-intensive (the most expensive control option available in the fleet) and also the most effective (Richards 2000). Chapter 2 of this dissertation found that organizational structure leaves its mark less on whether a plant adopts control than on how cleanly the plant is run once equipped; the cost asymmetry nonetheless matters, because it is on the expensive margin that an operator's latitude to under-invest is conceivably at its widest. The distributional question posed here addresses that asymmetry directly. The issue is not whether disadvantaged communities tend to neighbor plants that have installed no controls whatsoever, but rather whether they tend to neighbor

plants that have adopted the inexpensive controls while declining to adopt the costly ones that do most to improve the air those communities breathe. Adoption is, however, only the first of the two ways an operator's conduct can vary: a plant that carries the costly control may still run dirtier than its equipment allows, and the design that follows tells the two apart.

Posing the problem in these terms has the effect of relocating the environmental justice question from the matter of where a hazard is sited to the matter of how the operator of that hazard subsequently behaves. The scholarly tradition that documents environmental inequality has built its evidentiary core upon siting and proximity — that is, upon the question of who lives in proximity to a hazardous-waste facility, a toxic-release source, or a power plant (Chakraborty, Maantay, and Brender 2011; Mohai and Saha 2015). A smaller and more recent body of work has begun, however, to extend the disparity beyond proximity to the question of emissions intensity itself: the racial gradient in ambient nitrogen dioxide has been shown to be both national in scope and considerable in magnitude (Clark, Millet, and Marshall 2014), and the urban neighborhoods that mid-twentieth-century federal mortgage policy designated as hazardous have been shown to be the same neighborhoods in which fossil-fueled plants were subsequently sited, and whose residents continue, in the present day, to bear those plants' nitrogen-oxide, sulfur-dioxide, and particulate emissions (Cushing et al. 2023). Most directly, Hernandez-Cortes, Meng, and Weber (2023) have decomposed the national trend in these disparities, finding that electricity's contribution to racial $PM_{2.5}$ gaps *converged* sharply between 2000 and 2018, as the coal plants that produced the worst of them retired. That accounting, however, measures modeled ambient exposure rather than the emission rate, and closes in 2018 — leaving open whether the fleet's continuing recomposition has, in the years since, produced a fresh disparity in how cleanly the surviving plants are run. This chapter takes up that question, shifting the emphasis from where a plant is sited to how it is subsequently run.

This shift draws on the same power-resources logic that the environmental justice tradition has long applied to the siting decision. Where a host community is able to mobilize comparatively little in the way of economic, political, or organizational resources, the cost to a firm of deferring or forgoing expensive control is correspondingly low, such that firms (and, of equal importance, the agencies charged with regulating them) may be expected to follow what the tradition terms the path of least resistance (Mohai and Bryant 1992; Bullard 1990). Because that least-resistance

logic follows from the relative cost of community resistance rather than from anything peculiar to the siting decision as such, there is little reason to suppose it ceases to operate once a plant has been built. On the contrary, it should continue to bear upon the subsequent and recurring decisions an operator makes regarding the abatement it is willing to fund and (once funded) operate. The remainder of this chapter develops that argument at length, enters race and class separately so that the racial gradient may be assessed net of class as the environmental justice tradition's central claim requires, and tests the argument on the operating fleet. Spatial dependence, the inferential hazard a spatially clustered predictor raises, is addressed in the main outcome equations with Conley spatial heteroskedasticity- and autocorrelation-consistent (spatial-HAC) standard errors; Section 3.4.11 sets out the diagnostics behind that choice.

3.1.1 Research Questions

- **RQ1.** Does neighboring-community demographic composition predict plant-level NO_x intensity, net of organizational structure?
- **RQ2.** Does community composition predict which NO_x controls a plant adopts, and does any gradient concentrate on the capital-intensive post-combustion margin rather than the inexpensive combustion one?
- **RQ3.** Does the demographic association with intensity run through adoption, or does it persist among plants carrying the same controls?

3.2 Theory and Conceptual Framework

This section locates the study in the environmental justice tradition and its finding that racial disparity is not reducible to class, carries the tradition's path-of-least-resistance mechanism from the siting decision to the abatement decision, and concludes with a discussion of where the theory says a demographic gradient should concentrate: on the capital-intensive margin of control.

3.2.1 The environmental justice tradition: race, class, and distributive disparity

The environmental justice tradition is organized around the proposition that the burdens of environmental hazards — often grave, even life-threatening — fall unequally across the population, primarily along lines of race and class. The tradition's empirical core emerged from the movement-era studies of hazardous-waste siting conducted in the 1980s and early 1990s (Bullard 1990; Mohai and Bryant 1992), and was consolidated, over the two decades that followed, into a large and methodologically heterogeneous body of scholarship (Mohai, Pellow, and Roberts 2009). What lends the tradition its analytical coherence is a distributive conception of environmental harm: a hazard is produced at some particular location, and the costs of its operation come to rest upon the particular communities in its vicinity rather than upon the diffuse public that consumes whatever benefit that operation confers. This asymmetry — between a harm that is geographically concentrated and a benefit that is broadly dispersed — is the distributive counterpart to the disproportionality framework (Freudenburg 2006) from which Chapter 2 proceeded. This chapter therefore asks the distributive question that the previous chapter, having confined its attention to the organizational determinants of emission intensity, set aside.

Of the several empirical generalizations the tradition has produced, the one most consequential for present purposes — and among the most robust — concerns the persistence of racial disparity once class has been taken into account. The recurring finding, across narrative reviews of the early evidence (Mohai and Bryant 1992), a meta-analysis of the quantitative literature that followed (Ringquist 2005), and the first national longitudinal test of the siting question (Mohai and Saha 2015) alike, is that the association between racial composition and environmental burden cannot be reduced to the correlation between race and income: the racial gradient persists net of class, and in a considerable number of studies it proves the stronger of the two associations. So consistently has this result recurred that it now functions less as a contested finding than as a premise upon which subsequent theory is built, and it is treated as such here. Consequently, rather than collapse the racial and the economic attributes of the host community into a single index of disadvantage, as is on occasion done, the analysis enters them as separate quantities, so that the association attaching to racial composition may be assessed net of the associations attaching to income and to poverty. The

two attributes are, moreover, given different theoretical footings. Racial composition's association is treated as the empirical regularity just established above — a disparity that survives net of class study after study. Economic standing's association runs through a narrower, specific channel, developed in Section 3.2.3: the comparatively limited economic, political, and organizational resources of lower-income communities attenuate the pressure they can bring to bear upon operators to abate.

3.2.2 From proximity to abatement: relocating the environmental justice outcome

The proposition that environmental burden is distributed along lines of race and class has, for the greater part of the tradition's history, been examined as a question of proximity: the typical study asks how the demographic composition of the area surrounding a hazard compares with that of areas more distant from it, and so measures, in effect, the populations that find themselves situated near hazards (Mohai, Pellow, and Roberts 2009). The present study inverts the unit of analysis. Rather than ask which populations live near a given class of hazard, it assigns a demographic context to each hazard — to each power plant — and takes as its object of explanation the conduct of that hazard's operator. The quantities it estimates therefore describe the demographic surroundings of emitters, and not, it bears emphasizing, the exposure borne by populations. The distinction is consequential, for the two are not interchangeable, and several of the tradition's recurring disputes hinge upon which of them a given design has in fact measured (Downey 2005; Mohai and Saha 2006). This choice follows from the question the chapter opened with: because the object of interest is the abatement decision, the plant that makes it is the natural unit of analysis.

There is already reason to look past siting: the disparity the tradition documents is observable, as well, in what plants subsequently emit. The over-exposure of people of color is a feature of very nearly every major emitting sector, electricity generation among them (Tessum et al. 2021), and the racial gradient in ambient nitrogen dioxide (the pollutant nearest to this chapter's own outcome) is national in scope (Clark, Millet, and Marshall 2014). For fossil-fueled power plants in particular, the disparity reaches the emissions themselves: the neighborhoods designated hazardous by mid-century mortgage policy both drew such plants and continue to bear their nitrogen-oxide, sulfur-dioxide, and particulate emissions (Cushing et al. 2023); proximity to power-plant emissions

falls disproportionately upon poorer and more nonwhite communities (Declet-Barreto and Rosenberg 2022); and the retirement of the oldest and dirtiest units, even as it lowers emissions in the aggregate, has been found to redistribute the residual burden toward those same communities (Richmond-Bryant et al. 2020). That the co-pollutant intensity of combustion varies considerably from one source to another (Boyce and Pastor 2013), and that the gas-fired peaking units of growing importance to the fleet are themselves disproportionately sited near poor and minority populations (Krieger, Casey, and Shonkoff 2016), extends the concern from the coal fleet to the gas fleet with which this study is equally occupied.

This body of work establishes that the disparity reaches the emissions a plant produces. It does not yet establish whether the disparity reaches the control-technology decision from which those emissions issue. That question is the present study's investigatory task. The distinction is more than one of emphasis. Emissions are a jointly determined outcome, with fuel, vintage, utilization, and installed controls all entering, so a disparity in emissions could originate anywhere in that joint determination. This analysis instead isolates one mechanism out of that outcome for specific examination: the control-adoption decision, the deliberate investment through which an operator does or does not abate. It is this decision with which the mechanism developed in the following section is concerned, and it is one that the emissions-focused literature, for all that it is suggestive, has not yet directly examined to the knowledge of the author. The tradition's theory — the distributive frame, the race-net-of-class regularity, and the power-resources mechanism set out below — transfers to the present problem; its findings, drawn from siting and emissions, do not. How demographic context is assigned to each plant is developed in Section 3.4.5.1.

3.2.3 Path of least resistance: from siting to abatement investment

Environmental justice frameworks also offer an account of how the unequal distribution of environmental burden comes about, one centered on the distribution of political and economic power. Applied to the siting of facilities, the environmental justice account treats the process as more political than its formal, technocratic criteria would suggest — a contest among three sets of actors with markedly unequal resources: 1) the developers, who seek to minimize cost and to secure some measure of social acceptance for the facility; 2) the state agencies, which must weigh a conception

of the public interest against the opposition a siting provokes[1]; and 3) the host communities, which will bear the facility's burden and differ very widely in their capacity to resist it (Hamilton 1979). The community's resistance capacity is taken to govern the outcome of that contest: the economic resources, the political access, the social capital, and the organizational capacity that a community is able to mobilize against an unwanted facility (Walsh 1981; McAdam and Boudet 2012). The burden tends to settle where that capacity is low, for the reason that developers and the agencies that regulate them may be expected, other things being equal, to follow the path along which resistance is least (R. S. Liévanos, Greenberg, and Wishart 2018; Greenberg 2017).

The argument of this chapter is that the same logic, framed by the siting literature as an account of where a facility is placed, applies with little modification to the distinct question of how cleanly a facility, once placed, is operated. The controls installed at a power plant are not fixed irrevocably at the moment of its construction; whether to install capital-intensive post-combustion equipment is, rather, a recurring investment decision that an operator confronts and reconsiders as regulatory standards are tightened, as equipment ages and is retired or refurbished, and as the relative costs of compliance and non-compliance shift. That discretion is not an assumption; it is built into how NO_x compliance has been regulated across the years the analysis samples.[2] From the NO_x Budget Trading Program of the early 2000s through the Clean Air Interstate Rule (2005) and its successor the Cross-State Air Pollution Rule (2015–), compliance has run on a cap-and-trade logic rather than a command-and-control one: a source must hold allowances equal to its emissions, and it can meet that requirement either by installing abatement equipment, which shrinks the emissions it must cover, or by purchasing allowances on the permit market. Installing the capital-intensive post-combustion margin is thus only one of several routes to compliance — EPA's own compliance analysis lists operating existing controls, upgrading them, purchasing allowances, and shifting generation as the

[1] 'The state' here is the political-science ideal type theorized in the Marxist tradition, not the U.S. state governments that appear elsewhere in this chapter as concrete regulatory actors. On that account the state is no more a neutral arbiter of 'the public interest' than the firms it regulates: its own growth-oriented agenda intertwines with the political power corporations mobilize to redefine corporate-state relations in their favor (Przeworski 1991). As Prechel and Istvan put it, "when corporations' resource dependence constrains their capacity to accumulate capital, they mobilize politically and pressure the state to redefine corporate-state relations in ways that facilitate their capital accumulation agendas" (Prechel and Istvan 2016, 507). This point is noted here as background only; nothing in the analysis tests a state-growth-imperative mechanism.

[2] Program history and dates per EPA's Clean Air Markets documentation: https://www.epa.gov/power-sector/nox-budget-trading-program and https://19january2021snapshot.epa.gov/airmarkets/nox-budget-trading-program_.html.

options an owner weighs (EPA 2015). Which route an operator takes, moreover, has been shown to track its ownership and regulatory status as much as the emissions cap itself (Fowlie 2010). The pressure that bears upon an operator to undertake the more expensive forms of abatement — the pressure exerted by regulatory enforcement, by challenges to permits, by the threat of litigation, and by the political and reputational costs of being seen to pollute — is itself, to a considerable degree, a function of the resources and the standing of the surrounding community. An operator whose plant adjoins a community poorly equipped to bring such pressure to bear may therefore defer or decline the expensive margin of control at comparatively little cost to itself. There is, in short, little evident reason that the path of least resistance, having shaped the siting of the plant, should cease to shape the subsequent and recurring decisions through which the plant's operator determines how much of its pollution to abate.

The demographic patterning of environmental burden admits of more than one explanation, and the explanations carry different implications, so it matters which of them this argument invokes. The literature commonly distinguishes three: explicit discrimination; economic sorting, in which cheap land, or the property-value consequences of proximity, draws lower-income residents toward the hazard; and the sociopolitical path of least resistance described above — explanations that, as has frequently been emphasized, are not mutually exclusive (Mohai and Saha 2015; McCoy 2017). Of the three, it is the sociopolitical mechanism that the present study engages, for it is the one that operates upon the abatement decision itself, rather than upon the *which came first* problem of how a plant's demographic surroundings came to assume the form they now exhibit. That problem — whether the disadvantaged community was present before the plant or assembled around it afterward — is one that a cross-sectional design of the kind employed here cannot adjudicate, and the formal consequences of that limitation are reserved for Section 3.4.10.

This design cannot resolve how a plant's demographic surroundings formed. The evidence nonetheless weighs against the most benign explanation, that residents simply sorted toward cheap housing near an existing plant. The candidate mechanisms are well catalogued (Banzhaf, Ma, and Timmins 2019). Subnational studies of hazardous-facility siting find disparate siting to outweigh any subsequent in-migration (Pastor, Sadd, and Hipp 2001). The historical record links discriminatory mortgage policy to both the siting of power plants and their present-day emissions (Cushing et

al. 2023). And although a much-cited national study of waste-facility siting found, on a unit-based method, little evidence of disparate siting at all (Been and Gupta 1997), that finding proved controversial and was subsequently refuted: the null has been attributed in large part to the method itself, distance-based designs recovering the disparity that unit-based comparisons had obscured (Mohai and Saha 2007, 2015; Mohai, Pellow, and Roberts 2009). Where sorting around an existing plant does occur, it is constrained by discriminatory housing markets in a manner that leaves even comparatively affluent minority households more exposed than poorer white ones (Crowder and Downey 2010; Banzhaf and Walsh 2008). And a structural reading would in any case decline to treat the gradient as the product of discrete and deliberate choice, regarding it instead as the sedimented outcome of industrial and residential geographies that developed in tandem, race and class operating as social relations rather than as fixed attributes of place (Pulido, Sidawi, and Vos 1996).

The most recent power-plant evidence finds comparatively little disparate siting by race alongside a pronounced increase in the neighboring nonwhite population in the decades after a plant is built (Clay et al. 2025). The direction of formation is, in short, contested, and these studies are best read as establishing a pattern rather than an intent. What this chapter characterizes, and all it claims to characterize, is the condition that Downey (2005) terms "disparate exposure" — a concentration of pollution among particular groups exceeding what a spatially random distribution would yield, however that concentration came historically to arise.

The closest precedent for this argument is the smaller literature on enforcement and compliance behavior, which looks for demographic patterning not in where hazards sit but in how firms and regulators conduct themselves. That literature returns a more equivocal verdict on race than the siting canon. State enforcement of the major environmental statutes runs consistently weaker in poorer jurisdictions, with race a fainter and less consistent predictor than class (Konisky 2009; Konisky and Schario 2010). The undetected non-compliance through which firms evade abatement eases where the surrounding community is politically mobilized (Konisky and Reenock 2013), a result that gives the community-resistance mechanism a firm-and-regulator micro-foundation. Where regulation does bind, it reshapes the racial distribution of pollution: differential Clean Air Act enforcement has been identified as the principal driver of the post-2000 narrowing of the Black–white gap in particulate exposure (Currie, Voorheis, and Walker 2023), even as disparities in the pace of emission

reductions have persisted within the energy sector (Nunez et al. 2024). The implication is direct: if enforcement and the abatement it compels are both patterned by the demographic composition of the affected community, the control-adoption decision at the center of this chapter should be patterned likewise.

3.2.4 Two margins of abatement: bridging Chapter 2

Chapter 2 established that the control of NO_x is, as a technical matter, not a single decision but at least two, and that the two examined here are separated above all by their cost. Combustion modification (low-NO_x burners, overfire air, and the dry-low-NO_x combustor designs) prevents NO_x from forming in the course of combustion. It is relatively inexpensive, often little more than an incremental retrofit or an adjustment to operating practice. Post-combustion treatment (selective catalytic reduction and its non-catalytic counterpart) removes NO_x after it has formed. It is capital-intensive, a discrete and substantial investment, and selective catalytic reduction is the most costly control option the fleet affords.[3] The two margins are, moreover, as unequal in their effectiveness as in their expense: combustion modification abates on the order of a fifth to a quarter of a plant's NO_x, selective non-catalytic reduction something closer to a third or a half, and selective catalytic reduction the large majority of it (Richards 2000). The margin that costs the most is thus also the margin that accomplishes the most. That conjunction of high cost and high effectiveness in a single instrument is what raises the distributional stakes: the control an operator has the most latitude to decline is also the one that would do the most for the surrounding community's air.

If, as the foregoing mechanism supposes, the resistance capacity of the surrounding community conditions the abatement an operator undertakes, then that conditioning ought not to fall evenly across the two margins. The inexpensive margin is installed widely and is difficult to forgo, so comparatively little room remains for a demographic gradient to express itself there. The expensive margin is where an operator's latitude to underinvest is greatest, and is therefore where any gradient ought most plausibly to appear. The prediction that follows is accordingly not about the absence

[3] The asymmetry is roughly an order of magnitude. For an illustrative 500 MW coal unit, EPA's estimates put the capital cost of an SCR retrofit at about $292 per kilowatt and an SNCR retrofit at $46; combustion-control upgrades on the same unit remove NO_x at $430–1,200 per ton, against the SCR's $5,000 (EPA 2015). Utility SCR retrofits of the 2000s ran roughly $270–570 per kilowatt in 2011 dollars, with operating costs near 0.1 cents per kilowatt-hour for catalyst and reagent — a continuing commitment, not a one-time purchase (Sorrels et al. 2019).

of control but about its composition: plants near disadvantaged communities should carry the inexpensive controls while going without the expensive ones — the controls that would have done the most for the surrounding air. The operative question, in other words, concerns not whether a plant abates, but which of the two margins of abatement it elects to adopt.

The relation of this expectation to Chapter 2 is one of contrast as much as inheritance. The previous chapter looked for organizational structure on this same capital-intensive margin and largely did not find it there: government owners aside, adoption proved weakly patterned by ownership, and the organizational gradient surfaced instead downstream, in how cleanly plants run once equipped. That result does not carry the demographic prediction down with it, because the two mechanisms are different in kind. The organizational account turns on the cost of capital confronting each class of owner (Fowlie 2010); the power-resources account developed above turns on the resistance a community can mount (Grant, Jorgenson, and Longhofer 2018), and it is on the costly, discretionary margin that such resistance, or its absence, has the most room to tell. The question is therefore whether community composition patterns adoption where organization did not, and whether it also reaches operation where organization did, both asked net of the organizational structure the analysis carries forward as a control. The empirical logic follows directly: if the association between community demographics and emission intensity is carried by the abatement decision, it should diminish once adoption is held constant, and should diminish most upon the post-combustion margin; were it instead to persist undiminished, the implication would be that demographics bear upon intensity downstream of adoption, through channels of operating and dispatch behavior that conditioning on equipment cannot capture. That conditioning step, together with the directional predictions to which it gives rise, is set out in the hypotheses that follow.

3.3 Hypotheses

The theory developed above yields directional predictions about how the demographic composition of a plant's surroundings is related to its abatement choices and, by way of those choices, to the intensity of its NO_x emissions. The predictions are derived in reference to a single mechanism, namely the path of least resistance, operating most forcefully upon the capital-intensive margin of

control.

Three conventions govern how these predictions should be read, each developed in full below. 1) The predictions are associational rather than causal: the cross-sectional design cannot separate a plant's siting beside a disadvantaged community from the sorting of residents around it (Section 3.4.10). 2) Race and class enter separately, the racial association read net of class by nested entry, in keeping with the tradition's claim that racial disparity survives the introduction of class (Ringquist 2005; Mohai and Saha 2015) (Section 3.4.5.1). 3) The predicted signs concern the *level* of each association, not its trajectory: the demographic-by-year interactions test the trajectory without a directional prior, since the regulatory tightening of these years could as readily compress the gradient as steepen it (Section 3.4.8). Table 3.1 assembles the predicted signs; the hypotheses that follow supply the reasoning behind them.

Table 3.1: Predicted signs by demographic measure across the two adoption margins (Eq 1a, 1b), their differential, and the total and conditional intensity equations (Eq 3a, 3b); signs were fixed before estimation.

Measure	(1a) Combustion	(1b) Post-comb.	Differential	(3a) Total	(3b) Conditional
% nonwhite (race)	− / 0	−	−	+	+ atten.
% poverty (class)	−	−	−	+	+ atten.
Median HH income (class)	+	+	+	−	− atten.

Note. Signs are theoretical predictions on the log-odds of adoption (Eq 1a, 1b) and on log NOx intensity (Eq 3a, 3b). The differential is the post-combustion minus the combustion prediction: a negative value is the path-of-least-resistance signature for a disadvantage measure — the expensive margin is suppressed more than the cheap one. "atten." marks a coefficient predicted to shrink toward zero from Eq 3a to Eq 3b as adoption is held constant, the attenuation that would locate the mechanism in abatement. In that column the sign is the coefficient's predicted direction, carried over from the total column, and "atten." says only that its magnitude is expected to fall toward zero once adoption is held constant, not that it changes sign — so "+ atten." is a positive coefficient that shrinks, "− atten." a negative one. Higher income is protective, so its signs invert. The race measure is tested net of class by nested entry.

3.3.1 H1 — Adoption

The path of least resistance, as the previous section argued, does not bear evenly upon the two margins of control (Section 3.2.4): combustion modification is adopted widely and is difficult to forgo, while the capital-intensive post-combustion margin leaves an operator the widest latitude to

underinvest, and is therefore where the resistance capacity of the surrounding community ought to register most distinctly. Where that community is more heavily nonwhite or poorer (and so, on the tradition's account, commands a smaller share of the economic, political, and organizational resources by which costly abatement is compelled (Bullard 1990; Banzhaf, Ma, and Timmins 2019; Konisky and Reenock 2013)), the operator may defer or decline the expensive control at comparatively little cost to itself. The expectation is accordingly lower adoption, concentrated upon the post-combustion margin: a coefficient more sharply negative on the post-combustion logit (Eq 1b) than on its combustion counterpart (Eq 1a), and hence a negative differential between the two, net of the organizational structure that Chapter 2 showed to bear upon the same margin. Median income, as a measure of advantage rather than of disadvantage, is expected to run in the opposite direction, its signs inverted throughout.

H1a. Plants whose surrounding communities are more disadvantaged — higher in percent nonwhite and percent poverty, lower in median income — are less likely to adopt NO_x control, net of organizational structure.

H1b. The adoption shortfall of H1a falls hardest on the expensive control: for each disadvantage measure, the negative association is larger for post-combustion adoption (Eq 1b) than for combustion modification (Eq 1a), yielding a negative post-minus-combustion differential, net of organizational structure.

3.3.2 H2 — Mechanism

The mechanism this chapter proposes may be put as a question of where, along the chain that runs from community to emission, the demographic gradient does its work. The total gradient of Eq 3a is the headline quantity: whether plants neighboring disadvantaged communities run dirtier at all. It stands on its own, whatever the answer to the mechanism question turns out to be. That question is a sorting question: if such plants do run dirtier, do they run dirtier *because* they adopted less control, or *despite* the control they carry? The conditioning contrast sorts the gradient into exactly those two parts. If the gradient operates, as the theory holds, through adoption, it ought to weaken once the two adoption margins are entered into the intensity equation, in the passage from the unconditional specification (Eq 3a) to the specification that holds adoption constant (Eq 3b). And it ought to

weaken most where the post-combustion margin does the abating. Whatever attenuates is, in other words, an equipment gap: dirtier because less control was adopted. Whatever persists is a gap among plants carrying the same equipment on paper, and so locates the gradient downstream of the adoption decision the data observe, whether in operating and dispatch behavior, in how hard an installed control is actually run, or in influences on intensity the design does not measure. The attenuation is read as suggestive of a pathway rather than estimated as a formal indirect effect, for reasons set out alongside the estimating equations (Section 3.4.9).

H2a. The demographic association with emission intensity runs in part through adoption: it attenuates when control adoption is held constant (Eq 3a → Eq 3b), and whatever portion persists locates the gradient downstream of the adoption decision.

H2b. The attenuation is concentrated upon the post-combustion margin — the association falls most where the capital-intensive control does the abating — mirroring the adoption differential of H1b at the intensity stage.

3.3.3 H3 — Race net of class

With race and class modeled as separate quantities, the environmental justice literature yields a clear expectation: the association attaching to percent nonwhite should survive the introduction of the two measures of class (Ringquist 2005; Mohai, Pellow, and Roberts 2009). That pattern has held, where it has been examined, even for ambient nitrogen dioxide, the pollutant nearest this chapter's own outcome (Clark, Millet, and Marshall 2014). Whether it will survive on this fleet is, however, an open question. While the siting canon finds the racial association robust to controls for class, the nearer literature on enforcement and compliance behavior, reviewed above, finds class the steadier of the two (Section 3.2.3). Should the racial association instead collapse toward zero once the class measures are entered, the implication would be that racial composition had been standing proxy for class all along, its apparent effect an artifact of the residential segregation that makes racial and economic disadvantage coincide.

H3. The percent-nonwhite association with abatement and emission intensity persists net of poverty and income.

3.4 Data & Methods

3.4.1 Data sources

Table 3.2 lists the public datasets the analysis draws on. It carries forward Chapter 2's plant-side sources unchanged and adds the community-demographic and air-quality-regulation data the analysis introduces (EIA 2025; EPA 2026, 2025; EPA eGRID 2009, 2018, 2021; Bureau 2026a, 2026b; US EPA 2016; Gonzalez and Carlotto 2026; NCSL 2026; Grossmann, Jordan, and Mccrain 2021); each variable's source is given in full in the variable inventory below.

Table 3.2: Data sources for the analysis.

Dataset	Provider	Version / coverage	Role in the analysis
Form EIA-860 (Annual Electric Generator Report)	U.S. Energy Information Administration	2009, 2018, 2021	Generators, ownership, and control-equipment records; the plant spine
Form EIA-923 (Power Plant Operations Report)	U.S. Energy Information Administration	2009, 2018, 2021	Fuel and generation records (Page 1): the coal/gas sample screen and the fuel-mix and capacity-factor controls
eGRID (Emissions & Generation Resource Integrated Database)	U.S. Environmental Protection Agency	8th, 13th, 16th editions (2009, 2018, 2021 data)	Plant NOx emission rate (the dependent variable)
Clean Air Markets Program Data (CAMD/CEMS facility attributes)	U.S. EPA Clean Air Markets Program	2009, 2018, 2021	Simple-cycle gas-turbine NOx-control flags EIA does not record
American Community Survey, 5-year estimates	U.S. Census Bureau	2007–2011, 2016–2020, 2019–2023 windows	Community race, ethnicity, and income (the primary predictor)
TIGER/Line shapefiles	U.S. Census Bureau	2010 and 2020 tract vintages	Census-tract geography for population-weighted buffer apportionment

Green Book (Nonattainment Areas for Criteria Pollutants)	U.S. Environmental Protection Agency	8-hour ozone, snapshot years	County ozone-nonattainment status (RACT-driven abatement control)
State climate-policy dashboard	ClimateXChange	2009, 2018, 2021	State climate-policy count, RPS targets, carbon pricing
State partisan composition	NCSL; Correlates of State Policy Project	2008–2024	Governor and legislative party control (political-economic context)

3.4.2 Unit of Analysis

The unit of analysis is the plant-year: a single electricity-generating plant observed in one of three snapshot years — 2009, 2018, and 2021. The estimating sample is the continental-U.S. fossil fleet that burns coal or natural gas, restricted to plants of at least 25 megawatts that match to the U.S. Environmental Protection Agency's continuous-emissions-monitoring records (CAMD/CEMS). The capacity floor and the monitoring match do the same work from two directions: together they ensure that every plant in the sample is a non-trivial emitter whose NO_x is metered at the stack rather than estimated from engineering factors. Pooled across the three snapshots, the sample comprises 4,818 plant-years.

This is deliberately the same fleet Chapter 2 analyzed: the analysis inherits its sample and mediator and carries its organizational-structure measures forward as controls, so the demographic gradient is estimated net of organizational determinants on the same plants where the previous chapter documented them.

3.4.3 Scope of Analysis

The design is cross-sectional in three snapshots rather than a continuous panel. Each snapshot is a full cross-section of the qualifying fleet in a single year, and the three are spaced to bracket distinct regulatory and market regimes rather than to trace year-to-year dynamics. The 2009 snapshot observes the fleet under the Clean Air Interstate Rule, before the Cross-State Air Pollution Rule and the post-combustion retrofit wave it drove; the 2018 snapshot observes it once that rule (effective

2015) had matured and the bulk of that installation was complete; the 2021 snapshot observes it after the pandemic's demand shock and the continued displacement of coal by gas. The three are, in other words, three different policy environments sampled on the same fleet. This regime shift is what makes the movement between the periods legible: the pooled specification and its demographic-by-year interactions, developed in Section 3.4.8, turn that deliberate spacing into a test of whether the demographic gradient itself shifts across regimes.

The geographic scope is the continental United States; Alaska, Hawai'i, and Puerto Rico are excluded.

3.4.4 Dependent Variables

The dependent variable is NO_x emission intensity: pounds of nitrogen oxides emitted per megawatt-hour of net generation (`nox_rate_lb_per_mwh`), drawn from EPA eGRID. Intensity, rather than mass, is the primary outcome because the question is about how cleanly a plant runs, not how much it runs: a rate divides emissions by output and so isolates the abatement behavior of interest from the plant's scale, which a mass figure would confound. The rate enters in logs (`ln_nox_rate_lb_per_mwh`) — it is sharply right-skewed, its determinants act multiplicatively rather than additively, and a logged coefficient reads as a proportional change in intensity, the quantity in which the hypotheses are stated. This is the same outcome, in the same form, as Chapter 2, so the demographic and organizational gradients are estimated on a common scale.

A second, mass-based outcome — total NO_x in short tons, also logged (`nox_short_tons`) — is carried as a robustness check, and for this chapter it is more than a functional-form variation. What a neighboring community actually breathes is a quantity of pollution, not a rate; a clean-running but enormous plant and a dirty-running but small one can impose the same burden. The rate isolates the operator's decision; the mass measures the exposure that decision produces. The primary specification is built on the rate because the chapter's mechanism is the abatement choice, but the mass outcome keeps the distributional stakes in view.

3.4.5 Primary Independent Variable: Community Demographics

The primary independent variable is the demographic composition of the community surrounding each plant: the percent nonwhite, the percent in poverty, and the median household income of the population within a 3 km buffer, built by population-weighted areal apportionment of ACS 5-year tract estimates (Section 3.4.5.1), with the 5 and 10 km buffers retained as radius robustness. In the ACS's own terms, the three are built from tables B03002 (Hispanic or Latino Origin by Race), B17001 (Poverty Status in the Past 12 Months by Sex by Age), and B19013 (Median Household Income in the Past 12 Months): percent nonwhite is the share of the population that is not 'White alone, not Hispanic or Latino'; percent in poverty is the share of the population for whom poverty status is determined whose income falls below the poverty threshold; and median household income is B19013's inflation-adjusted median. The plain names are used hereafter. Race and class are modeled *separately* rather than collapsed into a single disadvantage index, so the nonwhite association can be read net of poverty and income — see Section 3.3.

Each measure takes the functional form its distribution dictates. Percent nonwhite and percent poverty are bounded proportions and are modeled linearly: percent nonwhite spreads across the full range (skewness 0.61), and 28.7% of plants neighbor a majority-nonwhite community — a margin the robustness analysis revisits as a majority-minority threshold; percent poverty is more compressed (skewness 1.2), but its zero floor precludes a log, so it too is linear. Median household income is sharply right-skewed in levels (skewness 1.52) and is modeled in logs, which restores near-symmetry (skewness 0.15) and lets its coefficient read as a semi-elasticity — the same rationale that logs the outcome. Tertile specifications of all three measures, and a majority-minority indicator for percent nonwhite, are retained as robustness checks on the linear form.

3.4.5.1 Demographic Measurement

Translating "the surrounding community" into a variable requires a rule for deciding which people count as surrounding a given plant, and the choice of rule matters. The environmental justice literature has shown repeatedly that assigning a hazard the demographics of its single host tract (the unit-hazard coincidence method) systematically understates disparity, because it ignores the populations just across the tract line who are no less exposed (Chakraborty, Maantay, and Brender

2011; Mohai and Saha 2006). The distance-based alternative takes the exposed community to be everyone within a fixed radius of the plant and weights them by population. The analysis adopts that approach as its primary measure: around each plant it draws circular buffers of three, five, and ten kilometers and computes the population-weighted demographic composition of the people inside (Figure 3.10).

Because census tracts do not respect a buffer's edge — most buffers cut through tracts rather than around them — the population inside must be apportioned. Each intersecting tract contributes its demographic rates in proportion to how much of it falls within the circle, under the working assumption that population is spread uniformly across the tract's area. For a buffer intersecting tracts indexed by i, tract i's weight is its total population scaled by the share of its area inside the buffer,

$$w_i = \text{pop}_i \times \frac{\text{area}_i^{\text{inside}}}{\text{area}_i},$$

and the buffer's value on a demographic rate r is the weighted mean $\bar{r} = \sum_i w_i r_i \,/\, \sum_i w_i$. All geometry is computed in the CONUS Albers equal-area projection (EPSG:5070), so distances and areas are metric and undistorted, and each snapshot is matched to its contemporaneous census geography and five-year American Community Survey window — 2010 tracts and the 2007–2011 ACS for 2009, 2020 tracts and the 2016–2020 and 2019–2023 windows for 2018 and 2021. A geonormalized alternative, tract series standardized to a single boundary vintage, was set aside because the standardized products carry only bare demographic counts, without the income and poverty measures this chapter requires entered alongside race.

Two coarser choices serve as robustness comparisons. The radius is varied across the three buffers, so that no single definition of "nearby" drives the result; and the population-weighted buffer is set against the nearest-centroid tract measure — the demographics of the lone tract whose centroid sits closest to the plant — the cruder linkage the apportionment is meant to improve upon. Table 3.13 sets the two example plants side by side and shows that the radius gradient has no fixed direction: widening the buffer makes the urban plant's surroundings look more nonwhite but less poor, and the rural plant's both more nonwhite and poorer. Figure 3.10 and Figure 3.11 bracket the two regimes the design must straddle: a dense city, where the apportionment of many small tracts dominates

and the radius matters less, and a sparse rural area, where a single tract fills the inner buffer so that apportionment is moot and the radius does all the work. The apportionment itself rests on the uniform-density assumption noted above; a dasymetric refinement, redistributing population to inhabited land within each tract, would sharpen it further and is left as an extension.

These two cases represent the ends of a continuum, which raises the question of how much the choice of radius actually matters. On the whole, reassuringly little. Widening the buffer from three to ten kilometers leaves the typical plant's estimated share nonwhite essentially unchanged — a median shift of under half a percentage point, as likely to fall as to rise — so the radius does not systematically tilt the estimate toward more or less disadvantaged communities; and the three bands rank plants almost identically, the three- and ten-kilometer measures correlating at 0.91. Yet that average concordance conceals real divergence in the tail: for the most affected tenth of plants the three-to-ten-kilometer estimate moves by 18 points or more, and about 11% of plants — nearer 16% among the densest, most urban ones — cross the majority-minority threshold between the two radii as a wider circle sweeps in adjacent, demographically different tracts (cf., Figure 3.11). Theoretically, the two radii are also measurements of two different ideas of the relevant community: three kilometers approximates the residential vicinity of the plant, though it is already far wider than the pedestrian-scale neighborhood of roughly 500 m that the spatial-demography literature describes (Lee et al. 2008; Raoul S. Liévanos 2015); ten kilometers is a broader 'shed' over which the plant's NO_x disperses. That they coincide on average makes the choice defensible either way; that they diverge for dense-city plants is why the analysis fixes the primary radius at three kilometers while retaining all four estimates — the three-, five-, and ten-kilometer buffers together with the nearest-centroid tract — in the analytical dataset, so that every result can be reported against the full set as a robustness check. A further reason favors the narrow radius: because the fleet is physically clustered (Section 3.4.11), wider buffers increasingly overlap those of neighboring plants, blending one plant's measured community into the next and inflating the apparent spatial clustering of the predictor — distortions the three-kilometer ring largely avoids. Across time, finally, the measure is comparatively stable: neighborhood composition shifts only modestly between snapshots for plants present throughout.

3.4.6 Mediator: Functional NO_x-Control Adoption

The mediator is whether the plant operates NO_x control, coded as two functionally distinct margins that differ in their capital intensity — the dimension on which the chapter's mechanism turns. The first, `has_combustion_control`, marks any in-furnace combustion modification — low-NO_x burners, overfire air, flue-gas recirculation, fuel reburning, and the dry-low-NO_x combustors and water or steam injection used on gas turbines — the comparatively inexpensive class of control that suppresses NO_x as it forms. The second, `has_postcombustion_control`, marks selective catalytic or non-catalytic reduction (SCR or SNCR), the capital-intensive flue-gas treatment that removes NO_x after it has formed and over which an operator retains the widest discretion to underinvest (Section 3.2.4). The two are estimated as separate adoption equations (Eq 1a, 1b), and they re-enter the conditional intensity equation (Eq 3b) as the abatement through which the demographic gradient may run.

Both are fuel-agnostic binaries, constructed exactly as in Chapter 2: each is set to one when a control of that function appears in either the EIA-860 equipment and boiler-strategy records or the EPA CAMD facility attributes; both sources are needed because the boiler-keyed EIA flags miss the controls on simple-cycle gas turbines, which the CAMD records catch.[4] Carrying the mediator over unchanged keeps the adoption margins measured identically across the two chapters, so that what Chapter 2 established about organizational structure and what this chapter asks about community demographics bear on the very same two decisions.

3.4.7 Control Variables

The controls are the factors that bear on a plant's emission intensity and would, if omitted, risk masquerading as a demographic effect. They fall in three blocks. The first is the plant's *capacity to emit* — the physical conditions that fix its unabated NO_x potential independently of any abatement

[4] "Boiler-keyed" means that EIA-860 records a plant's NO_x controls against its *boilers* — the chambers in which fuel is burned to raise steam — so a control only registers if it is attached to a boiler unit. A simple-cycle combustion turbine has no boiler: it drives its generator directly and vents the exhaust without a steam cycle, so it carries no boiler record for a control to attach to, even when the turbine is in fact fitted with a dry-low-NO_x combustor or with water or steam injection. The EPA CAMD facility attributes instead record controls at the emitting unit and cover combustion turbines directly. Treating a control in *either* source as evidence of adoption therefore matters most for the gas peaking fleet, much of which is simple-cycle: on the EIA boiler flags alone these units would read as uncontrolled, understating combustion-control adoption across a large part of the gas fleet.

choice: generating scale, utilization, generator age, prime-mover configuration, and fuel-mix composition. These hold the engineering baseline constant, so the demographic coefficient is not left to absorb the brute fact that, say, older coal units in poorer regions burn dirtier before anyone decides whether to control them.

The second is the *regulatory and political-economic context*: a composite count of state climate policies, the partisan lean of state government, participation in an organized wholesale market, and — the one addition this chapter makes to Chapter 2's control set — whether the plant's county was in ozone nonattainment. The last warrants its own justification. Ozone nonattainment triggers Reasonably Available Control Technology requirements on NO_x sources, so a plant in a nonattainment county may abate not in response to its surrounding community but because the Clean Air Act compels it; and because nonattainment areas are disproportionately urban, and so more heavily nonwhite, the designation would otherwise let a federal mandate be misread as a demographic gradient. Conditioning on it separates the regulatory pressure from the community signal (Section 3.4.10).

The third is *organizational structure*, the apparatus that was Chapter 2's primary predictor and is here demoted to a control: ownership and regulatory type, ownership concentration, firm scale, market position, and the cogeneration and small-power certifications. It enters because the kinds of firms that own plants are themselves unevenly distributed across communities and, as Chapter 2 showed, organizational type bears on abatement in its own right; Section 3.4.10 details the pattern. Conditioning on it lets the demographic coefficient be read net of which kind of firm owns where, rather than letting the sorting of firms across places stand in for a community response. Each control is catalogued with its source and construction in the variable inventory below.

3.4.8 Analytical Strategy

The 2009, 2018, and 2021 snapshots are pooled into a single cross-section estimated with year and NERC-region fixed effects and standard errors clustered on the operating utility. The choice to pool follows from where the identifying variation lies. The research questions guiding the analysis turn on differences across plants (plants in different communities abating differently), so the design uses that between-plant variation rather than differencing it away. Employing a plant-fixed-effects

panel would absorb the cross-sectional signal of interest to this chapter, and would do so only to recover within-plant demographic change that is both slow-moving and partly an artifact of shifting census-tract boundaries. Pooling the cross-sections thus preserves the between-plant comparison. It does not, however, assume the demographic gradient is fixed over the decade the snapshots span: the demographic-by-year interactions below let it vary and test whether it moves. Pooling does make a milder assumption: that the control relationships and functional form are stable enough across the three regimes to be estimated once.[5] Year fixed effects absorb the shocks that move the whole fleet at once: gas-price swings, demand shifts, and the secular regulatory tightening across the three regimes. NERC-region fixed effects absorb the regional differences in fleet composition and air-quality regulation that sort plants and communities together. The demographic coefficients then estimate variation in abatement and intensity across communities facing the same year and the same grid region: the comparison the question intends, not an artifact of when or where a plant sits. Clustering on the utility accounts for the correlation among plants run by a common operator.

The mechanism is read from a conditioning contrast rather than a formal mediation decomposition. The outcome equation is estimated twice, once without the two adoption margins and once with them, so that the passage from the total gradient (Eq 3a) to the conditional gradient (Eq 3b) sorts the demographic association into the two parts set out in Section 3.3: the part carried by lesser adoption, which attenuates, and the part that persists among plants with the same equipment. The paired adoption equations locate which margin does the carrying, the cheap or the capital-intensive. The attenuation is read as suggestive of the pathway, not as an identified indirect effect, a restraint Section 3.4.10 makes explicit.

Finally, the demographic measures are interacted with snapshot year. The main demographic coefficients then describe the gradient in 2009, the reference year, and each interaction term measures how the 2018 or 2021 gradient differs from it. Because the design pools regimes that a decade of regulation and fuel-switching separates, these terms ask whether the gradient itself moved: narrowing, if the tightening regulatory floor forced abatement even where communities have the least political weight, or widening, if the decade's abatement investment flowed disproportionately to plants near advantaged communities. And because the interaction is an estimated coefficient with

[5] Estimating on the full pooled sample, rather than three thinner annual cross-sections, is also what gives those interactions enough observations to be informative. On a single snapshot's fleet they would be badly underpowered.

a standard error, the question is posed as a single test (whether the gradient in one year differs from another by more than sampling noise) rather than left to estimating the three snapshots apart and judging by eye.

3.4.9 Estimating Equations

As in Chapter 2, the analysis estimates a system of four equations on the same fleet. D_{it} is the vector of community demographics around plant i in snapshot t, X_{it} the control blocks defined in Section 3.4.2, and y_{it} the plant's log NO_x emission intensity; the race and class components of D_{it} enter separately (Section 3.3). Every equation carries year fixed effects γ_t and NERC-region fixed effects $\delta_{r(i)}$, with standard errors clustered on the operating utility.

Two adoption equations — one per control function — model whether the plant operates that control, as logits:

$$\Pr(A^c_{it} = 1) = \Lambda\Big(D'_{it}\beta^c + X'_{it}\theta^c + \gamma_t + \delta_{r(i)}\Big) \qquad (1a)$$

$$\Pr(A^p_{it} = 1) = \Lambda\Big(D'_{it}\beta^p + X'_{it}\theta^p + \gamma_t + \delta_{r(i)}\Big) \qquad (1b)$$

where A^c_{it} and A^p_{it} are the combustion- and post-combustion-control binaries and $\Lambda(\cdot)$ is the logistic function. The levels of β^c and β^p — whether disadvantaged communities adopt less control on each margin — are the adoption-side test of H1a. Their contrast, $\beta^p - \beta^c$ — whether community demographics bear more heavily on the capital-intensive margin than on the cheap one — tests H1b. Because it carries the chapter's central claim, it is assessed formally rather than by inspection of two separately fitted models: the two adoption equations are stacked, with each plant-year entered once per margin, and the demographic vector is interacted with a post-combustion indicator, so that $\beta^p - \beta^c$ is recovered as a single estimated coefficient per demographic measure, with a standard error clustered on the operating utility. So that this differential is a level quantity rather than a year-specific one, the demographic terms in the stacked specification are entered without the year interaction that the per-margin equations carry.

Two outcome equations model emission intensity, the first unconditional on abatement and the

second conditional on it:

$$y_{it} = D'_{it}\beta^{\text{tot}} + X'_{it}\theta + \gamma_t + \delta_{r(i)} + \varepsilon_{it} \qquad (3a)$$

$$y_{it} = D'_{it}\beta^{\text{cond}} + \lambda^c A^c_{it} + \lambda^p A^p_{it} + X'_{it}\theta + \gamma_t + \delta_{r(i)} + \varepsilon_{it} \qquad (3b)$$

Equation 3a recovers the total demographic gradient in intensity; Equation 3b adds the two adoption indicators and so holds abatement fixed. The movement from β^{tot} to β^{cond} is the conditioning contrast that tests H2a (Section 3.4.8): the portion of the demographic gradient that travels through the abatement decisions the data can observe, read as suggestive rather than as an identified indirect effect. Because Equation 3b conditions on both margins jointly, H2b — whether that attenuation concentrates on the post-combustion margin — is read from where the gradient falls rather than recovered as a single coefficient. In each of the four equations the demographic vector is interacted with snapshot year, so that β is estimated year by year and the trajectory across the three regimes is a tested quantity rather than an assumption.

3.4.10 Identification and Causal Language

The design is cross-sectional and observational, and its coefficients are conditional associations rather than causal effects. Community demographics are not assigned at random: whatever places a plant beside a particular community shapes much else about the plant and its surroundings. The analysis conditions on the organizational, technical, and political-economic factors set out above and reads each demographic coefficient as the difference that survives them: in other words, the association between community composition and abatement net of what the model can hold fixed.

This critique is most pointed in the environmental-justice literature's *which came first* problem (Mohai and Saha 2015; Pastor, Sadd, and Hipp 2001). In brief, the problem is that a single snapshot cannot say whether a disadvantaged community grew up around an already-dirty plant or a dirty plant was placed beside an already-disadvantaged community. The question this chapter asks does not turn on resolving it: the origin of demographic gradients is not adjudicated here. The outcome is the operator's ongoing abatement decision, not the historical siting decision — and whatever

produced the demographic surroundings, the claim is only that operators abate less where those surroundings are more disadvantaged. The siting-versus-sorting ambiguity therefore bounds only what can be said about how the disparity *formed* (Section 3.2.3); it does not bound claims about whether plants neighboring disadvantaged communities run dirtier today, the focus of this chapter.

Two features of the design restrict interpretations of the coefficients further. Claims about change across the period rest on the demographic-by-year interactions (Section 3.4.8), not on narrating three separately estimated cross-sections as though their differences were a trend — a reading the slow drift of neighborhood composition and the decennial revision of tract boundaries make especially treacherous (Section 3.4.5.1). In addition, the kinds of firms that own plants are themselves unevenly distributed across communities — cooperatives in sparse, mostly white areas, merchant and government generators in denser, more nonwhite ones (Table 3.3). Because organizational type is in this way correlated with community demographics — and, as Chapter 2 showed, with abatement itself — an uncontrolled demographic gradient could in part reflect the distribution of firm types across places rather than any response to the community. The primary specification therefore conditions on organizational structure, carried over from Chapter 2, so that the demographic coefficient is read net of firm type.[6]

Table 3.3: Community context by organizational type: median neighborhood density (census tracts within 3 km) and median 3 km buffer composition, pooled over the three snapshots.

Organizational type	Plant-years	Tracts in 3 km	% nonwhite	% poverty	Median income
Cooperative	314	2	14%	12%	$58,669
Regulated IOU	1,229	3	23%	13%	$57,549
Merchant	1,828	5	33%	13%	$59,918
Host-serving	749	5	39%	17%	$52,620
Government	669	6	34%	15%	$55,676

[6] A firm-fixed-effects specification, reported among the robustness checks, pushes this further by posing the within-firm question — whether a given firm abates less at its plants in more disadvantaged communities. Because it is identified only off firms that own several plants across demographically different communities, it speaks to that subset rather than to the full fleet, and a null result would itself be informative: it would point to any emissions gap as a product of the sorting of firms across places — which firms own where — rather than of within-firm behavior across the communities a single firm spans.

3.4.11 Spatial Autocorrelation

Stated plainly at the outset: the main analysis addresses spatial autocorrelation with Conley spatial-HAC standard errors, computed on a twenty-five-kilometer band with fifty kilometers as robustness. The estimator remains ordinary least squares; the correction adjusts only the standard errors. The rest of this section sets out the reasoning: the threat, the diagnostic that measures it, and the alternatives held in reserve.

Plants are not scattered independently across the country, and neither are their emissions. Facilities close together share airsheds, weather, regional fuel markets, and the regulators who oversee them, and near things are, as a rule, more related than distant ones (Tobler 1970). The unexplained part of one plant's emissions intensity is thus liable to resemble its neighbors', so that spatial dependence in the residuals is a plausible threat to inference. The utility-clustered standard errors used throughout absorb the correlation among a single operator's plants but not this distinct, geographic dependence among plants of *different* owners. Left uncorrected, this dependence would understate the standard errors and invite over-rejection (Franzese and Hays 2007).

The research design does not correct beforehand: it first asks whether spatial dependence is present, and in what form, before any correction is applied. Spatial weights W are built on a k-nearest-neighbor graph, block-diagonal by snapshot year so that neighbors are contemporaneous and no plant is tied to another year's fleet, the orthodox structure for a spatial panel (Elhorst 2014; Anselin 2009). The graph is row-standardized with $k = 8$ in the primary specification and $k = 5$ and $k = 10$ reported as robustness. This k-nearest-neighbor graph is the *connectivity structure*: it is used to compute the residual Moran's I and, if needed, to estimate the spatial error model, and both uses require every plant to have neighbors. An adaptive graph is therefore preferred to a fixed-distance band. The fleet is too unevenly spaced for any single band to reach the most isolated plants; a band wide enough to do so (several hundred kilometers) would collapse much of the dense East into one neighborhood. A distance band has its place for the distance-based corrections below, where a plant with no near neighbor *should* stand alone rather than be forced into a graph. The predictor's global Moran's I is nearly flat across this range of k, so the analysis is reported for a single graph ($k = 8$), with the $k = 5$ and $k = 10$ graphs run only to confirm the findings do not hinge on that choice. Moran's I is then computed on the baseline model's residuals to test whether

spatial dependence is present and how strong it is (Anselin et al. 1996).

This step-wise approach matches the correction to the severity of the spatial dependence rather than imposing one by default. The likely primary remedy is a set of Conley spatial heteroskedasticity- and autocorrelation-consistent standard errors (Conley 1999), computed over a fixed distance band (twenty-five kilometers in the primary, fifty as a wider-bandwidth robustness), which leave the point estimates untouched and adjust only their precision for the dependence Moran's I detects. Should that dependence prove large enough that corrected errors are an inadequate account on their own, a spatial error model is estimated in its place (Anselin 2009). That model folds the autocorrelation into the disturbance structure, and it reverts to the connected k-nearest-neighbor graph because it needs every plant to have neighbors. Both target the same inferential threat: dependence among the residuals.

Residual dependence is not the only way geography enters. Because the demographic predictor is itself strongly clustered (Figure 3.1), a plant's emission rate may track not only its own community's composition but its neighbors' — a spatially lagged covariate, the WX term, rather than a nuisance in the disturbances (LeSage 2009; LeSage and Pace 2014). (Moran's I is an index of spatial clustering, 0 under spatial randomness and rising toward 1 as like values increasingly neighbor like; computed on the predictor and the plant geometry alone, it establishes that the precondition for residual dependence is present, not that the residuals themselves are dependent.) Whether that is plausible turns on how physically close plants actually are. The data show they are close: in every snapshot about half the fleet has another plant within ten kilometers and roughly three-quarters within twenty-five (Table 3.4), so neighboring plants routinely share a metropolitan area, an airshed, and a regulator. The design therefore adds a spatially-lagged-covariate specification as a planned robustness check (each plant's own demographic block entered alongside the population-weighted demographics of its neighbors) on the same twenty-five-kilometer band as the Conley correction, so one definition of "neighbor" governs both the nuisance and the substantive readings of spatial structure. Two limits apply. The reading stays associational: entering neighbors' demographics checks whether the own-community association survives regional context; it is not a claim that neighboring communities causally affect a plant's emissions (Betz, Cook, and Hollenbach 2018). And, because the buffers overlap once plants sit within a few kilometers, the own- and neighbor-

demographic measures partly describe the same tracts for the closest plants, so the lagged term is a coarse regional-context check rather than a clean decomposition.[7]

Table 3.4: Plant co-location by snapshot: distance to the nearest same-year neighbor and the share of plants with at least one neighbor within 5, 10, 25, and 50 km.

Year	Plants	Median NN (km)	Within 5 km	Within 10 km	Within 25 km	Within 50 km
2009	1,656	9.0	38%	53%	74%	91%
2018	1,622	10.0	36%	50%	73%	90%
2021	1,530	10.0	35%	50%	72%	89%

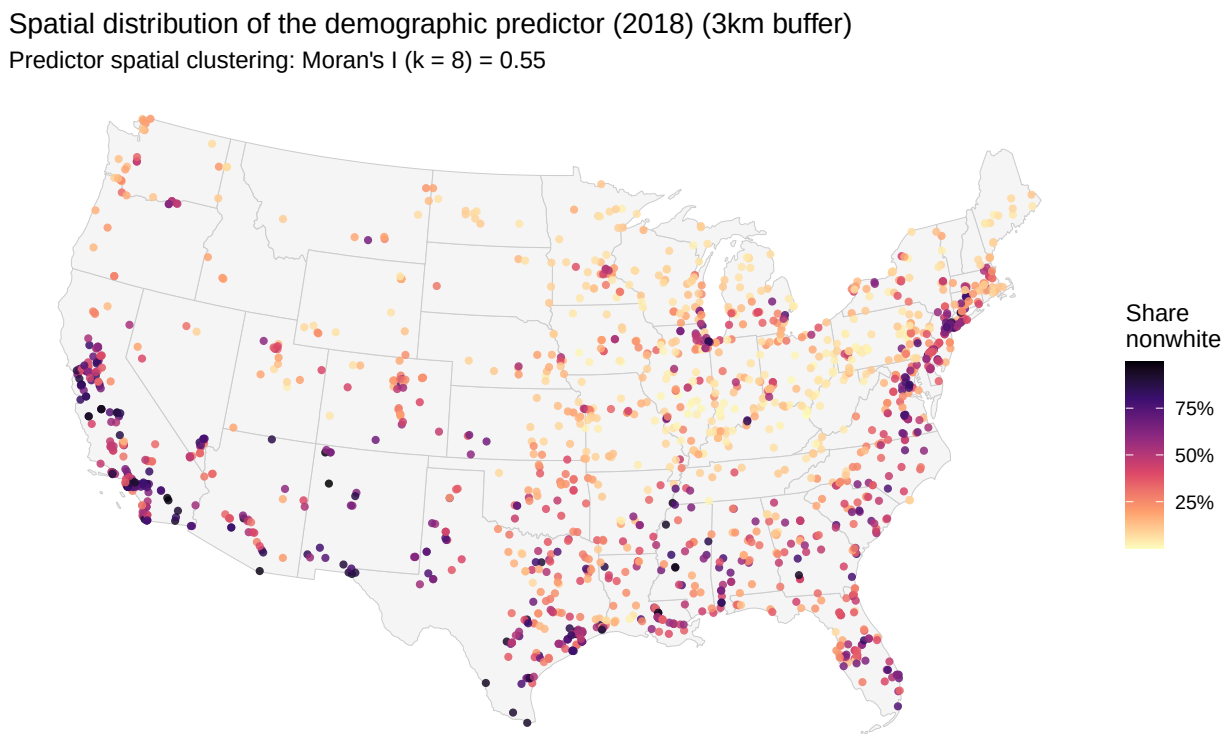

Figure 3.1: The spatial distribution of the demographic predictor: each coal- or gas-fired plant in the 2018 analytical sample, shaded by the share nonwhite of its 3 km population-weighted buffer, with the predictor's Moran's I.

[7] A third channel, distinct from both the residual dependence corrected above and the spatially lagged covariate, is spatial nonstationarity: the possibility that the demographic gradient itself varies across the country, the question a geographically weighted regression addresses (Fotheringham 2009; Gilbert and Chakraborty 2011). No such analysis is conducted here; the design estimates a single national gradient.

3.4.12 Estimation Procedure

The two adoption equations (1a, 1b) are logistic regressions fit by maximum likelihood; the two outcome equations (3a, 3b) are ordinary least squares on `ln_nox_rate_lb_per_mwh`. The demographic block enters in the functional forms fixed above (income logged, the share measures linear; Section 3.4.2), each term interacted with snapshot year (Section 3.4.9). Year and NERC-region fixed effects enter as factors throughout.

Standard errors are cluster-robust (CR1) on the operating utility (`vcovCL`/`coeftest`, sandwich). Clustering absorbs the dependence among a single operator's plants; the distinct geographic dependence among plants of *different* owners is handled separately by the spatial corrections of Section 3.4.11. Those corrections act on the variance side alone: the Conley spatial-HAC errors re-estimate the precision of the same OLS point estimates, so the estimator remains ordinary least squares throughout, and only the standard errors change. Each equation is fit on its own complete-case sample — the adoption logits on plant-years with non-missing mediator, demographics, and controls, the outcome equations on the further subset that also report a measurable NO_x rate.

Comparisons of logit coefficients can be deceiving: a difference can reflect unequal unobserved heterogeneity across the two outcomes rather than a genuine difference in effect (Allison 1999; Mood 2010). However, the deceptiveness of logit coefficient comparisons attaches only to a particular reading: it arises when a logit coefficient is taken as an effect on an unobserved *continuous* response — a latent propensity to abate, e.g. — and matters only if that latent variable is the substantive target. Where the object of interest is instead the *observed* outcome (here whether a plant adopts control), effects on the probability scale remain comparable across margins and models so long as they are correctly estimated (Kuha and Mills 2020). Two particular features of the research design keep the comparisons on that footing: the post-combustion-minus-combustion differential is reported as a difference in average marginal effects (Breen, Karlson, and Holm 2018), and the race-net-of-class nested entry within the logits is decomposed with the KHB method, which separates true confounding from rescaling (Karlson, Holm, and Breen 2012). The outcome equations, being linear, are not subject to this rescaling.

3.5 Results

3.5.1 Descriptive Statistics

Table 3.5 reports descriptive statistics for the analytical sample — the 4,763 plant-years that enter the outcome equations, pooled across 2009, 2018, and 2021 — organized by the role each variable plays. The demographic block carries both the pooled percent-nonwhite measure of the primary analysis and the four component shares (Black, Hispanic, Indigenous, Asian) of the exploratory decomposition; the latter make plain that the Indigenous share, and to a lesser degree the Asian share, is small at the three-kilometre buffer, which bears directly on how much weight a per-group coefficient can carry.

Table 3.5: Descriptive statistics for the analytical sample (pooled over 2009/2018/2021), by variable role.

Variable	Mean	SD	Min	Max
Outcome				
NOx rate (lb/MWh)	1.42	1.64	0.00	9.98
Log NOx rate	-0.44	1.48	-6.91	2.30
Demographics (3 km)				
% nonwhite (pooled)	0.354	0.259	0.001	0.999
% Black	0.115	0.168	0.000	0.990
% Hispanic	0.170	0.207	0.000	0.991
% Indigenous	0.010	0.048	0.000	0.984
% Asian	0.029	0.056	0.000	0.639
% poverty	0.155	0.089	0.000	0.792
Median HH income ($)	$62,062	$24,334	$12,390	$245,950
Mediators				
Combustion control adopted	0.77	0.42	0.00	1.00
Post-combustion adopted	0.46	0.50	0.00	1.00
Controls				
Log nameplate capacity	5.57	1.27	3.26	8.71
Capacity factor	0.317	0.265	0.000	0.999

Generator age (yrs)	28.0	16.5	0.0	88.1
Has steam turbine	0.48	0.50	0.00	1.00
Has gas turbine	0.42	0.49	0.00	1.00
Coal heat share	0.207	0.391	0.000	1.000
State climate-policy count	2.8	2.0	0.0	7.0
Democratic lean index	-0.01	0.82	-1.87	2.01
In ISO/RTO	0.64	0.48	0.00	1.00
Third-party ownership	0.12	0.33	0.00	1.00
Number of owners	1.2	1.0	1.0	16.0
Firm scale index	0.01	0.97	-1.99	2.05
State market-share index	-0.00	0.94	-0.55	9.17
FERC cogeneration	0.17	0.38	0.00	1.00
FERC small power	0.04	0.19	0.00	1.00
Ozone nonattainment	0.33	0.47	0.00	1.00

Note. Analytical sample: N = 4,763 plant-years, pooled 2009 / 2018 / 2021. Shares and binary indicators are proportions (0–1). The five-level organizational-structure control is omitted (see the research design).

The three primary demographic measures distribute as the functional-form choices in Section 3.4.2 anticipate (Figure 3.2): the two proportion measures are bounded and roughly continuous, while median household income is sharply right-skewed in levels and approximately symmetric once logged, so the figure plots it on the log scale, where that symmetry is apparent. Plotting each measure across the four linkage geographies the analysis carries — the nearest-centroid tract and the 3, 5, and 10 km population-weighted buffers — also makes the radius-robustness rationale visible: because a wider buffer averages over more tracts, the larger rings are smoother and less dispersed than the nearest-tract measure while preserving the same central tendency.

Figure 3.2 compares the measures across the four geographies: widening the buffer averages over more tracts, so the larger rings are smoother and less dispersed while holding their center.

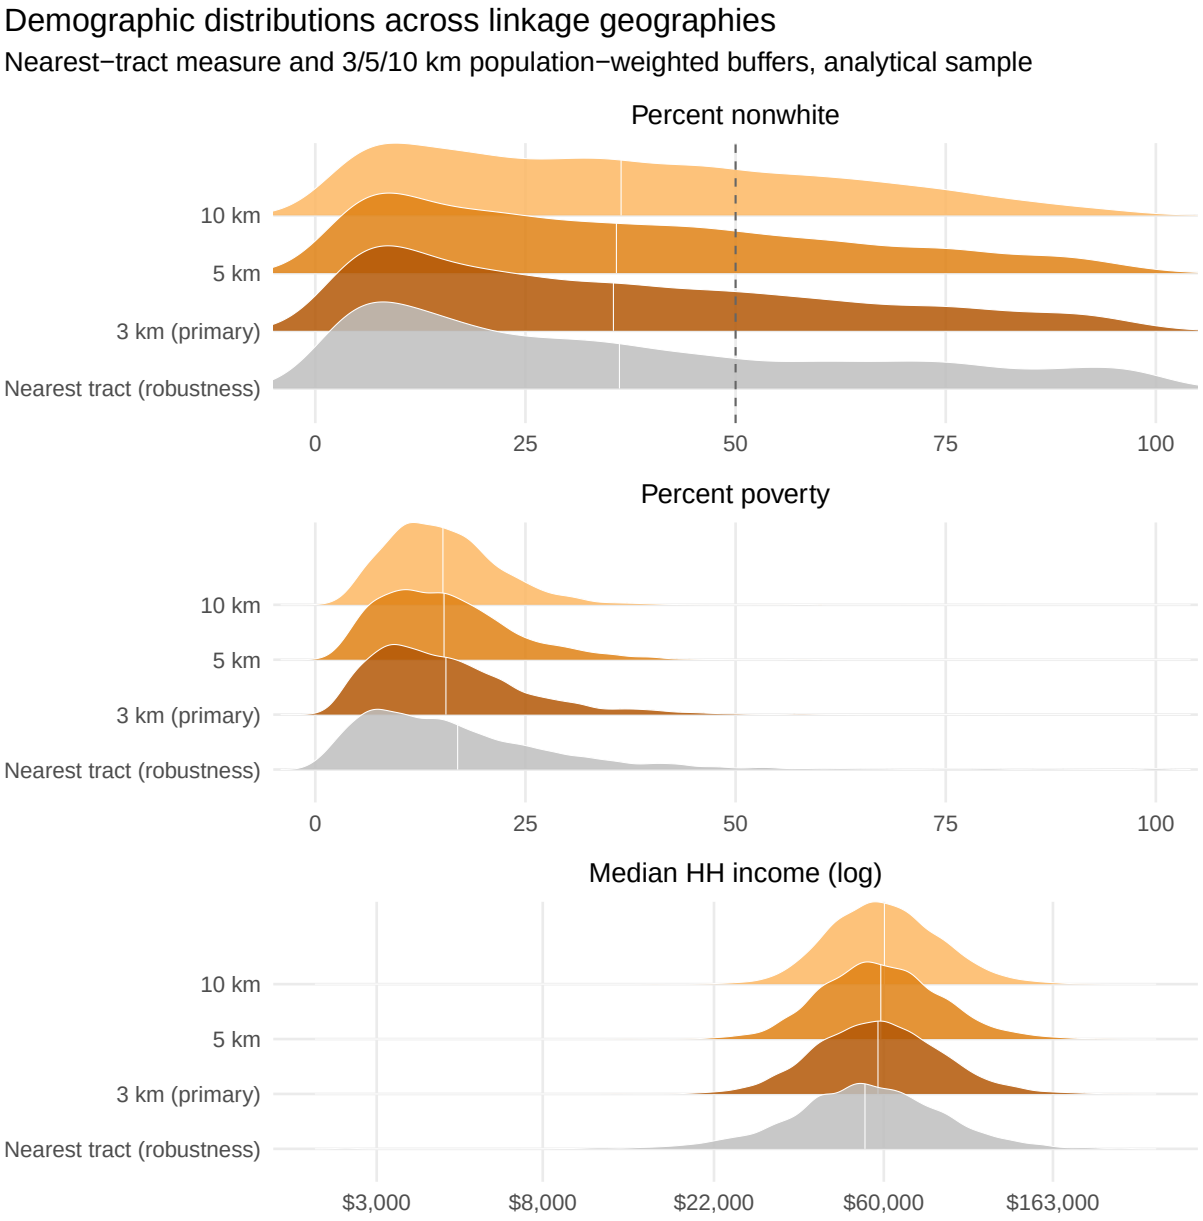

Figure 3.2: Distributions of the three demographic measures across the four linkage geographies: the nearest-centroid tract measure (robustness) and the 3 (primary), 5, and 10 km population-weighted buffers. Percent nonwhite and poverty are on a 0–100 scale; income is in logs. The dashed line marks the majority-minority (50%) threshold used in robustness.

That the rate and the mass are genuinely distinct outcomes, rather than two views of the same quantity, is visible in the fleet itself (Figure 3.3). At any given tonnage, emission rate spans orders of magnitude: the high-tonnage baseload runs at a moderate rate, while the highest rates belong to small, older, less-efficient units. A demographic gradient can therefore surface in intensity without

surfacing in mass.

For NO_x in particular, the rate may be the more telling of the two where *local* harm is concerned. The small, low-capacity-factor units that post the highest rates emit in concentrated bursts when they fire to meet peak demand, often on the hot, high-ozone afternoons when NO_x most readily forms ozone (Richards 2000). It is those episodic pulses, not a plant's averaged annual tonnage, that shape the acute exposure a neighboring community actually breathes. In the peaker fleet, that burden falls disproportionately on nonwhite and low-income urban neighborhoods (Krieger, Casey, and Shonkoff 2016; Alliance et al. 2020; Ramanan 2026).

The two outcome measures diverge (Figure 3.3): at any given tonnage the emission rate spans orders of magnitude, with high-tonnage plants predominantly coal running at moderate rates and the highest rates belonging to small, often older and less efficient units. Intensity and mass are therefore distinct exposures, which is why a demographic gradient can appear in the rate without appearing in mass.

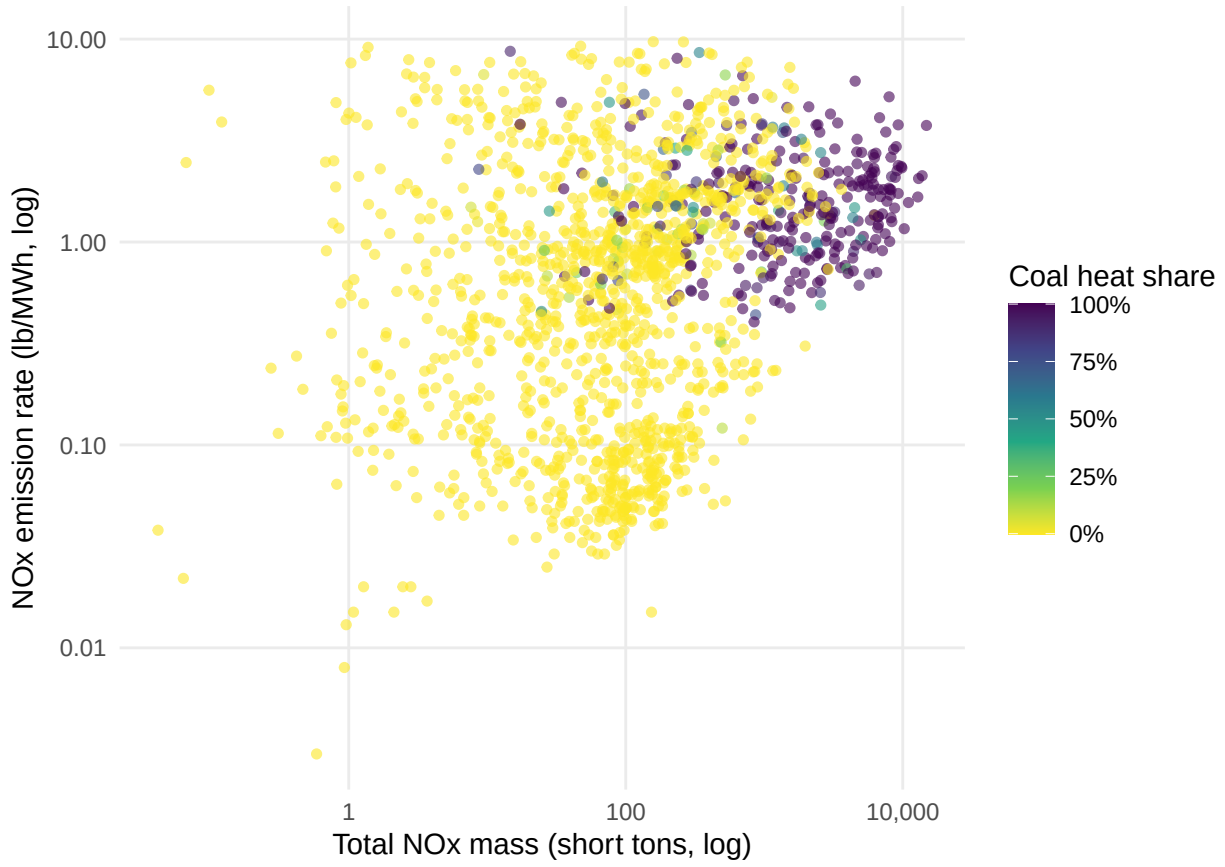

Figure 3.3: The 2018 fleet by its two outcome measures, total NO_x mass (x) and emission rate (y), both log-scaled, colored by the plant's coal heat share.

3.5.2 Primary Results

3.5.2.1 Adoption: no demographic gradient

Do plants in disadvantaged communities install fewer or cheaper NO_x controls? Across 4,788 plant-years, no (Figure 3.4; Table 3.6). Percent nonwhite predicts the adoption of neither combustion modification (Eq 1a) nor post-combustion control (Eq 1b), and the core claim of H1b, that disadvantage concentrates on the expensive post-combustion margin, finds no support: the percent-nonwhite differential is essentially nil, a 0.005-point change in adoption probability per tenth of community nonwhite share ($p = .423$). With the racial association estimated tightly around zero on both margins, this is a precise null rather than an underpowered one — the analysis rules out a meaningful racial gradient in adoption, not merely fails to detect one.

Nor do the class measures assemble into the predicted pattern. On the post-combustion margin, higher income (0.020 per $10,000, p = .010) *and*, against the logic that disadvantage suppresses investment, higher poverty (0.041 per tenth of population share, p = .019) each predict *more* adoption, not less. That income runs in the expected direction points not to any firm-side cost of capital but to community resistance capacity — wealthier communities command the regulatory and political leverage to press for abatement (Section 3.2.3) — and the deeper anomaly is poverty running the same way, which that logic would send in the opposite direction.[8] In other words, whatever produces the demographics-to-emissions disparity the analysis documents is not a product of differential adoption or cheaper technology procurement, but is emergent in the day-to-day operation of particular, emissions-intensive plants.

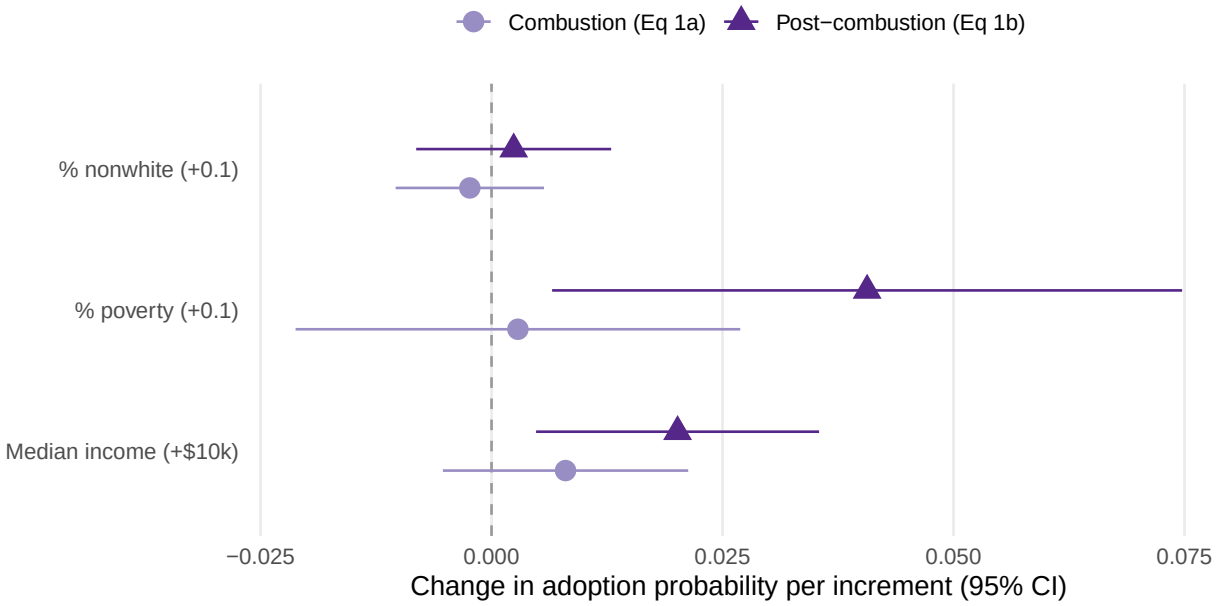

Figure 3.4: Average marginal effects of community demographics on the probability of adopting each NO_x control, by margin (combustion, Eq 1a; post-combustion, Eq 1b), utility-clustered. Effects are scaled per increment (percent nonwhite and poverty per +0.1 of population share; income per +$10,000) and reported on the probability scale (Breen 2018).

[8] One plausible resolution: post-combustion retrofits were federally driven onto the largest emitters under the CAIR and CSAPR programs, and those big coal units skew rural and poor — so the poverty coefficient may be proxying a large, federally-mandated emitter rather than community power. Testable against plant size and mandate status, but not adjudicated here.

Table 3.6: Average marginal effects of community demographics on adoption of each NO_x control and the post-minus-combustion differential, utility-clustered; scaling as in Figure 3.4, on the probability scale.

Demographic	Combustion (Eq 1a)	Post-combustion (Eq 1b)	Differential (post − comb.)
% nonwhite (per +0.1)	-0.002 (0.004)	0.002 (0.005)	0.005 (0.006)
% poverty (per +0.1)	0.003 (0.012)	0.041* (0.017)	0.038† (0.020)
Median income (per +$10k)	0.008 (0.007)	0.020** (0.008)	0.011 (0.010)

Note. Cells: AME (SE). Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

3.5.2.2 Emission intensity: an emerging racial gap given controls

The gradient that is absent in adoption reappears in operation. Conditional on everything that fixes a plant's baseline emissions (fuel, capacity, vintage, prime mover) and on the controls it carries, plants are, by 2021, sorted by the race of their surrounding communities: a more-nonwhite community's plants run measurably dirtier per megawatt-hour than a whiter community's. The demographic-by-year design locates the gap in time. The percent-nonwhite × 2021 term is positive and significant in both the total association (Eq 3a, 0.42, p = .012) and conditional on adoption (Eq 3b, 0.38, p = .018). The magnitude, though, is modest — on the scale the enforcement-disparity literature reports for race and class (Konisky and Schario 2010). The net percent-nonwhite slope in 2021 is 0.21 log-points of NO_x rate per unit of share, so a plant whose surrounding community sits at the 75th percentile of nonwhite share (54%) is associated with roughly 9% higher NO_x intensity than one at the 25th percentile (13%) — about 2.1% for every ten-percentage-point increase in nonwhite share (Table 3.7).

The racial gradient identified above is a development, not a constant. In 2009, and again in 2018, the gradient is statistically indistinguishable from zero and, if anything, faintly negative at the start of the period (-0.21, p = .203). Only by 2021 does it turn positive and significant (Figure 3.5). The trajectory also shows (this matters most for the mechanism) that conditioning on abatement does almost nothing to it: moving from the total association to the model that holds the two control binaries fixed (Eq 3a → 3b) barely moves the 2021 term (0.42 → 0.38). The controls themselves behave as expected: combustion modification and post-combustion treatment each lower the NO_x

rate, the latter the more steeply (-0.47 and -0.80 in log-points). Yet the racial gap survives their inclusion almost untouched. The same holds for utilization: capacity factor enters the outcome equations with a strong, precisely estimated negative coefficient, lower-output units running dirtier per megawatt-hour. The racial gap is thus not an artifact of the low duty cycle of peakers; it persists with how hard a plant runs held fixed. There is a reason a peaker runs dirty even when it runs: post-combustion controls scrub NO_x only within a specific catalyst temperature window, and a unit that cycles to follow peak demand spends much of its operating time at output well below rated capacity, producing exhaust too cool to reach that window, so its SCR or SNCR underperforms during the hours it is actually running (Richards 2000; Sorrels et al. 2019). The disparity, in other words, is not in which controls a plant has installed but rather in how dirtily the plant runs given them.

The identified gap survives the two most immediate challenges to it. The first is that percent nonwhite might merely proxy class: entering community poverty and income alongside it leaves the 2021 term essentially unchanged (race-only to full, 0.40 to 0.42 in Eq 3a and 0.35 to 0.38 in Eq 3b), so percent nonwhite is not standing in for community class. The second is that the significance might be an artifact of spatial dependence among neighboring plants: under Conley spatial-HAC standard errors on a twenty-five-kilometer band, the 2021 term stays significant in both equations ($p = .035$; $p = .041$). The full suite of robustness measures, including exposure radius, functional form, the mass outcome, and the firm-level structure that bears on what produces the gap, is taken up in Section 3.5.2.3.

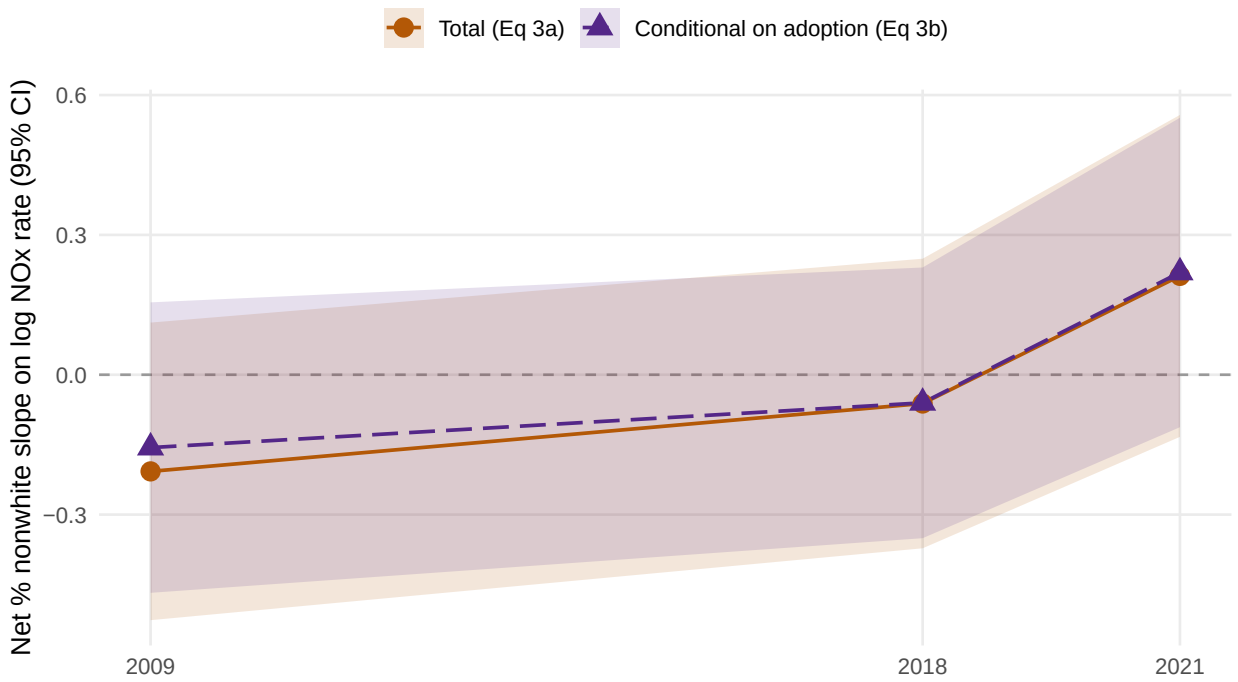

Figure 3.5: The percent-nonwhite gradient in log NO_x intensity by snapshot year: the net marginal slope (main effect plus year interaction) from the primary outcome equations, for the total (Eq 3a) and adoption-conditional (Eq 3b) series, with utility-clustered 95% confidence intervals. Slopes are per unit of nonwhite share (0–1).

Table 3.7: Community demographics and plant-level NO_x emission intensity (log lb/MWh): total (Eq 3a) and adoption-conditional (Eq 3b) associations, pooled with demographic × year interactions; main effects are the 2009 reference.

Term	Total (Eq 3a)	Conditional (Eq 3b)
Demographics (2009 reference)		
% nonwhite	-0.207 (0.163)	-0.156 (0.159)
% poverty	-0.869 (0.536)	-0.401 (0.511)
Log median income	-0.271† (0.147)	-0.083 (0.137)
Demographics × 2018		
% nonwhite × 2018	0.146 (0.163)	0.096 (0.157)
% poverty × 2018	1.051 (0.956)	0.847 (0.943)
Log income × 2018	0.275 (0.219)	0.132 (0.212)
Demographics × 2021		

% nonwhite × 2021	0.419* (0.167)	0.376* (0.159)
% poverty × 2021	0.167 (0.800)	-0.045 (0.783)
Log income × 2021	0.285 (0.182)	0.161 (0.171)
Adoption (Eq 3b)		
Combustion control	—	-0.474*** (0.080)
Post-combustion control	—	-0.803*** (0.064)

Note. Year and NERC-region fixed effects and all controls included (not shown). Utility-clustered SEs in parentheses. Eq 3a: n = 4,763, R^2 = 0.46. Eq 3b: n = 4,763, R^2 = 0.53. The % nonwhite × 2021 term remains significant under Conley spatial-HAC standard errors (3.6). Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

3.5.2.3 Robustness and bounds

The 2021 racial gradient withstands the choices that could most easily have manufactured it. Widening the exposure buffer from three kilometers to five and then ten leaves it intact. Abandoning the linear-in-share form for a coarser majority-minority indicator or a top-versus-bottom-tertile contrast also leaves it intact. And, under Conley spatial-HAC standard errors, the correction the residual diagnostics call for, the 2021 racial gradient remains significant at both the twenty-five- and fifty-kilometer bands (Figure 3.6; Table 3.8). Across every one of these specifications the total and adoption-conditional estimates sit almost on top of one another (the two panels of Figure 3.6 are nearly identical); conditioning on a plant's controls again does little to the gradient.

Two qualifications to the results are in order here. The first is that this is a disparity in rate, not in mass. Carried to the outcome of total NO_x tonnage, the 2021 gradient vanishes (p = .705; p = .814): more-nonwhite communities sit beside plants that run dirtier per unit of generation, not beside plants that emit more in aggregate. The finding is an efficiency disparity, then, rather than a claim about the total pollution burden a community bears. The second qualification concerns where the gradient comes from. Replacing the organizational controls with firm fixed effects (which absorb every fixed difference between operators and identify the demographic term only off firms that own plants in more than one community) attenuates the 2021 gap by about 49 percent (0.38 to 0.19, p = .290), its precision scarcely changed. The within-firm gradient that remains is positive but no longer significant (0.19, p = .290), so a substantial share of the association is between firms rather than within them. That attenuation does not, however, pin a mechanism. The gradient is broad-based, of

comparable magnitude among single-plant operators (0.46, p = .092) and multi-plant firms (0.38, p = .049), and its between- and within-firm components are similar in size and each underpowered, so whether the disparity reflects the sorting of dirtier operators into more-nonwhite communities or differential treatment within firms is left unresolved (Section 3.6). The gap softens further, to marginal, under the nearest-tract measure (p = .060) and under two-way clustering on utility and state (p = .059; p = .040).

The final exercise reweights plants by the people who live beside them, and it strengthens the finding rather than bounding it. Every estimate to this point counts plants, each weighted equally as one siting outcome. Weighting instead by the population each plant's buffer contains (shifting the estimand from the average plant toward the average exposed person) *widens* the 2021 racial gradient in intensity by 56 percent (to 0.66, p = .024; and 0.59, p = .032 conditional on adoption; Table 3.8). The standard errors roughly double (the person-weighted sample is dominated by the dense urban plants that host the most people), but the gradient survives the loss of precision. The plant count, in other words, understates what the exposed population meets: more people live beside the disproportionately dirty plants in more-nonwhite communities than counting plants alone reveals. The same reweighting lifts the mass gradient from near zero toward the positive (0.35), though too imprecisely to disturb the rate-not-mass reading above. This remains a plant-side exercise: the weight is population near a plant, not the dispersion-modeled exposure a tract-side design would compute. But it is a first step from siting toward incidence.

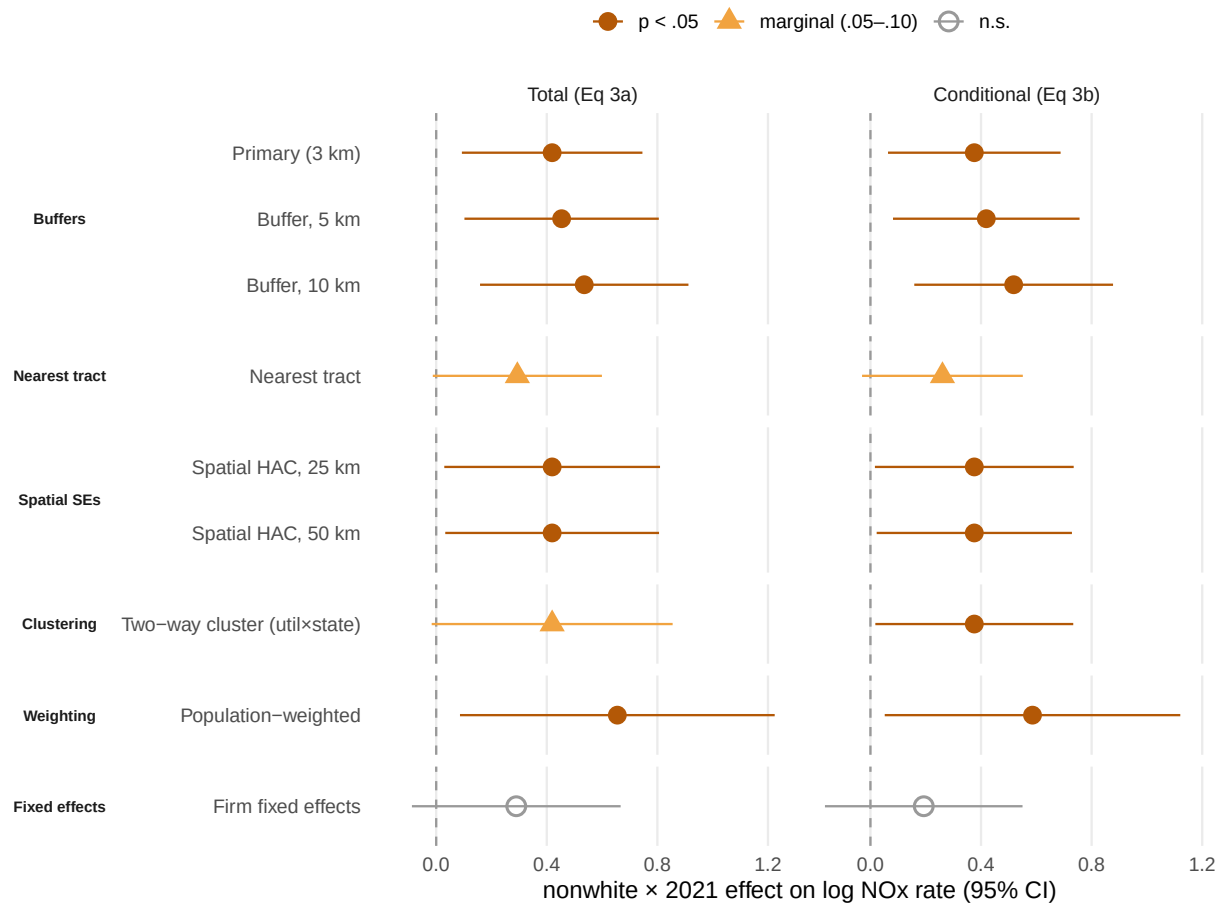

Figure 3.6: The 2021 percent-nonwhite gradient in log NO_x intensity under varied inference, exposure geography, fixed-effects structure, and population weighting, holding the estimand fixed. Each point is the `nonwhite × 2021` coefficient with its 95% confidence interval, for Eq 3a and Eq 3b; color marks significance (solid, p < .05; marginal, .05–.10; n.s., > .10). Functional-form and mass-outcome robustness appear in Table 3.8.

Table 3.8: Robustness of the 2021 percent-nonwhite gradient: each cell is the nonwhite × 2021 coefficient (standard error) from a re-estimation changing one element of the primary specification, for Eq 3a and Eq 3b.

Spec	Total (Eq 3a)	Conditional (Eq 3b)
Buffers		
Primary (3 km)	0.419* (0.167)	0.376* (0.159)
Buffer, 5 km	0.454* (0.179)	0.419* (0.172)
Buffer, 10 km	0.536** (0.192)	0.518** (0.183)

Nearest tract		
Nearest tract	0.294† (0.156)	0.260† (0.148)
Functional form		
Majority-minority dummy	0.189* (0.089)	0.181* (0.086)
Top nonwhite tertile	0.225* (0.099)	0.205* (0.096)
Spatial SEs		
Spatial HAC, 25 km	0.419* (0.199)	0.376* (0.183)
Spatial HAC, 50 km	0.419* (0.197)	0.376* (0.180)
Clustering		
Two-way cluster (util×state)	0.419† (0.222)	0.376* (0.183)
Weighting		
Pop-weighted, rate	0.655* (0.290)	0.586* (0.273)
Pop-weighted, mass	0.352 (0.283)	0.306 (0.261)
Outcome		
Mass outcome (ln tons)	0.063 (0.167)	0.038 (0.162)
Fixed effects		
Firm fixed effects	0.290 (0.193)	0.193 (0.182)

Note. Cells: % nonwhite × 2021 coefficient (SE). The **Functional form** rows rescale the predictor (a majority-minority dummy; top vs. bottom nonwhite tertile) and the **Outcome** row is log NOx short tons, so those magnitudes are not directly comparable to the continuous-share-on-rate rows. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

3.5.2.4 Which communities? A group decomposition (exploratory)

The pooled percent-nonwhite measure treats every nonwhite community alike. A post-hoc disaggregation — exploratory, outside the registered design — replaces it with separate Black, Hispanic, and Asian shares, each entered by year and net of the same class terms and controls. (Indigenous share is near-zero within three kilometers — a median of 0.1 percent, above five percent for only 157 of 4789 plant-years — and is left out of the model.) The three shares are only mildly collinear (variance-inflation factors between 1.58 and 1.98), so their coefficients separate cleanly. The exercise

asks what the pooled gradient cannot: to which communities does the 2021 emergence belong?

It belongs, almost entirely, to the Hispanic share — and is better read as the erosion of an advantage than the onset of a disadvantage. Of the three groups, only the Hispanic share carries a statistically significant demographic-by-2021 interaction (0.52, p = .004, in Eq 3a; 0.42, p = .013, conditional on adoption), and of every per-group year-specific slope in the decomposition only one is individually distinguishable from zero: the Hispanic slope in *2009*, which is significantly *negative* (-0.47, p = .014). In 2009, that is, plants in more-Hispanic communities ran measurably cleaner than the fleet average; by 2021 that gap had closed to zero. The Black share's gradient drifts upward over the decade, but its 2021 interaction never reaches significance (0.21, p = .409); the Asian share, estimated on very small shares, is too imprecise to read in either direction. Split three ways, the pooled signal is too thin to fix a significant 2021 *level* of disadvantage on any single community — what the decomposition locates is the trajectory's source: Hispanic communities losing the cleaner-plant edge with which they began the period (Figure 3.7). (The aggregate series is not the arithmetic sum of the component series: the models differ and pose different counterfactuals, so the overlay compares trajectories rather than decomposing them additively.)

That pattern is the relocation account in demographic form. In 2009 the dirtiest-per-megawatt-hour units were the large rural coal plants, sited away from the dense Hispanic neighborhoods of the metropolitan Southwest and the coastal cities; the gas capacity those communities did neighbor was, comparatively, clean. As the coal fleet retired and the dirtiest unit migrated to the small urban peaker, it migrated toward those same communities, and the 2009 advantage dissolved. The disaggregation thus refines rather than overturns the headline: at this stage of the fleet's transformation, the emerging racial gradient in NO_x intensity is a Hispanic-community phenomenon, and one of convergence from a cleaner starting point rather than a new, group-specific penalty. Because it is post-hoc and not part of the pre-registered design, it is offered as a direction for confirmatory work, not a tested claim.

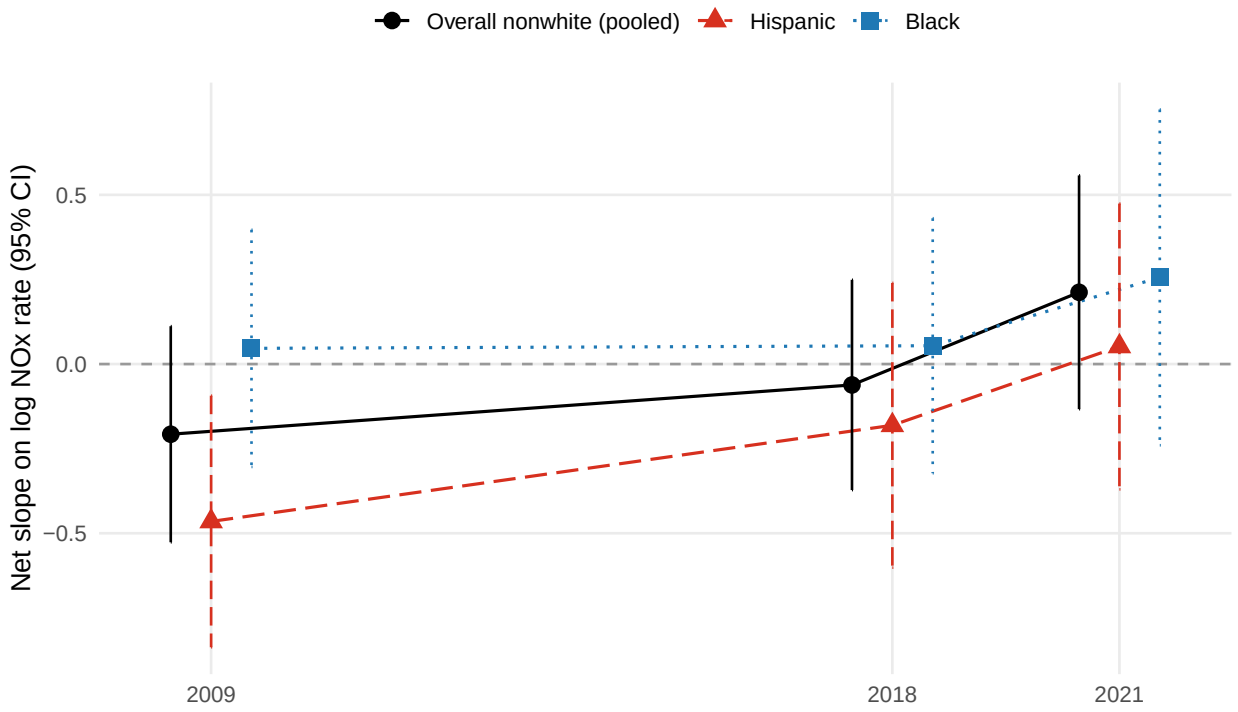

Figure 3.7: Net year-specific slope on log NO_x intensity (Eq 3a): the aggregate percent-nonwhite gradient against the exploratory Black and Hispanic component gradients, with utility-clustered 95 percent confidence intervals; slopes are per unit of population share (0–1). The Asian gradient, whose small shares produce far wider intervals, appears in Table 3.21. Exploratory, post-hoc.

3.6 Discussion

RQ3 asked whether the demographic association with intensity runs through adoption or persists beyond it. Put another way: where, in the process that turns fuel into electricity and pollution, does environmental disadvantage enter? Not, this analysis finds, at the abatement-investment stage. Plants in more-nonwhite communities are no less likely to adopt controls, on either the cheap combustion margin or the capital-intensive post-combustion one, and the adoption differential at the design's center is essentially nil. The gradient appears instead in operation: by 2021, plants in more-nonwhite communities run measurably dirtier per megawatt-hour, conditional on the controls they carry and on how hard they run. The disparity is not in whether a plant has been equipped to run clean, but in how cleanly it is run.

Because the outcome equations hold fixed a plant's fuel, capacity, age, and prime mover (and

the two abatement binaries themselves), the residual racial gap cannot be a matter of equipment; it must lie in how a plant is run. That raises a question beyond the three posed at the outset: what kind of plant runs dirty per megawatt-hour while abated? The profile is consistent: the top decile of 2021 NO_x intensity is not the baseload fleet but its opposite — small, old, and barely running. These units carry a median nameplate of 73 megawatts against 354 for the rest of the fleet, a median capacity factor of 0.03 against 0.32, and a median age of 40 years against 26. They are peakers: aging combustion turbines that fire on startup and at low load to meet demand spikes, in the operating regime where post-combustion controls lose their grip — selective catalytic and non-catalytic reduction both fall off outside a narrow temperature window, and startup emissions escape them altogether. Their rate is high largely because their generation is small, and the same small generation is why they leave no trace on the mass outcome. They sit, too, in modestly more-nonwhite communities than the rest of the fleet (40.7 versus 37.2 percent), and skew toward government and municipal ownership. The peaker is an object of environmental-justice concern in its own right: nationally, an estimated 84 percent of the 63 million Americans living within three miles of one are in higher-than-average-nonwhite communities (Ramanan 2026; Krieger, Casey, and Shonkoff 2016) (Figure 3.8). This is a descriptive cut on raw intensity — a portrait of the dirtiest plants, not a claim about what produces the estimated gradient.

If the disparity is operational, it is also recent — and its emergence has a mechanical, almost incidental, character. In 2009 the highest-intensity decile sat in communities that were *whiter* than the fleet average (a -9.1-point gap); by 2021 it had crossed into more-nonwhite ones (3.5). Over the same decade the dirty tail changed identity: from larger, partly-running units (median 167 MW, capacity factor 0.16) to the tiny idling peakers described above (73 MW, 0.03), as the coal fleet retired and the grid came to lean on small gas and oil units for peak capacity (Figure 3.9). The racial gradient, on this reading, is a byproduct of decarbonization's uneven geography. Cleaning up and closing the rural coal fleet did not so much create a disparity as relocate the dirtiest-per-megawatt-hour unit (from the large rural coal plant to the small urban peaker), and the urban peaker, by the logic of the grid it serves, sits in a metropolitan load pocket. Those units endure not because they are economic but because they are needed for reliability: capacity markets pay them to stand ready for the handful of stressed hours each year, and local transmission constraints can require

them to run regardless of cost. New York City is the emblematic case. Its roughly fifteen oil- and gas-fired peakers operate on the order of six percent of the year, drew some $4.5 billion in standby payments between 2010 and 2019, and sit overwhelmingly in low-income communities of color. They have become the target of a sustained grassroots campaign (the PEAK Coalition) and of the fossil-fuel-phaseout politics that carried Zohran Mamdani first to the State Assembly, from a western Queens district ringed by these units, and in 2025 to the mayoralty (Kaufman 2020; Alliance et al. 2020).

Three qualifications bound this relocation account. The first concerns the apparent conflict with the convergence literature. Hernandez-Cortes, Meng, and Weber (2023) find that the racial gap in electricity-generated $PM_{2.5}$ exposure narrowed sharply between 2000 and 2018, as the dirtiest coal plants retired. An emerging racial gradient might seem to contradict that narrowing but it does not for a few reasons. Their object is modeled ambient exposure, aggregated to secondary $PM_{2.5}$ and to mass; ours is the upstream NO_x *rate*, and the ozone pathway that makes NO_x acutely harmful lies outside a $PM_{2.5}$ transport model altogether. Their window closes in 2018; the gradient here surfaces only by 2021. And the compositional change they credit with the narrowing (coal giving way to gas) is the same change that, on the account above, left behind the small urban peakers in which the new rate disparity appears. The aggregate disparity shrank, in other words, even as what remained of it took a form that a series ending in 2018 could not register.

A second qualification concerns mechanism. The disparity could arise in two ways: dirtier operators could be concentrated in more-nonwhite places (a between-firm story, whether through siting or through the sorting of populations), or firms could run their plants in more-nonwhite communities dirtier than their other plants (a within-firm story). Replacing the organizational controls with firm fixed effects, which compare plants only within the same firm, roughly halves the 2021 term (by about 49 percent; Section 3.5.2.3), so a substantial share of the disparity is between firms. But the within-firm gradient does not vanish: it remains positive (0.19), only too imprecisely estimated to carry weight on its own. Nor is the disparity an artifact of firm scale: it is of comparable magnitude among single-plant operators (0.46) and multi-plant firms (0.38). The two components are similar in size and each underpowered, so the data do not settle which story dominates. Adjudicating between them would demand the operator- and dispatch-level detail this

design does not bring to bear, and is left to future work.

A third qualification narrows what the disparity is. It is a gradient in emission *rate*, not in emission *mass*: the association robust on intensity vanishes on total NO_x tonnage ($p = .705$ for Eq 3a, $p = .814$ for Eq 3b). More-nonwhite communities sit beside plants that run dirtier per megawatt-hour, not beside plants that emit more in aggregate — the high-tonnage baseload lies elsewhere. This bounds the claim. What a community breathes tracks mass more closely than rate, so the result is best read as a disparity in operational *efficiency* — in how cleanly the power serving these communities is made — rather than as a measure of total pollution burden. Even so, a small, frequently-cycled unit can dominate local concentrations during the very high-demand events when it fires, and the coincidence of such units with nonwhite communities is of a piece with the older, marginal, urban capacity that the historical-siting literature traces to redlining and disinvestment (Cushing et al. 2023).

This rate–mass qualification, however, is weaker for NO_x than it would be for most pollutants. For NO_x the harm is episodic rather than annual. Its gravest pathway is secondary: NO_x reacts with volatile organic compounds under sunlight to form ground-level ozone, a reaction that runs fastest on hot, stagnant, high-demand days — the very hours an aging peaker fires. The harm, in other words, arrives in particular hours rather than evenly across the year. When New York's gas peakers run, they can account for more than a third of the city's daily power-plant NO_x; ozone above background is estimated to cause some four hundred premature deaths and several thousand asthma emergencies there each year, a burden falling disproportionately on the high-poverty, nonwhite neighborhoods that host the plants (the South Bronx, Sunset Park, and other communities of color), where ozone-attributable asthma hospitalizations run roughly four times the citywide rate (Alliance et al. 2020). An annualized emission rate, averaging the clean hours against the dirty ones, understates just this — the short, acute window in which a community breathes the worst of what a plant emits. That the design carries a nonattainment control is a partial acknowledgment of the point: nonattainment marks the places where these hour-scale ambient exceedances bind, even where the rate measure cannot resolve the hour itself (US EPA 2016).

One question remains, outside the three posed at the outset and older than any of them: is the gradient racial, or is race standing in for class? Entering community poverty and income alongside

percent nonwhite leaves the 2021 term essentially unchanged, so percent nonwhite is not proxying poverty or income. The literature has long divided on the point: the siting canon tends to find race surviving controls for class (Ringquist 2005; Mohai and Saha 2006), while the enforcement literature finds class the steadier predictor (Konisky and Schario 2010). On this fleet, at this stage of the pollution process, race holds. That the signature is racial, that it is operational rather than a matter of siting or equipment, and that it surfaced only as the fleet began to decarbonize, together mark it as a distinct, and newer, face of environmental injustice in the power sector than the one the proximity literature has mapped.

The disparity documented here is unlikely to be static, because two forces work to keep the dirtiest units on the grid and to run them harder. First, energy transitions have historically been additions rather than substitutions: new capacity accumulates atop the old rather than displacing it, and aggregate fossil generation falls far more slowly than the growth of renewables implies (York and Bell 2019; York 2006). The present decade follows the pattern. Load growth from data centers and artificial intelligence is leading grid operators to delay coal and gas retirements, recommission shuttered plants, and run peakers "more often, turning them into baseload," with PJM, the nation's largest regional grid operator, projecting its own emission intensity to rise rather than fall through 2030 (Ramanan et al., n.d.). Second, climate change works in the same direction: each additional degree multiplies the hot-day cooling spikes that dispatch precisely these old, uncontrolled, urban units, on the stagnant summer days when ozone forms most readily. The two forces are mutually reinforcing, and their joint effect is distributional: rising demand stays the retirement of high-rate capacity, warming widens the high-NO_x windows, and the burden settles again on the communities least equipped to refuse it. The COVID-19 pandemic, which killed at higher rates where chronic air pollution was already worst, gave one measure of what that accumulated vulnerability costs (Alliance et al. 2020).

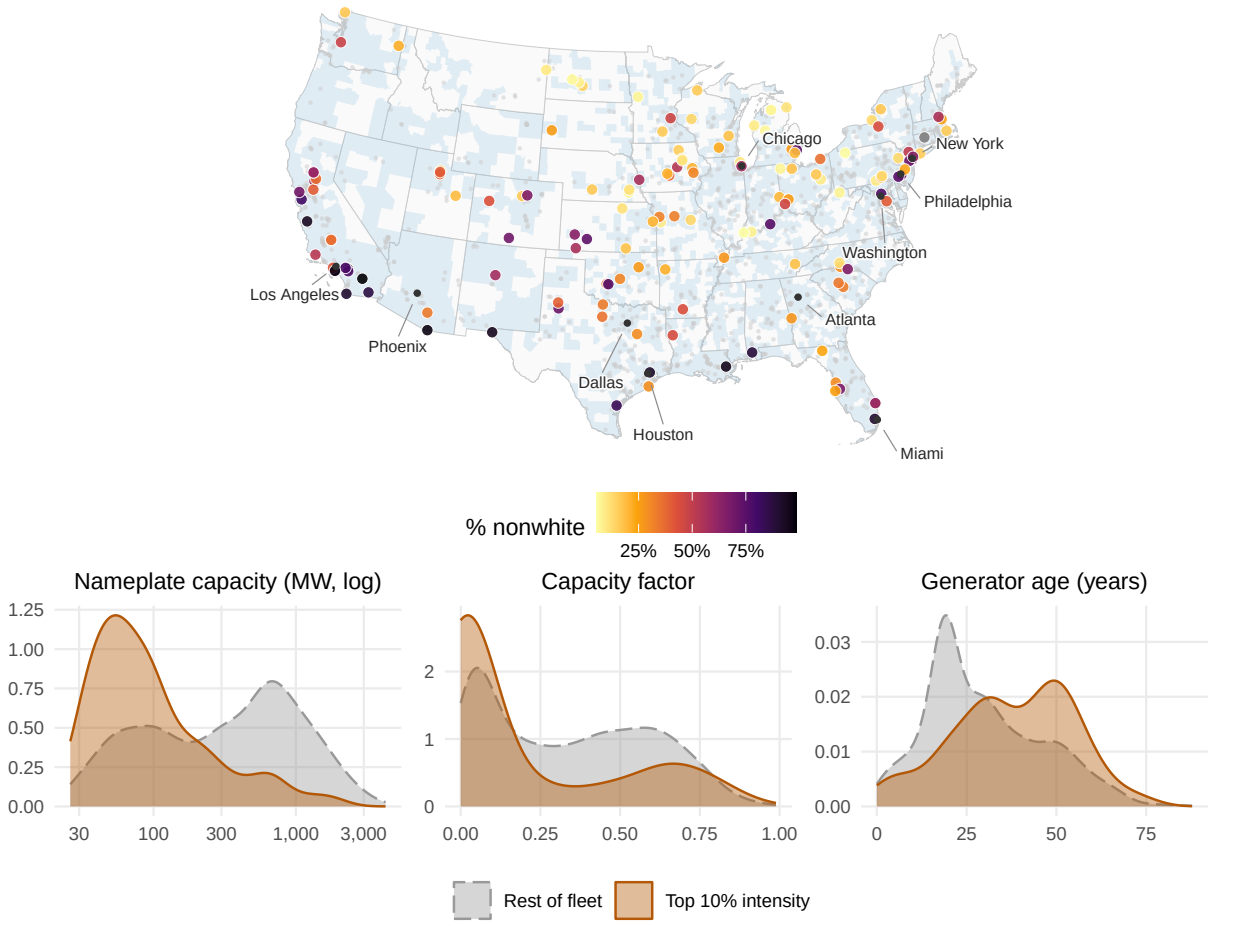

Figure 3.8: The top decile of plants by 2021 NO_x emission rate. Top: the decile (points colored by the surrounding community's percent nonwhite) against the rest of the fleet (gray), over the lower-48 states with metropolitan statistical areas shaded. Bottom: capacity-factor and generator-age densities for the decile against the remaining 90 percent, each normalized within its group. Descriptive (raw intensity), exploratory.

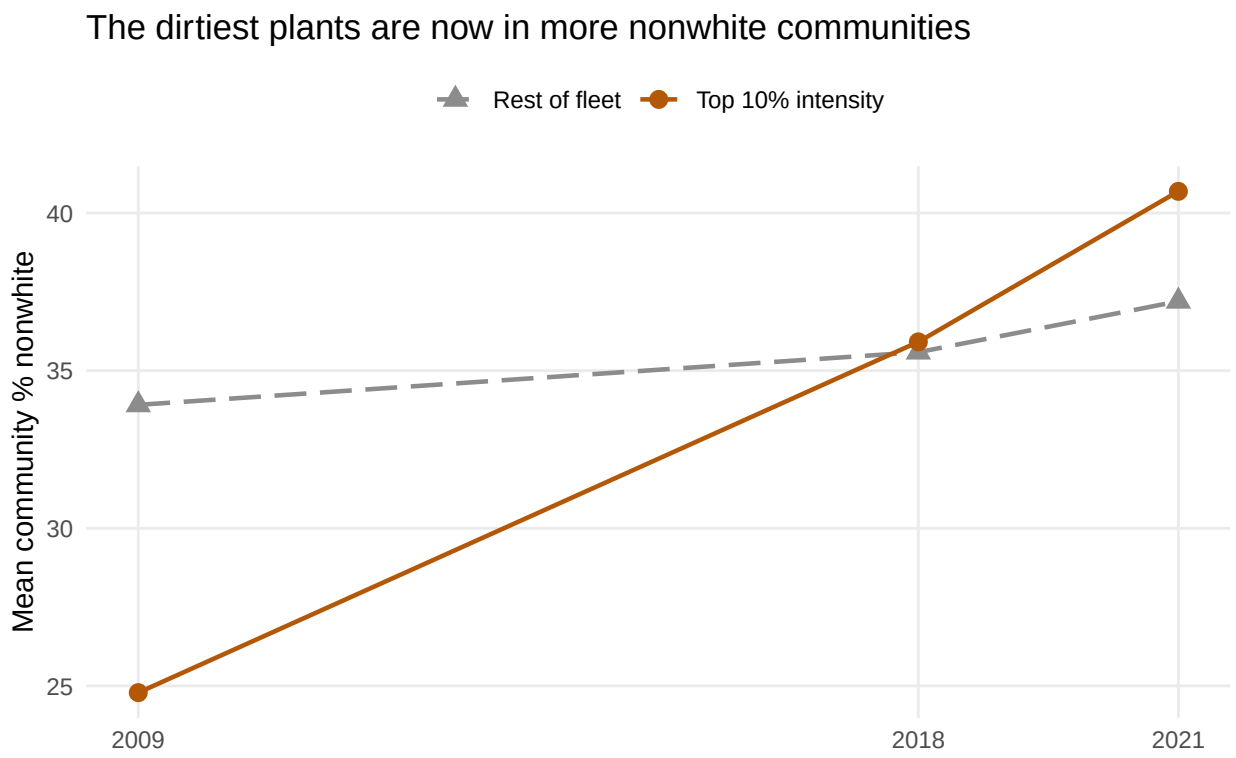

Figure 3.9: Mean community percent nonwhite for the top decile of plants by NO_x emission rate versus the rest of the fleet, by snapshot year. Descriptive, exploratory.

3.7 Limitations

Several limitations bound the claims made in this chapter. The research design is associational, not causal: three cross-sectional snapshots cannot follow a plant through time, and the 2021 emergence is a demographic-by-year interaction rather than a within-plant trajectory. Whether the gradient reflects the siting of dirtier operators in more-nonwhite communities or the sorting of populations around dirtier plants is not resolved here; firm fixed effects offer partial leverage but are too underpowered within firms to settle the question.[9] External evidence bears on the question, if obliquely: Hernandez-Cortes, Meng, and Weber (2023) find that changes in residential location contributed almost nothing

[9]Organizational structure conditions both the level and the distribution of plant emissions (Grant et al. 2010; Grant, Jorgenson, and Longhofer 2018), which raises the possibility that fixed effects defined at the *ultimate* corporate parent rather than the EIA-860 operating entity used here might more precisely isolate the within-corporate-family differential treatment this design cannot. Constructing that hierarchy would require a commercial corporate-affiliation database such as LexisNexis Corporate Affiliates to link operating subsidiaries to their ultimate parents, and is left to future work.

to the national trend in electricity-pollution disparities over 2000–2018, which weighs against the residential-mobility form of the sorting account — though it leaves untouched any cross-sectional sorting already in place at the period's start.

The gradient is also narrower than it may first appear. It is a disparity in emission *rate*, not in mass or ambient concentration, and so is not a direct measure of the pollution burden a community breathes: the high-tonnage baseload that dominates exposure lies elsewhere (Section 3.5.2.3). Nor can the analysis say what drives the residual rate gap. The outcome equations hold capacity factor fixed, so the gap is net of how hard a plant runs; but the channels that would explain why a plant runs dirty per megawatt-hour at a given utilization go unobserved here: the effectiveness and maintenance of installed controls, fuel quality, and the fine structure of dispatch (startup frequency, ramping, low-load operation) that the annual capacity factor cannot resolve. By the same token, holding utilization fixed nets out the dispatch-and-exposure channel on which the discussion's feedback loop turns: whether more-nonwhite communities host plants that are run more often, and breathe more as a result, is a question this rate-based design holds constant rather than answers.

The central result is robust but not invariant. The 2021 gradient softens to marginal significance under the nearest-tract demographic measure and under two-way (utility-by-state) clustering — robust at the core of the specification space, that is, but sensitive at its edges.

Finally, several measures are coarser than the constructs they stand for. Local air-quality regulation enters only as a binary ozone-nonattainment proxy for RACT stringency, leaving finer county- and airshed-level variation unmodeled. The demographic predictor is built by apportioning population across buffers whose underlying tracts shift vintage across the panel (2010 and 2020) and grow noisy at small radii, and it characterizes the neighborhoods of emitters rather than the exposure of populations — an ecological measure, not an individual one. And firm identity is the EIA-860 ownership entity, which does not capture the LLC and parent–subsidiary structures through which operators can shield liability — a gap that bears directly on the unresolved firm-level mechanism.

Two further caveats attach to the demographic measures themselves. The ACS is a sample survey, and its tract-level estimates carry sampling error large enough that the reliability literature recommends screening out estimates whose coefficient of variation exceeds 0.50 (Folch et al. 2016; Raoul S. Liévanos 2018, 2019); no such screen is applied here, and the documented geography

of ACS unreliability, concentrated in the southern United States and in low-income communities and communities of color, means the estimates may be least precise precisely where the gradient under study lives. And the decennial censuses of 2010 and 2020, which measure race and ethnicity with full-count precision, are not used in their place: they carry no income or poverty measures (the long form having been retired after 2000), and their enumeration dates align with none of the three snapshot years, so the ACS supplies all demographic measures, with a decennial-based sensitivity analysis left to future work. The three snapshots also draw their demographics from three different five-year windows (2007–2011, 2016–2020, 2019–2023), and combining distinct multiyear estimates raises measurement issues of its own (Raoul S. Liévanos 2019).

3.8 Conclusion

The environmental-justice signature in the coal-and-gas fleet is not, this study finds, a matter of *who gets pollution controls*. Plants in disadvantaged communities are no less likely to be abated. The disparity lies instead in how cleanly those plants run given the controls they carry. And it is a recent disparity, a byproduct of the fleet's own cleanup: as the dirtiest unit per megawatt-hour relocated from the rural coal plant to the small urban peaker, the gradient followed it into the metropolitan load pocket.

The finding moves the tradition's question. The distributive concern that began with siting, and that this chapter carried to the adoption decision, does not end where the equipment ends: the disparity survives the equipment and lives in operation. The implication for measurement is concrete. An annual emission rate was enough to detect the gradient; the harm arrives at finer resolution, in the hour a peaker fires and the window in which ozone forms. The mechanism, whether siting, sorting, or firms' treatment of their own plants, awaits operator- and dispatch-level detail. The urgency of that work is not hypothetical. As electricity demand climbs and the climate warms, the aging, high-rate units at the center of this story are not retiring but running more, concentrating a growing share of the nation's combustion pollution, in its most acute hours, on the communities least equipped to refuse it.

Computational environment

The analysis pipeline was implemented in R, spanning data construction (the ACS buffer apportionment and the EIA/EPA join), estimation, spatial inference, diagnostics, and rendering. The packages used are listed in Table 3.9; full citations appear in the References Cited.

Table 3.9: R packages used in the analysis pipeline (versions as rendered).

Package	Version	Purpose
Estimation		
sandwich	3.1.1	Cluster-robust (CR1) variance estimators
lmtest	0.9.40	Coefficient testing (coeftest)
car	3.1.5	Variance inflation factors (collinearity)
marginaleffects	0.32.0	Average marginal effects; net year-specific slopes
fixest	0.14.1	High-dimensional fixed effects (firm-FE robustness)
Spatial		
sf	1.1.0	Spatial features; buffer apportionment
spdep	1.4.2	Spatial weights and residual Moran's I
conleyreg	0.1.9	Conley spatial-HAC standard errors
tigris	2.2.1	TIGER/Line geographies
Data		
arrow	23.0.1.2	Read and write Parquet data files
here	1.0.2	Project-relative file paths
glue	1.8.0	String interpolation
dplyr	1.2.1	Data manipulation
tidyr	1.3.2	Data tidying and reshaping
tibble	3.3.1	Modern data frames
purrr	1.2.2	Functional iteration
Tables & figures		
ggplot2	4.0.2	Statistical graphics
ggridges	0.5.7	Ridgeline distributions
ggrepel	0.9.8	Non-overlapping map labels
patchwork	1.3.2	Figure composition

gt	1.3.0	Formatted tables
knitr	1.51	Dynamic report generation

Appendix — Variable inventory

Every variable in the analytical dataset is inventoried below by conceptual role: the dependent variable, the community-demographic predictor, the functional NO_x-control mediator, and the control blocks. Each row gives the variable name, a plain-English label, its units, the data source, and a one-line construction recipe. The fleet is the same coal- and gas-fired, $\geq$ 25 MW, CAMD-reporting sample defined for Chapter 2 (Section 3.4.2); the demographic block is what this chapter adds, and the organizational-structure apparatus that served as Chapter 2's primary predictor carries forward here as a control (Section 3.4.10).

3.8.1 Dependent Variables

The dependent variable is NO_x emission intensity (lb/MWh), logged for the outcome equations. Total NO_x mass (what a community actually breathes, the lived reality at the heart of environmental justice) is carried as the robustness outcome (Section 3.4.2).

Table 3.10: Codebook: dependent variables.

Variable	Label	Units	Source	Operationalization
`nox_rate_lb_per_mwh`	NOx emission intensity (primary DV)	lb/MWh	eGRID PLNT	Direct field from eGRID PLNT; primary outcome
`ln_nox_rate_lb_per_mwh`	Log NOx emission intensity	log(lb/MV	constructed	log(nox_rate_lb_per_mwh); used in OLS specifications

3.8.2 Primary Independent Variable: Community Demographics

The primary predictor is community demographics — percent nonwhite, percent poverty, and median household income (ACS five-year) — entered separately so the race gradient reads net of

class, and measured by population-weighted buffer apportionment at 3/5/10 km with a nearest-tract robustness variant (Section 3.4.2, Section 3.4.5.1).

Table 3.11: Codebook: community demographics, buffer measures.

Variable	Label	Units	Source	Operationalization
`pct_nonwhite_3km`	Share nonwhite, 3 km buffer	fraction 0–1	ACS 5-yr × TIGER tracts	Population-weighted mean of tract % nonwhite over tracts intersecting the 3 km buffer; weight = tract pop × (area inside ÷ tract area).
`pct_poverty_3km`	Share below poverty, 3 km buffer	fraction 0–1	ACS 5-yr × TIGER tracts	Population-weighted mean of tract poverty rate within 3 km (same apportionment weight).
`median_hh_income_3km`	Median household income, 3 km buffer	USD	ACS 5-yr × TIGER tracts	Population-weighted mean of tract median household income within 3 km (same weight). Logged in the primary spec.
`pct_nonwhite_5km`	Share nonwhite, 5 km buffer	fraction 0–1	ACS 5-yr × TIGER tracts	As pct_nonwhite_3km, 5 km ring (radius robustness).
`pct_poverty_5km`	Share below poverty, 5 km buffer	fraction 0–1	ACS 5-yr × TIGER tracts	As pct_poverty_3km, 5 km ring (radius robustness).
`median_hh_income_5km`	Median household income, 5 km buffer	USD	ACS 5-yr × TIGER tracts	As median_hh_income_3km, 5 km ring (radius robustness).
`pct_nonwhite_10km`	Share nonwhite, 10 km buffer	fraction 0–1	ACS 5-yr × TIGER tracts	As pct_nonwhite_3km, 10 km ring (radius robustness).
`pct_poverty_10km`	Share below poverty, 10 km buffer	fraction 0–1	ACS 5-yr × TIGER tracts	As pct_poverty_3km, 10 km ring (radius robustness).
`median_hh_income_10km`	Median household income, 10 km buffer	USD	ACS 5-yr × TIGER tracts	As median_hh_income_3km, 10 km ring (radius robustness).

Table 3.12: Codebook: community demographics, nearest-tract measures.

Variable	Label	Units	Source	Operationalization

pct_nonwhite_tract	Share nonwhite, nearest tract	fraction 0–1	ACS 5-yr (Census)	% nonwhite of the census tract whose centroid is nearest the plant; nearest-centroid robustness measure.
pct_poverty_tract	Share below poverty, nearest tract	fraction 0–1	ACS 5-yr (Census)	Poverty rate of the nearest-centroid census tract; robustness measure.
median_hh_income_tract	Median household income, nearest tract	USD	ACS 5-yr (Census)	Median household income of the nearest-centroid census tract; robustness measure.

Community Demographics - buffer apportionment example: rural v urban

Tracts straddling the buffer boundary contribute proportionally rather than all-or-nothing (Figure 3.10), which is what separates the apportionment measure from assigning the plant a single host tract.

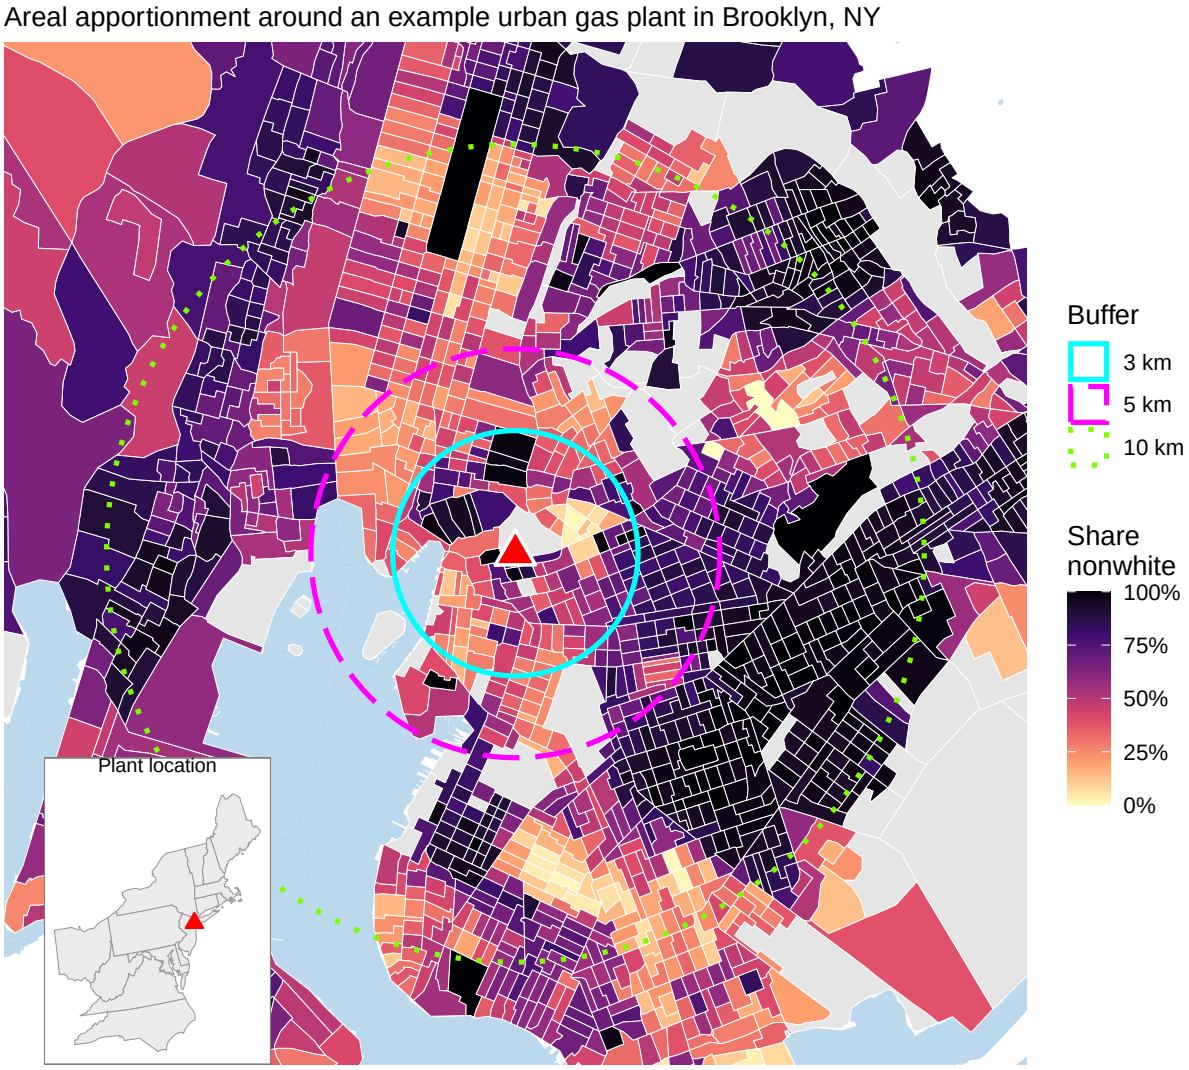

Figure 3.10: Population-weighted areal apportionment around a dense urban example: a natural-gas plant in New York City (inset: location within the U.S. Northeast). Tracts are shaded by share nonwhite; the plant is the red triangle, with rings at 3, 5, and 10 km. Each tract intersecting the 3 km buffer (outlined in black) enters the measure weighted by the share of its area falling inside; grey tracts have no resident population and receive zero weight; tracts beyond 10 km are shown for orientation only.

Table 3.13: Demographic estimates across the buffer radii (3 / 5 / 10 km) and the nearest-centroid tract for the two example plants (2018 snapshot).

	Nearest tract	3 km	5 km	10 km
Brooklyn — urban gas plant				

% nonwhite	78%	52%	52%	60%
% in poverty	11%	22%	18%	16%
Median HH income	$66,722	$88,741	$97,169	$85,718
Kansas — rural coal plant				
% nonwhite	76%	37%	68%	72%
% in poverty	14%	9%	13%	15%
Median HH income	$51,287	$67,462	$54,612	$53,589

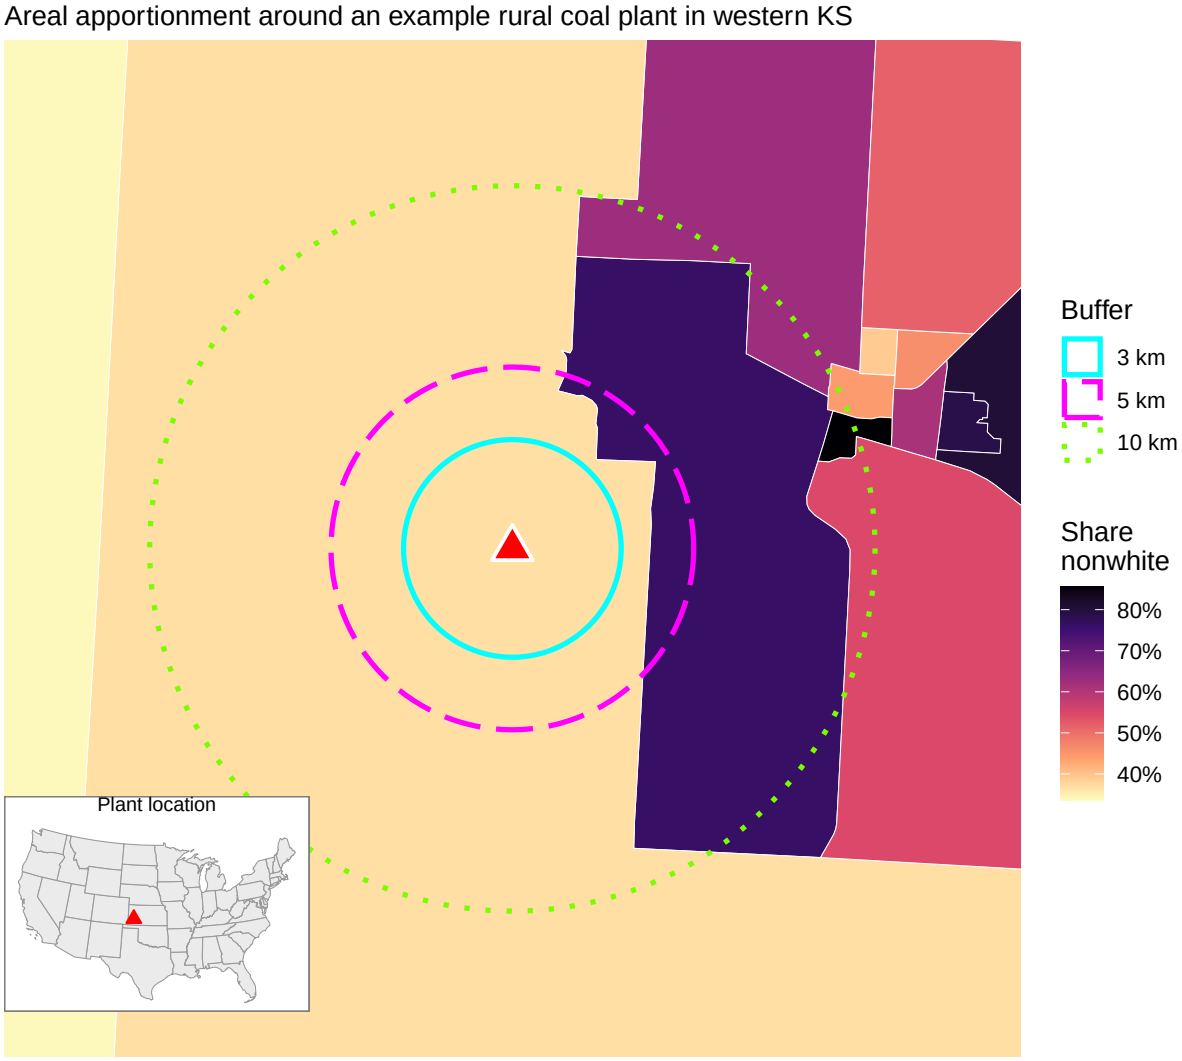

Figure 3.11: The same construction for a sparse rural plant: a coal plant in western Kansas (ORISPL 108; inset: location within the continental U.S.). A single census tract fills the 3 km buffer, so apportionment is immaterial and the radius is what moves the estimate (roughly 37% nonwhite at 3 km, 72% at 10 km). Rings, marker, and shading as in Figure 3.10.

3.8.3 Mediator: Functional NO_x-Control Adoption

The mediator is the pair of fuel-agnostic functional control binaries carried from Chapter 2 — combustion modification and post-combustion treatment (SCR/SNCR) — each flagged from the union of EIA-860 and EPA CAMD records (Section 3.4.2, Section 3.2.4).

Table 3.14: Codebook: functional NO_x-control adoption.

Variable	Label	Units	Source	Operationalization
has_combustion_control	Combustion-modification NOx control present (primary mediator A)	0/1	EIA-860 boiler strategy + EPA CAMD	1 if any in-furnace control present: LNB, OFA, LEA, FGR, steam/water injection, fuel reburn, staged combustion, DLN, cell burner. OR-merge of has_combustion_nox_control (EIA) and camd_has_combustion (CAMD). APTI 418 Ch. 6–8.
has_postcombustion_control	Post-combustion NOx control present (primary mediator B)	0/1	EIA-860 ECE + EPA CAMD	1 if SCR or SNCR present (capital-intensive flue-gas treatment). OR-merge of has_scr|has_sncr (EIA) and camd_has_postcombustion (CAMD). APTI 418 Ch. 11–12.

3.8.4 Control Variables

The controls fall in three blocks: the plant's capacity to emit; the regulatory and political-economic context (including ozone nonattainment); and organizational structure, demoted from Chapter 2's primary predictor to a control here (Section 3.4.2, Section 3.4.10).

Table 3.15: Codebook: capacity-to-emit controls.

Variable	Label	Units	Source	Operationalization
log_nameplate_cap_mw	Log nameplate generating capacity	log(MW)	constructed	log(total_nameplate_mw)
capacity_factor	Capacity factor (plant-level)	fraction 0–1	eGRID PLNT	Net generation ÷ (nameplate × 8,760 hrs).
avg_generator_age_yrs	Average generator age	years	EIA-860 Sched 3 GENERATOR	‘obs_year − mean(operating_year)‘ across plant generators

has_steam_turbine	Steam turbine (ST)	0/1	EIA-860 Sched 3 GENERATOR	1 if any generator has ‘prime_mover_code == ’ST’‘
has_gas_turbine	Combustion/gas turbine (GT)	0/1	EIA-860 Sched 3 GENERATOR	1 if any generator has ‘prime_mover_code == ’GT’‘
coal_heat_share	Share of heat input from coal	fraction 0–1	EIA-923 Sched 5	Sum coal-type heat input ÷ total heat input; 0 for plants burning no coal

Table 3.16: Codebook: state political-context controls.

Variable	Label	Units	Source	Operationalization
state_climate_policy_count	Composite state climate policy count (primary)	0–7 ordinal	State policy panel (constructed)	Sum of the seven binaries below (excluding ‘rps_target_pct‘); primary-spec covariate per §1.1 of paper1_analysis.qmd
dem_lean_index	Democratic-lean index (composite; primary partisan control)	z-score composite	constructed	Standardized mean of dem_governor, hou_dem_share, and sen_dem_share; the partisan control used in place of the five raw measures. Spans Republican-control to Democratic-control monotonically (divided government in the middle), so trifecta_dem and divided_govt are dropped from the spec. The five raw partisan vars are retained for descriptives and the disaggregated-partisan robustness spec.
in_iso_rto	In ISO/RTO territory	0/1	eGRID PLNT	1 if ‘ISORTO‘ is non-null and non-empty; join on ‘orispl‘

ozone_ nonattainment	County in 8-hr ozone nonattainment (snapshot year)	0/1	EPA Green Book PHISTORY × eGRID county FIPS	1 if the plant's county was in nonattainment for any 8-hour ozone NAAQS (1997/2008/2015) in the snapshot year (PHISTORY pw_ YYYY non-blank). County FIPS from eGRID PLNT (FIPSST+FIPSCNTY). Controls for RACT-mandated NOx abatement (Paper 2 confound: ozone-nonattainment areas are disproportionately urban, correlated with the demographic IV). Maintenance status not separated — clean binary. See R/greenbook/02_ plant_nonattainment.R.

Table 3.17: Codebook: organizational-structure controls.

Variable	Label	Units	Source	Operationalization
org_ structure_ type	Organizational type (typology factor; primary regulatory-status IV)	factor (5 levels)	EIA-860 Sched 1 UTILITY	Mutually-exclusive ownership/regulatory type from entity_type × regulatory_status: Merchant (reference), Regulated IOU (entity_type I & RE), Cooperative (C), Government (F/M/S/P), Host-serving (IND/COM). Only investor-owned (I) splits by RE/NR; the ~35 pre-2013 plants lacking entity_type are recoded from sector so none are unclassified. Supersedes has_govt_ownership, is_ regulated_utility, and ferc_ewg_ binary in the primary spec.

has_third_party_ownership	Third-party / leased ownership (code W)	0/1	EIA-860 Sched 4 OWNER	1 if any generator has ‘ownership_code == ’W’‘
n_owners	Number of distinct ownership entities	count	EIA-860 Sched 4 OWNER	Count distinct ‘ownership_id‘ values across all generators at plant
firm_scale_index	Firm-scale index (composite; H5 primary)	z-score compos-ite	constructed	Mean of the z-scored logs of owner_n_plants and owner_total_nameplate_capacity_mw; the H5 regressor in place of the two collinear scale measures. Mirrors state_market_share_index. Both raw measures retained for descriptives and robustness.
state_market_share_index	State market-share index (composite; H4 primary)	z-score compos-ite	constructed	Mean of the z-scored log1p of state_capacity_share and state_output_share; the H4 regressor in place of the two collinear (r = 0.87) shares. Mirrors firm_scale_index. Right-skewed (a few near-monopoly plants). Both individual shares retained for descriptives.
ferc_cogen_binary	FERC cogeneration status (PURPA)	0/1	EIA-860 Sched 2 PLANT	Direct field, binarized
ferc_small_power_binary	FERC small power producer certification	0/1	EIA-860 Sched 2 PLANT	Direct field, binarized

Appendix — Full regression tables

Table 3.18 and Table 3.19 report every coefficient in the four primary equations, for readers who want the complete specification behind the curated panels in the Results. Standard errors are CR1, clustered on utility; year and NERC-region fixed effects are included throughout (summarized as a

single line rather than printed).

Table 3.18: Full adoption-equation estimates (Equations 1a and 1b), logistic regression in log-odds.

Term	Combustion (Eq 1a)	Post-combustion (Eq 1b)
Intercept	-9.193* (4.148)	-16.093*** (3.791)
% nonwhite	0.038 (0.390)	0.389 (0.381)
% poverty	-0.035 (1.277)	3.439** (1.320)
Log median income	0.650† (0.365)	1.083** (0.338)
% nonwhite × 2018	-0.067 (0.368)	-0.492† (0.291)
% nonwhite × 2021	-0.617 (0.381)	-0.250 (0.313)
% poverty × 2018	0.040 (1.563)	-1.440 (1.414)
% poverty × 2021	0.758 (1.585)	-1.726 (1.450)
Log income × 2018	-0.664 (0.409)	-0.712* (0.347)
Log income × 2021	-0.375 (0.414)	-0.748* (0.347)
Log nameplate capacity (MW)	0.879*** (0.086)	0.695*** (0.085)
Capacity factor	-0.695* (0.292)	1.363*** (0.278)
Generator age (yrs)	-0.033*** (0.006)	-0.046*** (0.006)
Has steam turbine	0.694*** (0.204)	0.650** (0.203)
Has gas turbine	-0.250† (0.141)	-0.993*** (0.135)
Coal heat share	1.361*** (0.297)	0.240 (0.222)
State climate-policy count	-0.095 (0.058)	0.011 (0.061)
Democratic lean index	0.254* (0.106)	0.326** (0.105)
In ISO/RTO	-0.130 (0.165)	0.374* (0.179)
Org: Regulated IOU	-0.161 (0.246)	-0.036 (0.257)
Org: Cooperative	-0.111 (0.345)	-0.398 (0.318)
Org: Government	-0.449* (0.216)	0.534* (0.226)
Org: Host-serving	-0.162 (0.207)	-0.469† (0.241)
Third-party ownership	0.302 (0.211)	0.054 (0.170)
Number of owners	0.111 (0.095)	-0.082 (0.060)
Firm scale index	-0.061 (0.106)	-0.028 (0.096)
State market-share index	-0.156 (0.125)	-0.029 (0.121)
FERC cogeneration	-0.407* (0.176)	-0.648*** (0.191)
FERC small power	-0.347 (0.349)	-0.277 (0.377)
Ozone nonattainment	-0.309† (0.170)	0.428* (0.167)

Note. Log-odds coefficients; CR1 utility-clustered SEs in parentheses. Year and NERC-region fixed effects included. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

Table 3.19: Full outcome-equation estimates (Equations 3a and 3b), OLS on log NO_x intensity; Equation 3b adds the two adopted-control indicators.

Term	Total (Eq 3a)	Conditional (Eq 3b)
Intercept	4.117* (1.662)	2.021 (1.529)
% nonwhite	-0.207 (0.163)	-0.156 (0.159)
% poverty	-0.869 (0.536)	-0.401 (0.511)
Log median income	-0.271† (0.147)	-0.083 (0.137)
% nonwhite × 2018	0.146 (0.163)	0.096 (0.157)
% nonwhite × 2021	0.419* (0.167)	0.376* (0.159)
% poverty × 2018	1.051 (0.956)	0.847 (0.943)
% poverty × 2021	0.167 (0.800)	-0.045 (0.783)
Log income × 2018	0.275 (0.219)	0.132 (0.212)
Log income × 2021	0.285 (0.182)	0.161 (0.171)
Combustion control	—	-0.474*** (0.080)
Post-combustion control	—	-0.803*** (0.064)
Log nameplate capacity (MW)	-0.409*** (0.035)	-0.260*** (0.034)
Capacity factor	-0.948*** (0.132)	-0.796*** (0.117)
Generator age (yrs)	0.030*** (0.002)	0.022*** (0.003)
Has steam turbine	0.007 (0.095)	0.138 (0.102)
Has gas turbine	0.525*** (0.059)	0.358*** (0.057)
Coal heat share	1.354*** (0.084)	1.443*** (0.081)
State climate-policy count	-0.036† (0.021)	-0.039* (0.019)
Democratic lean index	-0.013 (0.038)	0.047 (0.035)
In ISO/RTO	-0.188** (0.068)	-0.161* (0.067)
Org: Regulated IOU	0.237** (0.087)	0.222** (0.074)
Org: Cooperative	0.282** (0.093)	0.226** (0.074)
Org: Government	-0.133 (0.093)	-0.092 (0.088)
Org: Host-serving	-0.089 (0.113)	-0.178 (0.112)
Third-party ownership	-0.057 (0.080)	-0.026 (0.080)
Number of owners	0.060** (0.022)	0.049** (0.019)
Firm scale index	-0.063 (0.043)	-0.067 (0.042)

State market-share index	0.057† (0.030)	0.044† (0.026)
FERC cogeneration	0.259*** (0.077)	0.134† (0.076)
FERC small power	0.439 (0.287)	0.396 (0.303)
Ozone nonattainment	-0.216** (0.069)	-0.181** (0.068)

Note. OLS coefficients on log NO_x intensity; CR1 utility-clustered SEs in parentheses. Year and NERC-region fixed effects included. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

Table 3.20 and Table 3.21 report every coefficient in the exploratory race disaggregation (Section 3.5.2.4) — the full adoption and outcome equations with the pooled percent-nonwhite share replaced by separate Black, Hispanic, and Asian shares. Controls and Year + NERC fixed effects mirror the primary equations above. Post-hoc, not part of the registered design.

Table 3.20: Exploratory race disaggregation of the adoption equations (1a, 1b), logistic log-odds, with separate Black/Hispanic/Asian shares replacing pooled percent nonwhite.

Term	Combustion (Eq 1a)	Post-combustion (Eq 1b)
Intercept	-10.721* (4.552)	-11.788** (4.087)
% Black	-0.059 (0.532)	-0.470 (0.508)
% Hispanic	0.008 (0.469)	1.185* (0.462)
% Asian	-1.736 (1.496)	3.368* (1.394)
% poverty	0.409 (1.286)	2.474† (1.350)
Log median income	0.790* (0.403)	0.699† (0.365)
% Black × 2018	0.727 (0.540)	-0.041 (0.416)
% Hispanic × 2018	-0.430 (0.404)	-0.702* (0.342)
% Asian × 2018	0.298 (1.590)	-2.105 (1.290)
% poverty × 2018	-0.148 (1.638)	-0.875 (1.457)
Log income × 2018	-0.632 (0.493)	-0.445 (0.399)
% Black × 2021	0.130 (0.584)	0.020 (0.456)
% Hispanic × 2021	-0.848* (0.420)	-0.479 (0.373)
% Asian × 2021	-0.958 (1.525)	-2.643* (1.303)
% poverty × 2021	0.706 (1.643)	-0.714 (1.528)
Log income × 2021	-0.249 (0.484)	-0.391 (0.396)
Log nameplate capacity (MW)	0.882*** (0.086)	0.685*** (0.085)
Capacity factor	-0.728* (0.295)	1.344*** (0.278)
Generator age (yrs)	-0.034*** (0.006)	-0.046*** (0.006)
Has steam turbine	0.684*** (0.204)	0.671** (0.204)

Has gas turbine	-0.254† (0.142)	-1.017*** (0.135)
Coal heat share	1.386*** (0.296)	0.246 (0.223)
State climate-policy count	-0.087 (0.059)	0.001 (0.061)
Democratic lean index	0.284** (0.109)	0.331** (0.107)
In ISO/RTO	-0.141 (0.173)	0.327† (0.186)
Org: Regulated IOU	-0.152 (0.245)	-0.040 (0.258)
Org: Cooperative	-0.093 (0.339)	-0.426 (0.315)
Org: Government	-0.429* (0.216)	0.508* (0.226)
Org: Host-serving	-0.139 (0.209)	-0.457† (0.244)
Third-party ownership	0.293 (0.209)	0.056 (0.173)
Number of owners	0.111 (0.094)	-0.080 (0.060)
Firm scale index	-0.061 (0.105)	-0.025 (0.095)
State market-share index	-0.175 (0.124)	-0.006 (0.120)
FERC cogeneration	-0.415* (0.178)	-0.646*** (0.192)
FERC small power	-0.342 (0.349)	-0.288 (0.373)
Ozone nonattainment	-0.275 (0.171)	0.384* (0.168)

Note. Log-odds coefficients; CR1 utility-clustered SEs in parentheses. Year and NERC-region fixed effects included. Post-hoc, not registered. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

Table 3.21: Exploratory race disaggregation of the outcome equations (3a, 3b), OLS on log NO_x intensity, with separate Black/Hispanic/Asian shares replacing pooled percent nonwhite.

Term	Total (Eq 3a)	Conditional (Eq 3b)
Intercept	2.509 (1.747)	0.921 (1.607)
% Black	0.046 (0.179)	-0.010 (0.171)
% Hispanic	-0.465* (0.190)	-0.316† (0.186)
% Asian	-1.337 (0.817)	-1.009 (0.791)
% poverty	-0.526 (0.545)	-0.160 (0.523)
Log median income	-0.127 (0.155)	0.016 (0.144)
% Black × 2018	0.008 (0.194)	0.061 (0.187)
% Hispanic × 2018	0.284 (0.188)	0.187 (0.181)
% Asian × 2018	0.452 (1.094)	0.130 (1.101)
% poverty × 2018	0.900 (1.027)	0.780 (1.020)
Log income × 2018	0.212 (0.268)	0.118 (0.264)
% Black × 2021	0.210 (0.254)	0.249 (0.251)

% Hispanic × 2021	0.517** (0.179)	0.419* (0.169)
% Asian × 2021	1.580† (0.885)	1.203 (0.887)
% poverty × 2021	-0.174 (0.822)	-0.263 (0.815)
Log income × 2021	0.117 (0.195)	0.049 (0.186)
Combustion control	—	-0.479*** (0.080)
Post-combustion control	—	-0.796*** (0.064)
Log nameplate capacity (MW)	-0.407*** (0.035)	-0.258*** (0.034)
Capacity factor	-0.939*** (0.132)	-0.794*** (0.116)
Generator age (yrs)	0.030*** (0.002)	0.022*** (0.003)
Has steam turbine	0.004 (0.095)	0.136 (0.102)
Has gas turbine	0.531*** (0.059)	0.362*** (0.056)
Coal heat share	1.344*** (0.084)	1.437*** (0.081)
State climate-policy count	-0.032 (0.021)	-0.036† (0.019)
Democratic lean index	-0.019 (0.038)	0.043 (0.035)
In ISO/RTO	-0.167* (0.071)	-0.148* (0.070)
Org: Regulated IOU	0.238** (0.087)	0.223** (0.074)
Org: Cooperative	0.287** (0.092)	0.229** (0.074)
Org: Government	-0.125 (0.093)	-0.086 (0.087)
Org: Host-serving	-0.094 (0.113)	-0.179 (0.112)
Third-party ownership	-0.057 (0.080)	-0.026 (0.080)
Number of owners	0.059** (0.021)	0.048** (0.019)
Firm scale index	-0.064 (0.043)	-0.067 (0.042)
State market-share index	0.052† (0.030)	0.040 (0.026)
FERC cogeneration	0.258*** (0.077)	0.133† (0.076)
FERC small power	0.441 (0.289)	0.398 (0.303)
Ozone nonattainment	-0.202** (0.069)	-0.170* (0.067)

Note. OLS coefficients on log NO_x intensity; CR1 utility-clustered SEs in parentheses. Year and NERC-region fixed effects included. Equation 3b adds the adopted-control indicators. Post-hoc, not registered. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

4 The Unequal Geography of the Green Transition

The Changing Demographic Composition of U.S. Electricity-Generation Communities, 2010–2023

Abstract

Does the decarbonization of the American power system unwind the fossil era's environmental inequities, or carry them forward under a green banner? Drawing on an annual plant-year panel covering the full continental U.S. generation fleet, fossil, nuclear, and renewable alike, observed from 2010 to 2023, this chapter documents how the demographic composition of each energy type's host communities is redrawn across the transition decade, assigning community demographics to plants by population-weighted buffer apportionment and reading every quantity as divergence from the national trend. Three descriptive analyses — a trajectory of host demographics by fuel, a dynamic decomposition of each fuel's change into entry, exit, and neighborhood-drift channels, and a within- versus between-state analysis of where new capacity lands — return one consistent picture. The transition's spatial demography changes fuel by fuel rather than uniformly, through the turnover of the fleet rather than the churn of neighborhoods, and, for renewables, at the scale of the state rather than the locale: solar's host communities cross from more to less nonwhite than the nation as the buildout disperses into whiter states, while coal's relative whitening is the residue of retirement, the plants that close having stood in more nonwhite communities than those that remain. The account

is deliberately descriptive and plant-side: it characterizes where capacity is sited, not the absolute burden any population bears, and a person-weighted view turns the same fleet systematically more nonwhite. The causal and exposure questions would require a tract-side design, and lie outside this chapter's scope. The demographic map of the clean transition, the analysis shows, is being redrawn less by who moves in beside power plants than by which plants a decarbonizing fleet builds and retires, and where.

4.1 Introduction

Electricity is the lifeblood of modernity: the lights in homes, the hum of industry, the pulse of global trade, and the machinery of war above all, run on nothing more solid than electrons in motion. Yet where electricity's consumption is placeless, its generation is a resolutely material act, embedded in a particular place, with particular neighbors. In the case of fossil fuels, neighbors suffer for their proximity: burning coal, oil, and gas concentrates particulate matter, sulfur dioxide, and nitrogen oxides on the communities nearest the stack (Ebenstein et al. 2017). Though the burden of a wind farm or a solar field is lighter, and different in kind — not smoke and soot but land taken, horizons altered, and open country turned industrial — it is still a categorical burden all the same, and one that communities contest no less fiercely (Mueller and Brooks 2020). Here is the central tension of energy politics: everyone wants the power, and no one wants the plant. The 'NIMBY' reflex this produces is not an aberration of the energy transition but a permanent condition of it, rooted in an attachment to place (Devine-Wright 2009).

What the transition changes is not that tension but its geography. Fossil fuels are compact, transportable, and temporally flexible — sunlight stored in the carboniferous remains of past life, dense enough to cross an ocean or ride a railroad and burn on demand (Smil 2010; Malm 2016). Industrial renewables are none of these things. They are spatiotemporally contingent: harvested by the square meter, consumed as they are generated, and reliant on new transmission and storage to bridge the distance between a windy night and a still afternoon (Fischer-Kowalski and Haberl 2007). The move from the one to the other is therefore not a simple swap of fuels but a reorganization of energy's spatial form. From the vertical regime of centralized extraction toward a horizontal one

of decentralized, land-hungry generation (Bridge et al. 2013), decarbonization does not shrink the footprint of energy, it spreads it out.

That spreading-out carries a political corollary the engineering literature tends to pass over. Because renewable facilities are more numerous, larger in their terrestrial footprint, and sited nearer the communities they serve, the transition draws far more places into the work of hosting power. Both the number and the size of potential conflict sites multiply. Moreover, wherever it happens, siting is never the purely technocratic exercise its proponents describe. It is a contest among unequally-resourced actors: developers with capital and counsel; states balancing climate ambition against local opinion; and host communities whose capacity to refuse an unwanted facility varies enormously (Walsh 1981; Carley et al. 2020). Where that capacity is thin, the facility tends to land (Mohai and Saha 2015). This is the durable finding of environmental-justice scholarship, built largely on the study of coal: communities of color and low-income residents have borne disproportionate proximity to the dirtiest infrastructure of the fossil era (Bullard 1990; Mohai and Bryant 1992; Cushing et al. 2023).

That siting contest is not restricted to the fossil fleet; it extends now to the buildout of renewables. *What* communities resist, though, has changed in kind: a wind farm or a solar field is not a smokestack, and the burden it imposes is land taken, viewsheds altered, and local control surrendered rather than air and water fouled — a 'locally unwanted land use' whose costs are real but categorically distinct (and objectively less burdensome) from the emissions of a fossil plant (Mueller and Brooks 2020). Renewable energy enjoys broad support in the abstract and meets fierce resistance in the particular. The proposed Gemini Solar Project in Nevada, poised to be the country's largest, drew sustained local opposition despite the economic benefits it promised its rural hosts (Milman 2022). In California's Mojave, residents describe utility-scale solar as turning a desert 'oasis' into a 'dead sea' (Wainwright 2023). By the end of 2024, at least 459 counties and municipalities across 44 states had adopted severe restrictions on siting renewables (a 16% rise in a single year), with 498 projects actively contested across 49 states (Eisenson et al. 2025). The political response has been to move the decision upward: Illinois, Michigan, and Minnesota have stripped localities of the power to bar wind and solar, routing approval through state commissions over local objection. In one Michigan township, the approval overrode residents who had voted the project down and recalled the officials

who backed it (Meadows 2024). That reallocation of authority has itself become the contest, litigated to the steps of the state supreme court by townships arguing not against renewables as such but for 'a meaningful voice' in where they land ("Appeals Court Issues Split Ruling in Michigan Solar Permitting Suit: A State Appeals Panel Has Upheld Most of Michigan's Regulations Limiting Local Control over Renewable Energy Projects" 2026; Hornier 2026). This is the transition's distinctive political form: as the land take grows toward the tens of millions of acres a decarbonized grid will require (Milman and Witherspoon 2024), the locus of siting authority migrates from the town hall to the statehouse, and with it the terrain on which communities must mount whatever resistance they can.

Does the green transition unwind its inheritance of environmental injustice, or reproduce it? The honest answer is that we do not yet know, because the question has seldom been put to the whole fleet at once. The evidence is fossil-shaped and fuel-by-fuel: coal hosts communities that are 'Whiter but poorer' (Cranmer et al. 2023); wind clusters in rural, predominantly white counties, sorting along axes of age, education, and rurality more than race (Mueller and Brooks 2020); utility-scale solar lands in places that are less white but economically bifurcated (Cranmer et al. 2023); and the determinants of renewable siting are, on the best available evidence, overwhelmingly technoeconomic (resource quality and open land), with demography a residual of that screen, the variation left after viability has done the sorting (O'Shaughnessy et al. 2023). That there are distinctive energy-specific geographic patterns is apparent, in short, but it is not yet clear how these patterns relate to one another, let alone the broader contextual patterns of states and the nation at large.

Drawing on a plant-year panel of the full continental US generation fleet, fossil, nuclear, and renewable alike, observed annually from 2010 to 2023, this chapter attempts a documentation of how the demographic composition of each energy type's host communities evolves across the decade of the transition. It decomposes that change into its parts: **fleet turnover**, the *siting of new capacity* and the *retirement of old*, set against the slower drift of the communities around **plants that simply persist**. These are the two canonical mechanisms of environmental inequality formation (Pellow 2000) — disparate siting versus post-siting 'sorting,' the long-running 'which came first' question (Pastor, Sadd, and Hipp 2001; Been and Gupta 1997; McCoy 2017). The contribution

is descriptive, and deliberately so. The aim is not to identify the cause of any disparity, which demographic data measured in overlapping five-year averages cannot credibly do (Banzhaf, Ma, and Timmins 2019), but to establish the stylized facts that any causal account would have to explain: whether the transition's demographic map is being redrawn chiefly by where the new plants go or by who lives around the old ones, and whether that answer differs, as the fuel-specific evidence implies it must, across technologies.

Three choices undergird the design. First, it is plant-centric: community demographics are assigned to each facility by population-weighted apportionment within distance buffers — the linkage the measurement literature endorses over cruder host-unit comparison (Chakraborty, Maantay, and Brender 2011; Mohai and Saha 2006). Second, it does not collapse generation into a single 'hazard.' The burdens differ in kind: an emitting plant imposes a pollution externality that a wind farm does not, so the question worth asking is not whether 'energy' tracks demography but whether 'clean' and 'dirty' capacity come to rest in different communities. Third, it carries race and class together: percent nonwhite, Black, and Hispanic alongside poverty and median income, entered separately (Ringquist 2005).

The stakes are not only scholarly. Suppose the transition's burdens are being relocated rather than retired, the clean buildout flowing toward communities the fossil system spared, the emitting remnant concentrating among those it did not. Decarbonization would then carry environmental inequality forward under a 'green' banner even as its aggregate harms decline (Carley and Konisky 2020; Richmond-Bryant et al. 2020). Whether it does, and by which channel, is what this chapter sets out to establish.

4.2 Background and Theoretical Framing

4.2.1 Environmental inequality formation: power resources and the path of least resistance

That the burdens of environmental hazard fall unequally, and in the main along lines of race and class, is the founding proposition of the environmental-justice tradition and, by now, less a contested finding than a premise upon which subsequent work is built (Bullard 1990; Mohai and Bryant 1992;

Ringquist 2005). The question this chapter takes up is not whether such disparity exists but how it forms — and, in a generating fleet being rapidly remade, how it is redrawn. The most useful point of departure is therefore not the distributive snapshot but the literature on environmental inequality *formation*, which treats a demographic disparity not as a static fact to be measured but as an outcome continuously produced and reproduced by the unequal distribution of social power (Pellow 2000, 2016). On this account a disparity is best understood as what Downey (2005) terms disparate exposure: a concentration of harm among particular groups exceeding what a spatially random distribution would yield, produced by hierarchies that come to rest the harm where the capacity to resist it is weakest.

What governs where that harm settles is, on the tradition's dominant account, a political contest rather than a technical one. Facility siting pits three sets of actors of markedly unequal endowment against one another — developers seeking to minimize cost and to secure some measure of social acceptance for the facility, the state agencies that must weigh a public interest against the opposition a siting provokes, and the host communities that will bear the burden and that differ enormously in their capacity to refuse it (Greenberg 2017; R. S. Liévanos, Greenberg, and Wishart 2018). That capacity, comprising the economic resources, political access, social capital, and organizational structures and reach a community can mobilize against an unwanted facility (Walsh 1981; McAdam and Boudet 2012), is what the tradition calls a community's power resources, and the prediction that follows is familiar: other things equal, the burden follows the path of least resistance, settling where those resources are least (Mohai and Saha 2015). What makes that capacity so often decisive is a structural feature of energy-facility politics the social-acceptance literature has named the 'social gap': the routine divergence between a diffuse public that supports generation in the abstract and the concentrated opposition a specific proposal calls forth in the particular (Wüstenhagen, Wolsink, and Bürer 2007). That divergence throws a facility's fate onto exactly the local arena in which power resources are most unevenly held. Two corollaries organize the analysis below. The first is that the demographic composition of the communities hosting generation is not the product of chance, and so is a quantity worth explaining in its own right. The second is that the same logic, run forward, predicts that *new* capacity, like old, will come to rest where resistance is weakest. For renewable capacity that prediction requires qualification, since siting there is governed less by the

host community's resistance than by technoeconomic constraint; the section on material profiles takes this up.

4.2.2 Siting, sorting, and exit: the formation and redrawing of disparity

If the path of least resistance explains where burdens settle, it does not by itself explain how a community came to be disadvantaged in the first place — whether the hazard was sited in a community already poor and nonwhite, or whether the community assumed that composition only after the hazard arrived. This is the tradition's long-running 'which came first' question, and the distinction it draws is between disparate *siting* and post-siting *sorting* (Been and Gupta 1997; Pastor, Sadd, and Hipp 2001; Mohai and Saha 2015). The two are distinct social processes. Siting is the political contest just described; sorting is a matter of residential mobility, the slower rearrangement of households around a fixed facility — driven, on one account, by the Tiebout logic of residents 'voting with their feet' toward the cheaper land that proximity to a hazard affords (Banzhaf and Walsh 2008), and constrained, on another, by the discriminatory housing markets that leave even comparatively affluent minority households more exposed than poorer white ones (Crowder and Downey 2010).

That binary, however, was built upon a fleet of 'hazards' that only ever grew or held constant — the accreting stock of waste facilities and fossil plants on which the canonical studies were conducted. The energy transition violates that premise. Here the fleet of 'hazards' is being remade from both ends at once: new capacity sited on one margin and, on the other, the wholesale *exit* of the incumbent fleet as coal-fired plants are retired across the decade. Exit is not a special case of siting or of sorting; it is a third channel through which the demographic map is redrawn, and one for which the inherited binary has no category. The analysis therefore performs a ternary decomposition rather than a binary one: siting, sorting, and exit. The exit channel carries its own theoretical freight, drawn from the literature on fossil decline and just transition. The retirement of the oldest and dirtiest units, even as it lowers emissions in the aggregate, entails distributional consequences of its own: shifting the residual pollution burden across racial and economic lines (Richmond-Bryant et al. 2020), reordering the map of advantage and disadvantage the fossil era bequeathed (Carley and Konisky 2020; Raimi, Carley, and Konisky 2022), and measurably altering conditions in the

neighborhoods the retired plants leave behind (Casey et al. 2020). Within this chapter's relative, plant-side accounting, though, the exit channel measures something narrower than any redistribution of exposure — that absolute-burden question would require a tract-side design. The exit channel measures the *demographic signature of departure*: the degree to which the communities around the coal that retires differ from those around the coal that survives, and so the extent to which a fuel's mean host composition moves by subtraction alone. That coal's demographic change should be exit-driven and renewables' entry-driven is, on its face, an accounting artifact: coal is what retires, renewables what gets built. Extending the 'which came first' debate to a fleet shrinking on one end and growing on the other is what gives that artifact content: the political-economy account developed above predicts not merely that each fuel moves through the channel available to it, but that turnover should outweigh the drift of standing communities, and in a direction each fuel's political economy anticipates. However, the decomposition cannot adjudicate cause. Which of the three channels *produced* a community's disadvantage is a question that demographic data measured in overlapping five-year averages cannot answer (Banzhaf, Ma, and Timmins 2019). The analysis accordingly reads its decomposition as a descriptive accounting of the transition's moving geography, not as a causal test of its formation.

4.2.3 Material profiles and the limits of least resistance: why the transition's geography is fuel-specific

Why the three channels should sort communities differently by fuel returns the argument to the material profile of generation with which the chapter began. The distinction that matters is not 'energy' against its absence but the kind of burden a given technology imposes. A fossil plant's burden is an emissions externality: particulate matter, sulfur dioxide, and nitrogen oxides concentrated on the communities nearest the stack and downwind of it. Its siting has historically followed the path of least resistance toward precisely the race-and-class disparity the tradition documents. A wind farm or a solar field imposes no such externality. Its burden is a 'locally unwanted land use,' categorically distinct in kind from a pollution externality (Mueller and Brooks 2020). Moreover, its siting is governed far less by the political resistance of the host than by technoeconomic constraint, by the resource quality, open land, and transmission access that determine where a project is viable at all

(O'Shaughnessy et al. 2023). The claim is about which filter binds first, not that resistance never prevails: organized opposition can still delay or defeat an individual project (Eisenson et al. 2025). The path-of-least-resistance prediction is, for renewables, both weakened and redirected: weakened because demography enters only conditionally, after the technoeconomic filter has done its work; and redirected because land availability selects for rurality and the economic profile that comes with it, not for the racial composition on which fossil siting sorted. Demographic composition is, throughout, a proxy for resistance capacity rather than resistance itself: the tradition inherited the practice of reading sociopolitical capacity off census measures (Mohai and Bryant 1992), while newer work measures opposition directly, in ordinances passed, permits contested, and projects delayed (L. C. Stokes et al. 2023). The two need not coincide, and the demographic measures used here should be read as composition, not as a direct measure of mobilization. That land is not a blank surface: utility-scale wind and solar are land-intensive, so their entry is mediated by the existing political economy of rural land — its ownership and concentration, and its competing agricultural uses — and who hosts the buildout is filtered through who holds the land and on what terms (Knuth et al. 2022; Sareen et al. 2025). One demographic axis, however, resists this smoothing into rurality and class. The siting of energy infrastructure on and adjacent to tribal land (from the uranium and coal of the fossil century to the utility-scale solar and wind now proposed for the public lands of the interior West) is constructed by a distinctive, settler-colonial, history of its own. Dispossession reduced the Indigenous land base within the present-day United States by nearly 99 percent and relocated tribes to lands more exposed to climate hazards and less likely to hold subsurface oil and gas (Farrell et al. 2021). The resultant modern geography is thus one where proximity to Native land operates less as a continuous gradient than as a categorical condition. For that reason the analysis models Indigenous proximity as a threshold rather than a smooth covariate.

It would be a mistake, however, to read the recession of host-community resistance as the recession of politics as such. The political economy of the clean transition is not absent from renewable siting but *reorganized* — displaced from the local contest over a single facility onto a multi-scalar terrain of distributive conflict (Colgan, Green, and Hale 2021; Breetz, Mildenberger, and Stokes 2018; Aklin and Mildenberger 2020). Clean energy threatens the asset values of carbon-intensive incumbents, who have proven both willing and able to impose sustained obstruction on its deployment, through the

legislatures, regulators, and courts in which energy policy is produced (Meckling 2015; Mildenberger 2020). The conflict described in the introduction is one front in this larger struggle: the migration of siting authority from the town hall to the statehouse, as municipalities, counties, states, and the federal government contest the terms on which renewables are built. It is a politics of energy federalism whose decisive battles are waged at levels well above the host community itself (D. Stokes 2021; L. C. Stokes et al. 2023). The implication for the demographic question is paradoxical: the political economy of renewable siting is, if anything, more intense than that of the fossil fleet it displaces, but its principal axis of conflict (incumbent against entrant, one order of government against another) increasingly cuts across the axis of host-community race and class along which the fossil disparity was inscribed. That a renewable project's fate should depend less on the resistance capacity of its immediate neighbors than on a contest waged largely elsewhere is itself a reason to expect the fossil-era race gradient to attenuate even as the politics surrounding the transition grows fiercer.

This tension, between a political economy of siting that would have renewables reproduce the fossil pattern and the technoeconomic and supra-local political logics that would break it, yields three competing predictions, which structure the chapter's formal 'descriptive expectations' (as opposed to rigorously tested hypotheses) set out below. Under *replication*, the political economy dominates. The resistance-capacity mechanism operates on clean capacity as it did on dirty, and renewable hosts come to resemble fossil hosts in their racial and economic disadvantage. Under *transformation*, the technoeconomic and land-use logics dominate: the demographic associations weaken, or shift their axis from race toward rurality and class — the prior this chapter holds for wind in particular. Under the *null*, resource potential swamps demography altogether, and the siting of renewables bears no systematic relation to the composition of the communities that host them (O'Shaughnessy et al. 2023). The three are not mutually exclusive across fuels; the most likely outcome, on the evidence assembled here, is that different technologies fall under different predictions, and it is this expectation of fuel-specific divergence that the analysis is built to examine.

4.2.4 From a static cross-section to a trajectory: positioning and scope

The nearest precedent for the question, so framed, is Cranmer et al. (2023), whose national cross-section established that the host communities of different fuels diverge in composition — coal and wind 'Whiter but poorer,' gas and solar less white but economically bifurcated. That study furnishes this chapter's point of departure and the bulk of its priors. What it offers is a *snapshot*, and what this chapter adds is a *trajectory* — the same fuel-by-fuel contrast observed annually across the transition decade and decomposed into the channels that move it, the exit channel among them, which a fleet being retired from one end makes available and which the inherited siting-versus-sorting frame has no room for. Two scope conditions follow from the design. First, the analysis is plant-side and relative: it characterizes the demographic composition of the communities surrounding generators, and the divergence of that composition across fuels, rather than the absolute pollution burden borne by any population — the relative-versus-absolute distinction that Richmond-Bryant et al. (2020) illustrates, and the one that separates this chapter's question from the exposure question a tract-side design would pose. Second, it is associational: the causal siting-versus-sorting test that a credible identification strategy would require (Clay et al. 2025) is outside the scope of this analysis. The chapter's aim is accordingly descriptive: to document how the spatial demography of American electricity generation is redrawn as the system decarbonizes (building new gas and industrial renewables even as it retires coal), and by which of the three channels. It is to the measurement of that geography that the chapter now turns.

4.2.5 Research questions

- **RQ1 (trajectory).** How do the demographic compositions of energy-type host communities evolve over 2010–2023, net of the national trend?
- **RQ2 (decomposition).** How much of each fuel's demographic change is *fleet turnover* (entry/exit) versus *change around persisting plants*?
- **RQ3 (entrant composition).** How does the community composition surrounding newly built capacity differ from that surrounding incumbent capacity, by fuel, and does that difference arise within states or between them?

4.3 Data and Methods

4.3.1 Data sources

The analysis draws on the public datasets listed in Table 4.1, each freely available from the provider named; the table records the version or coverage used and the role each plays in the plant-year panel (EIA 2025; Bureau 2026a, 2026b; Rand et al. 2020; Fujita et al. 2026; Bank 2026).

Table 4.1: Data sources for the plant-year panel.

Dataset	Provider	Version / coverage	Role in the analysis
Form EIA-860 (Annual Electric Generator Report)	U.S. Energy Information Administration	Annual, 2008–2023	Plant and generator fleet spine; energy type; coordinates; operating and entry/retirement status
American Community Survey, 5-year estimates	U.S. Census Bureau	End-anchored 5-year windows, 2006–2010 through 2019–2023	Community race, ethnicity, poverty, and median household income (buffer outcomes)
TIGER/Line shapefiles	U.S. Census Bureau	2010 and 2020 tract vintages	Census-tract geography for population-weighted buffer apportionment
U.S. Wind Turbine Database	USGS, LBNL & ACP	v9.0 (2026)	Turbine coordinates for the wind-array measurement check
U.S. Large-Scale Solar Photovoltaic Database	USGS & LBNL	v4.0 (2026)	Solar facility footprints for the solar-footprint check
Consumer Price Index (US)	World Bank	Annual, 2010–2023	Deflation of median household income to constant 2023 dollars

Note. The American Community Survey window is end-anchored: each annual observation uses the five-year estimate ending in that year. Median household income is deflated to constant 2023 dollars using the Consumer Price Index.

4.3.2 Sample and unit of analysis

The analysis covers the full U.S. electricity-generation fleet, renewables included, as an annual panel of 112,782 plant-years (13,063 plants, 2010–2023). The unit is the plant-year, and the panel is annual rather than periodic — a departure from the snapshot design of the dissertation's earlier chapters that the question itself compels, since a trajectory cannot be read from a handful of points spaced years

apart. The spine is the U.S. Energy Information Administration's Form 860, an annual census of the nation's electric generators. Two of its properties recommend it here: its annual resolution, and a coverage of the utility-scale renewable fleet more complete than that of the emissions-reporting sources on which the earlier chapters leaned. The panel opens in 2010. The year 2009 is dropped not for want of generation data but for want of community data: the American Community Survey's five-year estimates reach back only so far, and their earliest usable window rests on tract boundaries drawn to the 2000 Census — a vintage the transition decade has since redrawn.

A plant is in the panel in the years it is operable and absent in the years it is not: it is observed across its active span, from the first year it reports an operable generator to the last, so that a plant idled for a season and returned to service is retained, while the years before it is built and after it is retired are not. The decomposition and the entrant-composition analysis each classify a plant by its place in the fleet's turnover, but they do so differently. The decomposition compares a plant's presence at the two years it spans: a *survivor* is present in both, an *entrant* in the later but not the earlier, an *exiter* the reverse. The entrant-composition analysis asks instead when a plant was built, dating it by the start year of its oldest generator: there an *entrant* is one built within the panel's span, an *incumbent* one built before it. The two senses of 'entrant' usually coincide, but not always: a plant built before 2010 that only returns to service within the panel is an entrant to the decomposition yet an incumbent to the entrant-composition analysis. The text marks which sense is meant where the distinction matters.

4.3.3 Community demographic measurement

Community composition is captured by seven measures of the population surrounding each plant, drawn from three ACS tables: B03002 (Hispanic or Latino Origin by Race), B17001 (Poverty Status in the Past 12 Months by Sex by Age), and B19013 (Median Household Income in the Past 12 Months). Five are of race and ethnicity, all built from B03002's mutually exclusive Hispanic-origin-by-race cells: the share nonwhite (the population that is not 'White alone, not Hispanic or Latino'), the non-Hispanic Black and non-Hispanic Asian shares, the Hispanic share (Hispanic or Latino of any race), and a binary marker for proximity to a substantially Native community (a surrounding population at least five, and alternatively ten, percent non-Hispanic American Indian or

Alaska Native). Because each single-race share is the ACS's not-Hispanic, single-race cell, the plain names used hereafter ('Black,' 'Asian,' 'American Indian or Alaska Native') all denote non-Hispanic categories; only the Hispanic share crosses race. Two are of class: the poverty rate (the share of the population for whom poverty status is determined whose income falls below the poverty threshold) and median household income, the latter deflated to constant dollars so that change across the decade registers real fortunes rather than inflation. Race and class enter separately throughout, the better to ask whether any racial gradient survives the conditioning of income (Mohai and Bryant 1992). Each measure is assigned to a plant by population-weighted areal apportionment: tract-level estimates from the American Community Survey are mapped to circular buffers of three, five, and ten kilometres around the plant (three kilometres the primary radius), and every intersecting tract is weighted by the share of its population falling inside the ring (Chakraborty, Maantay, and Brender 2011; Mohai and Saha 2015).

Each measure takes the functional form its distribution invites. The four shares of race and ethnicity are bounded proportions and enter linearly, a form the descriptive decomposition in any case requires, since it partitions a *mean*. Percent nonwhite is spread across the full range (skewness 0.89), the rarer shares more compressed about zero. The Indigenous measure is a case apart: near-zero for the overwhelming majority of plants and substantial only in the neighborhood of tribal lands, it is no smooth gradient, and it enters as the binary marker already described. Median household income, sharply right-skewed in levels (skewness 1.45), is deflated to constant dollars and modelled in logs, which restores near-symmetry (skewness 0.2) and lets its coefficient read as a proportional difference. Tertile cuts of the continuous measures, and a majority-minority threshold for percent nonwhite, are retained as robustness on the linear form.

This procedure presumes that a generating facility can be fixed by a single point — defensible for the compact infrastructure of the fossil era (Bridge et al. (2013)). Renewable generation is not compact but expansive and rambling; in a way, harking back to the pre-fossil age (cf., Fischer-Kowalski and Haberl (2007)). In other words, the geography of renewables is 'horizontal': because sun and wind arrive as a diffuse flow, harvesting them at scale means spreading capacity across the land rather than stacking it at a point — a wind 'facility' is dozens of turbines scattered over tens of square kilometres, a solar 'facility' a contiguous field of panels. The contrast is not rhetorical but

measurable, in the power density a technology wrings from its land (Figure 4.1): fossil generation out-densifies renewable generation by one to three orders of magnitude, so the single coordinate that faithfully locates the former will badly misplace the latter (Smil 2010).

Fossil generation is one to three orders of magnitude denser than renewable generation (Figure 4.1): it stacks capacity vertically at a point, where renewables spread it horizontally across the land.

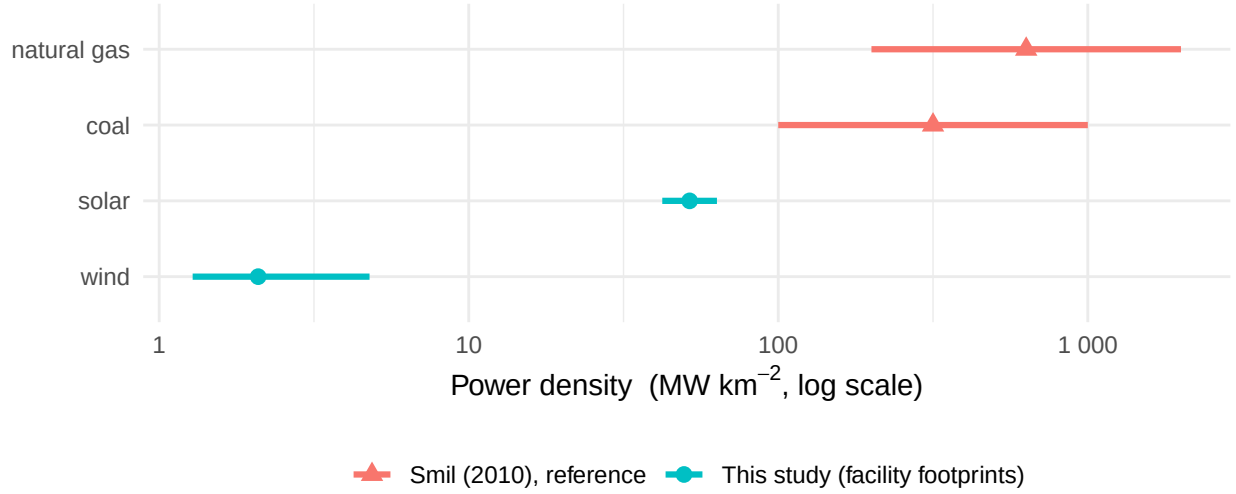

Figure 4.1: Power density, generating capacity per unit of land, by energy type, on a log scale. Wind and solar are computed from facility footprints (the U.S. Wind Turbine and Large-Scale Photovoltaic Databases; point = median, bar = interquartile range); coal and natural gas are reference ranges from Smil (2010).

Two consequences follow, and the analysis addresses each. A point that wanders contaminates any study of change over time: EIA-860's reported coordinates drift across years for a non-trivial share of plants (most of all the dispersed renewables), and an unpinned buffer would read that drift as neighborhood change; each plant is therefore fixed to the median of its observed coordinates before buffering. And a point may simply be a poor proxy for a spread-out facility; yet re-measuring wind installations as the union of buffers around their constituent turbines leaves the aggregate community composition essentially unchanged — the array reaches more people, but not systematically different ones — while solar's compact footprint sits comfortably within even the smallest buffer. Both checks are reported in Appendix Section 4.9.2.

4.3.4 Energy-type classification

Each plant-year is assigned a single energy type. A plant need not be a pure specimen of one technology: a site may pair gas turbines with a solar array, or keep an idle coal unit beside the gas that replaced it. The classification therefore rests on a capacity-weighted rule: a plant's energy type for a given year is the fuel category commanding the most operable nameplate capacity among its generators, with ties broken alphabetically and the type left undefined in the years a plant carries no operable generator at all. In practice the rule rarely has much to arbitrate: only about 5% of plant-years carry operable generators of more than one fuel category, and at the remainder a single fuel commands the entire nameplate.[1] The rule yields eleven types (coal, oil, gas, hydro, biomass, geothermal, nuclear, solar, wind, storage, and a residual 'other'), and these enter the analysis as eleven distinct fuels, since the fuel-specificity of the transition's geography is the very thing under examination. Coal serves as the reference category throughout, the legacy baseline against which the rest of the fleet is read. The least populous types (geothermal, storage, the residual 'other') are carried for completeness but reported with the caution their thin numbers warrant, and the figures that follow lead with the fuels of substantial weight.

4.3.5 Analytic strategy

The three research questions are addressed via three descriptive analyses. The first traces *trajectories*: the mean host-community composition of each energy type, computed across plants and plotted year by year over the panel, so that the divergence of one fuel's communities from another's becomes visible as a moving picture rather than a single frame. The plant, not the megawatt, is the unit of the average: each facility counts once, as one observation. The analysis makes the consequence of that choice explicit rather than burying it, setting the plant-weighted account against a population-weighted one so that the difference between counting plants and counting people can be read directly (the Discussion takes it up). And because the national population is itself diversifying across the decade, every trajectory is read net of that secular trend rather than in absolute level. To net out that trend, each fuel's yearly mean is compared with the demographic composition of the continental

[1] The commonest such pairings are dual-fuel fossil sites (gas with oil-fired startup or backup capacity, and legacy coal beside gas or oil), and, increasingly, the solar-and-storage hybrids of the transition era.

United States in the same year, population-weighted over the same five-year survey windows. The benchmark is deliberately external to the generating fleet: a baseline built from the fleet's own composition would itself shift as the buildout reweights the fleet toward the communities it favors, and would absorb part of the very geography the trajectory is meant to display. What the figure shows is therefore each fuel's divergence from the nation, not its raw level nor its position relative to a moving fleet average (which is reported only as a robustness). For fuel k in year t, the plotted quantity is that fuel-year mean net of the national benchmark, Equation 4.1,

$$\text{net}_{kt} \;=\; \bar{y}_{kt} - \bar{y}_t^{\text{US}}, \tag{4.1}$$

with $\bar{y}_{kt}$ the unweighted mean of the demographic measure over plants of type k operating in year t, and $\bar{y}_t^{\text{US}}$ the contemporaneous population-weighted national value.

The second analysis *decomposes* the shift each trajectory records into the fleet dynamics that produce it. It operates on raw host composition, not the net-of-national series the trajectory plots, so the common national diversification is not differenced away but surfaces within the survivor term. The research design sets out a dynamic productivity decomposition (Melitz and Polanec (2015)) to account separately for surviving, entering, and exiting units, effectively partitioning a fuel's change in mean host composition between two endpoints into three additive parts: the change among the communities of surviving plants (those present throughout), set against the contributions of entry and exit, the fleet turnover by which new plants enter the average and retired ones leave it. For a fuel k, writing S, X, and E for its survivors, exiters, and entrants across the endpoints t_0 and t_1, that change partitions exactly as Equation 4.2,

$$\Delta\bar{y}_k = \underbrace{(\bar{y}_{k,S}^{\,t_1} - \bar{y}_{k,S}^{\,t_0})}_{\text{within-survivor}} + \underbrace{\tfrac{n_{k,E}}{N_{k,1}}(\bar{y}_{k,E}^{\,t_1} - \bar{y}_{k,S}^{\,t_1})}_{\text{entry}} + \underbrace{\tfrac{n_{k,X}}{N_{k,0}}(\bar{y}_{k,S}^{\,t_0} - \bar{y}_{k,X}^{\,t_0})}_{\text{exit}}, \tag{4.2}$$

where $\bar{y}_{k,\cdot}$ denotes the mean of the demographic measure over the indicated subset at the indicated endpoint, and $N_{k,0} = n_{k,S} + n_{k,X}$, $N_{k,1} = n_{k,S} + n_{k,E}$. The partition is what gives the second research question a degree of empirical purchase, distinguishing a demographic map redrawn by *where the new plants go* from one redrawn by *who comes to live around the old ones*. The design is specified so as to expose the fuel-specific shape the theory anticipates — coal's change running

largely through the exit term, renewables' through entry, with the survivor term, most exposed to the survey smoothing discussed below, carrying what remains. That survivor term is also the one quantity sensitive to the redrawing of Census tract boundaries between their 2010 and 2020 vintages, since it alone differences a community against its own earlier self; the decomposition is therefore anchored on the single-vintage 2010–2019 window, with the full 2010–2023 span carried as an extension.

The third analysis turns to *entrant composition*, asking where new capacity comes to rest. Host demographics are regressed on the interaction of entrant status with energy type — incumbents present at the panel's outset furnishing the baseline, coal the reference fuel, race and class entered as separate outcomes so that any racial gradient is read net of the economic one, and standard errors clustered by state. For plant i observed in year t, with demographic outcome y, this is Equation 4.3,

$$y_{it} = \sum_{k} \alpha_k \, \mathbf{1}[\text{type}_i = k] + \sum_{k} \beta_k \, \mathbf{1}[\text{type}_i = k] \cdot \text{entrant}_i + \gamma_t + \delta_s + \varepsilon_{it}, \tag{4.3}$$

in which β_k is fuel k's entrant-versus-incumbent gap, γ_t are year effects, and δ_s are state effects (carried only in the within specification). The coefficients describe how the communities receiving *new* plants of each type differ from those that hosted the incumbent fleet — the entrant margin of the transition's changing geography, estimated as association and not as effect. The regression is run in two forms, and the distance between them is the object of interest. The first estimates each fuel's entrant-versus-incumbent gap across the country as a whole; the second adds state fixed effects, so that the comparison is made only within states. The within-state estimate is the part of the gap arising among communities within the same state (a coarse container for the metropolitan, county, and regional scales at which segregation is typically organized (Raoul S. Liévanos 2015)), new plants landing in systematically whiter or more nonwhite communities than the incumbents *in the same state*; the difference between the two, the between-state component, is the part explained by which states the new capacity is built in rather than which communities within them. The split distinguishes two processes that would otherwise leave the same fleet-wide footprint: a within-place sorting of technology onto community, or a relocation of the fleet across the map. For each fuel, that same distinction separates replication from transformation, a racial gradient arising within states from one drawn between them.

Two cautions bound what that split can claim. Because the demographic shift proves to live between states rather than within them, and because the estimator clusters on state, the standard errors on the entrant gaps are wide; the analysis makes no claim to statistical significance, reads the point estimates descriptively, and rests their magnitude on the decomposition's cleaner accounting. And locating the gradient at the scale of the state is not the same as explaining it: the between-state component documents *that* new capacity is built in whiter or more nonwhite states than the incumbent fleet, but the technoeconomic determinants that would explain why (the quality of the wind and solar resource, the proximity of transmission, the price of land) lie outside this project's scope, since they are not comparable across technologies; the relocation is therefore measured rather than accounted for. That finer separation, and with it any causal claim about where capacity is drawn net of where the resource and the grid permit it to be built at all, would require the resource and transmission data of a causal siting design, and is outside the scope of this analysis.

4.3.6 Identification and interpretation

Three features of the design fix what the analysis can and cannot claim. The first is that the estimand is associational throughout. The community measures derive from the American Community Survey's five-year estimates, and adjacent five-year windows share four of their five years, so the demographic series is heavily smoothed and no clean temporal lag can be cut from it — none that would set a community's composition *before* a plant arrived against its composition after. The entrant analysis accordingly compares the communities of new and incumbent plants as they are contemporaneously observed, not a community against its own earlier self; entry is finely dated, but the demographics it is matched to are not, and the design declines to regress a sharply timed event upon a smeared outcome as though the two kept a common clock. What the analysis estimates is therefore the demographic *context* in which the transition's plants are found, and the divergence of that context across fuels — not the causal effect of a plant upon the people who live near it.

The second feature follows from the first. Every quantity the analysis reports is read as a divergence *net of* the secular national trend rather than as a raw level, but the netting takes two forms suited to the two kinds of estimate. The trajectory differences each fuel's mean against the contemporaneous composition of the continental United States, population-weighted over the same

five-year windows so that the common survey smoothing cancels on both sides of the difference — a national benchmark held deliberately external to the generating fleet, lest a baseline drawn from the fleet's own evolving composition absorb the very transition the trajectory means to display. The entrant-composition regression, by contrast, carries year fixed effects throughout, which absorb the secular trend and the overlapping-window smoothing artefact common to every plant within a given year. A further seam runs through the community data themselves. The survey publishes its estimates on exactly two tract-boundary vintages across the panel: the 2010 Census geography through the 2015–2019 window, and the 2020 geography from the 2016–2020 window onward (Schroeder 2007). The redrawing at that break is substantial, the two vintages sharing 61,638 tract identifiers of the 95,832 in their union, with 11,418 retired and 22,776 newly created, so a plant observed across the full panel is read first on one partition of its surroundings and later on another. (A geonormalized alternative, tract series standardized to a single boundary vintage, was set aside because the standardized products carry only bare demographic counts, without the income and poverty measures the analysis requires alongside race.) Because the buffer apportionment integrates over whoever falls within a fixed radius, the re-tracting leaves the cross-sectional trajectory and the within-year entrant comparison essentially undisturbed; it is only the within-survivor term of the decomposition, which sets a community against its own earlier self, that the seam can move.

How much it moves was measured directly. At the 2019-to-2020 step, and only there, the year-over-year dispersion of buffer composition roughly doubles — for percent nonwhite at 3 km, a standard deviation of change near 0.040 against a 0.020–0.023 baseline in the nine preceding steps — while the mean change sits on the trend line between adjacent years. That is the signature of a partition change rather than of demographic change: re-tracting scatters estimates without moving their center. The seam therefore enters the decomposition as inflated variance in the within-survivor term, not as a biased level. One caution attends the reading: the 2016–2020 window spans the pandemic, so the 2020 step is confounded, and it is the mean-versus-dispersion contrast, not the timing, that separates artifact from real change.

The third feature is the scope condition already stated in the theory and repeated here for the record: the analysis is plant-side and relative. It characterizes the composition of the communities surrounding generators, not the absolute pollution burden those communities bear; the absolute-

exposure question, and the causal siting-versus-sorting test that would demand an identification strategy this design does not possess, are, by design, outside the scope of this analysis.

4.4 Expectations

Because the analysis is descriptive and its estimand associational, the expectations set out here are not hypotheses in the testing sense — there is no null to reject, no threshold to clear. They are directional priors, fixed before estimation in the manner of a pre-registration, and their purpose is to give the results something to be held against: to mark where the fuel-specific pattern conforms to what the theory anticipates and, more instructively, where it does not. They follow the structure the theory has already laid down: the master frame of inequality formation; the ternary anatomy of siting, sorting, and exit; and the trichotomy of replication, transformation, and null. And they are fuel-specific, since the theory rules out any single gradient running uniformly across the whole fleet.

4.4.1 E1: Trajectory (RQ1)

Host-community composition is neither random (this is the inequality-formation frame) nor uniform across fuels: the static cross-section that precedes this study finds coal and wind hosts Whiter but poorer, and oil, gas, and solar hosts less white and bifurcated on class (Cranmer et al. 2023). The first expectation is that these contrasts persist across the panel and that, read net of the national trend, the host communities of the several fuels *diverge* over the transition decade rather than converge. The expectation is pointedly not that a single gap widens; it is that the fuels travel different demographic paths, and the question worth putting to the trajectories is which fuels separate from the fleet mean, in which direction, and on which dimension — race or class. The Indigenous dimension is expected to behave unlike the rest, concentrating around particular technologies and locations rather than varying smoothly, which is the reason it is carried as a threshold rather than a gradient.

4.4.2 E2: Decomposition (RQ2)

The ternary anatomy of siting, sorting, and exit yields a sharper prior about *how* each fuel's trajectory moves. The expectation is of a fuel-specific shape. Coal's demographic change should run largely

through the *exit* term, since the retirements of the decade are overwhelmingly coal and the communities a retiring plant vacates need not resemble those its surviving peers retain (Richmond-Bryant et al. 2020; Casey et al. 2020). Renewables' change should run almost wholly through *entry*, the siting of new capacity, since they retire at near-zero rates across the window. Gas, building and retiring at once, should split its change between the two turnover terms. The *within-survivor* component, the slow drift of communities around plants that neither enter nor leave, is expected to be the smallest of the three and the least certain, for it is the quantity the survey's overlapping windows smooth most heavily, and the one the tract-vintage seam can disturb. That the demographic map of the transition should be redrawn by *fleet turnover* for some fuels and, if at all, by *neighborhood change* for others is itself the expectation; the asymmetry, not any single headline figure, is what the decomposition is built to surface.

Beyond which channel dominates, the design commits in advance to the *direction* each channel runs. On the exit margin (in practice coal, the one fuel retiring at scale), the prior is that the less-disadvantaged plants retire first: the host communities of the coal that exits are expected to be Whiter, higher in income, and less poor than those of the coal that survives, so that exit, by subtraction alone, *steepens* the disadvantage gradient across the coal fleet that remains (Richmond-Bryant et al. 2020). On the within-survivor margin the priors are deliberately modest and descriptive: around persisting legacy-fossil plants, communities are expected to drift weakly toward greater disadvantage (marginally more nonwhite and poorer, net of the national trend, in keeping with the sorting and constrained-mobility dynamics the tradition describes (Crowder and Downey 2010; Banzhaf and Walsh 2008)), while around the persisting clean fleet of hydro and nuclear no systematic drift is expected.

4.4.3 E3: Entrant composition (RQ3)

Where new capacity comes to rest, the expectation is that different fuels fall under different branches of the trichotomy of replication, transformation, and null. For the emitting entrants, the new natural-gas capacity that makes up the bulk of fossil additions over the window, the prior is *replication*: the path of least resistance that shaped the fossil fleet has no evident reason to spare its newest additions, and new gas capacity is expected to land nearer poorer and more nonwhite communities, the fossil

race-and-class gradient reproduced in the new build. The expectation is strongest for gas *peaking* units in particular, long observed to sit disproportionately near poor and minority populations (Krieger, Casey, and Shonkoff 2016). The capacity-weighted dominant-fuel classification used here does not separate peakers from combined-cycle baseload, however, so the prior is entered for natural gas as a whole, with a prime-mover split reserved for robustness. For wind, the prior is *transformation*: its siting is bound by resource and land in a way that channels it toward rural, and disproportionately Whiter, country, so that the racial association is expected to weaken, or to give way to an axis of rurality and class, rather than to repeat the fossil pattern (O'Shaughnessy et al. 2023; Mueller and Brooks 2020). Utility solar is expected to sit between the two — less white than the fleet it joins, but bifurcated on class, in keeping with the static evidence (Cranmer et al. 2023). Biomass, emitting but sited rurally near its feedstock (agricultural and forestry residues), is expected to follow wind toward *transformation*, with its racial association faint and its host communities rural and modestly poorer. Battery storage, sited to the grid rather than to a resource and frequently supplanting the gas peakers that preceded it, carries no confident directional prior and is registered under the *null*. Across the renewable fuels taken together, then, the expectation is that the racial gradient is weaker and less consistent than the fossil case would lead one to predict — both because the technoeconomic filter selects on land rather than on a community's demographic capacity to resist, and because the political economy of clean siting, as the third section argued, is waged at a scale above the host community itself. Where resource potential swamps demography outright, the *null* is the appropriate expectation and is not to be mistaken for an absence of findings (Downey 2005). Beyond entry, the exit channel carries its own prior: with the less-disadvantaged coal expected to retire first, its departure steepens the disadvantage gradient among the surviving coal fleet (Richmond-Bryant et al. 2020). Table 4.2 consolidates the priors for all three channels (entry, exit, and stayer) in the spirit of a pre-registration: directional expectations fixed before the coefficients are seen, against which the estimates can later be read.

Table 4.2: Pre-registered expectations by decomposition channel: directional priors committed before estimation for entry, exit, and within-survivor drift.

	Limb / role	% nonwhite	% poverty	Income

Entry — new capacity vs existing fleet				
Natural gas	Replication	+	+	−
Wind	Transformation	− (weak)	+	−
Utility solar	Mixed	+	+	?
Biomass	Transformation	− (weak)	+	−
Storage	Null / mixed	?	?	?
Exit — retiring vs surviving coal				
Coal	Advantaged exit first	−	−	+
Stayer — within-plant drift, net of trend				
Legacy fossil (coal, gas)	Sorting drift	+ (weak)	+ (weak)	− (weak)
Persisting clean (hydro, nuclear)	~ flat	~0	~0	~0

Note. Priors are fixed before estimation; race is taken net of class. **Entry** and **exit** signs give the host composition of the *moving* plants relative to the comparison fleet (entrants vs the existing fleet; retiring vs surviving coal); **stayer** signs give the direction of *within-plant* drift over 2010–2019, net of the national trend. "+" = more nonwhite or poorer (lower-income), "–" = less nonwhite or wealthier; "(weak)" marks a direction expected but an association expected faint; "?" an ambiguous, class-bifurcated prior; "~0" no expected drift. On exit, the less-disadvantaged coal is expected to retire first, so the exit channel *steepens* the disadvantage gradient among the surviving coal fleet. Coal is the entry-regression reference and rarely enters; hydro and nuclear rarely enter; the remaining small-N entrant fuels (geothermal, oil, the residual 'other') are carried with priors held loosely. Under the *null*, where resource potential dominates siting, no systematic relation to demography is expected.

4.5 Results

Before turning to change, Table 4.3 establishes the cross-sectional picture: the mean host-community characteristics of each energy type, pooled across the operating panel. The fuel-specific quadrants the analysis will trace are already legible — coal's communities the whitest and poorest of the substantial fuels, gas's markedly more diverse, and solar's, in line with the static evidence, less white yet *higher*-income than the fossil fleet it displaces. The small-sample fuels the analysis reads with caution (geothermal, storage, the residual category) are visible here in their plant counts.

Table 4.3: Host-community characteristics by energy type, pooled over operating plant-years, 2010–2023.

		Race / ethnicity (share)	Class

	Plants	Nonwhite	Black	Hispanic	Asian	Indigenous	Poverty	Median income
Coal	558	20.3%	7.8%	7.5%	1.1%	1.7%	15.2%	$66,622
Oil	769	22.5%	8.4%	9.3%	2.0%	0.6%	14.7%	$68,192
Gas	2,280	39.0%	11.1%	19.9%	4.4%	0.8%	16.5%	$73,554
Hydro	1,438	17.5%	3.4%	8.6%	1.4%	1.5%	13.9%	$69,779
Biomass	909	32.4%	12.0%	13.8%	3.4%	0.6%	14.7%	$74,732
Geothermal	74	43.8%	2.1%	34.1%	1.6%	3.6%	16.2%	$72,586
Nuclear	65	18.8%	9.3%	6.6%	0.8%	0.3%	11.6%	$78,624
Solar	5,647	38.1%	11.5%	19.2%	3.4%	1.0%	14.2%	$78,538
Wind	1,404	20.0%	1.7%	14.2%	1.0%	1.0%	12.4%	$69,435
Storage	295	40.5%	5.0%	26.0%	5.7%	0.6%	13.9%	$87,331
Other	64	39.9%	13.2%	19.7%	3.4%	0.8%	16.0%	$68,266

Note. Race and class cells are means of the 3 km buffer measures; income in real 2023 dollars; Plants is the count of distinct plants. Fuels ordered legacy to clean.

4.5.1 The changing geography of energy-type communities (RQ1)

The fuel-by-fuel trajectories in Figure 4.2 bear out the theoretical expectations: there is no single demographic gradient running across the generating fleet, but a family of fuel-specific ones, moving over the transition decade in different directions and at different speeds. Read against the national population (the benchmark every series is netted against, so that what the figure shows is divergence from the country rather than raw level), the fleet falls into the now-familiar quadrants of the fuel-and-demography literature, and then, over the decade, redraws them (Figure 4.3).

Coal anchors the white-and-poor corner, and sits more deeply in it at the end of the decade than at the start. Its host communities are whiter than the nation throughout, by roughly 15 percentage points in 2010, and that gap widens to some 21 points by 2023. The widening is real, but it would be a mistake to read it as coal's neighborhoods turning whiter in place. Coal's aggregate host composition holds roughly flat across the panel even as the national population diversifies around it, and a fleet that stands still while the country moves comes, by that fact alone, to sit ever further below the national line. That flatness, however, is not inertness. It is the near-cancellation of two opposing forces the decomposition below separates: the communities of surviving coal plants grow

more nonwhite with the diversifying nation, while the plants that retire, drawn from more nonwhite communities than those that remain, pull the surviving mean back toward white. Because the first of these is simply the national tide that every fuel's survivors ride, what remains once that tide is netted out is the second. In other words, coal's relative whitening is the residue of retirement, not of its standing communities changing character. On the economic axis those same communities sit well below the national median income throughout: coal occupies the corner the literature has long assigned it, whiter and poorer than the country it powers.

The transition's signature is most clearly evident in the trajectory of solar. Where coal remains relatively flat, solar's demographics shift widely as the result of the build out, crossing clean through the national line. Early utility-scale solar landed in communities more nonwhite than the nation, by some 8 points in 2010; by 2019 that gap closed, and by 2023 solar's host communities are roughly 5 points whiter than the nation. In other words, the median solar community crosses, over the course of the buildout, from the more-nonwhite to the less-nonwhite side of the country. The anatomy of that crossing (Figure 4.13) is not a uniform drift but the emergence of a new mode: as the buildout disperses across the country, a dense and growing concentration of low-nonwhite, rural host communities that forms beneath solar's early, more demographically diverse, Californian footprint.

Completing the picture, gas, like early solar, sites in communities more nonwhite than the nation, and most distinctively in Hispanic ones. Wind, by contrast, is as white as coal, well below the nation on nonwhite share, and on class the two whiter fuels look much alike: both sit far below the national median income, yet neither carries an elevated poverty rate, coal's running near the national level and wind's a little below it. Theirs is a rural, low-wage profile: low incomes without a deep poor tail, not the concentrated distress a high poverty rate would mark. The upshot is that race and class do not move together across the fleet. The whitest fuels host the lowest-income communities, while the more diverse ones sit closer to the national median: gas and early solar are more nonwhite than the nation yet, in solar's case especially, *higher*-income than the fossil fleet it displaces.

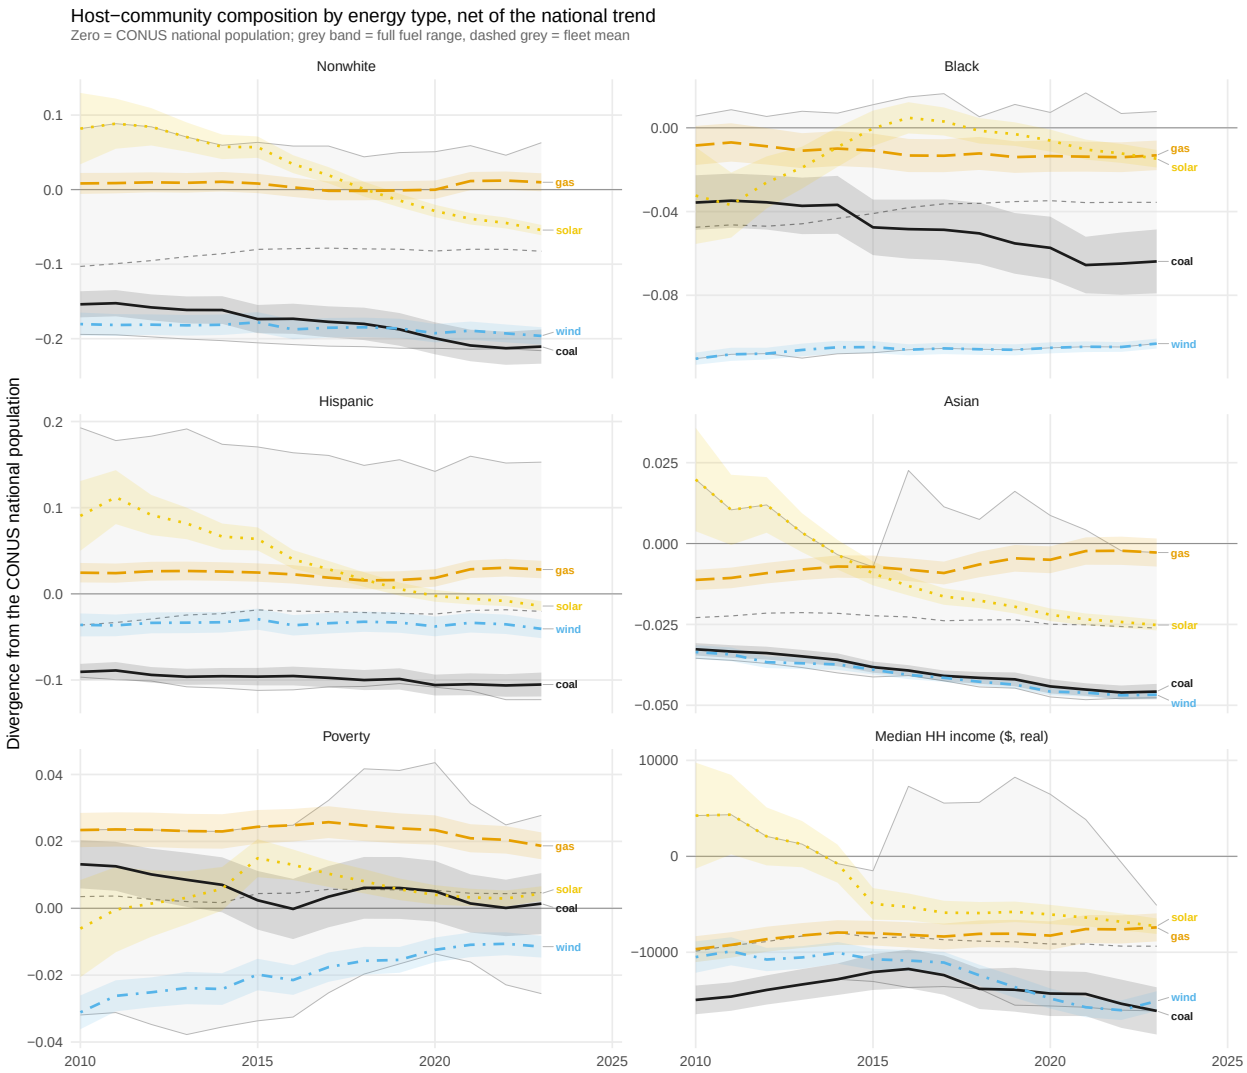

Figure 4.2: Host-community composition by energy type, net of the national trend, 2010–2023. Each colored line is a fuel's mean host-community value minus the contemporaneous CONUS national value (population-weighted over the same five-year ACS windows); zero is the national level. The grey band spans all eleven fuels, the dashed grey line is the fleet mean, and colored bands are 95% CIs for the four highlighted fuels. Income is real 2023 dollars.

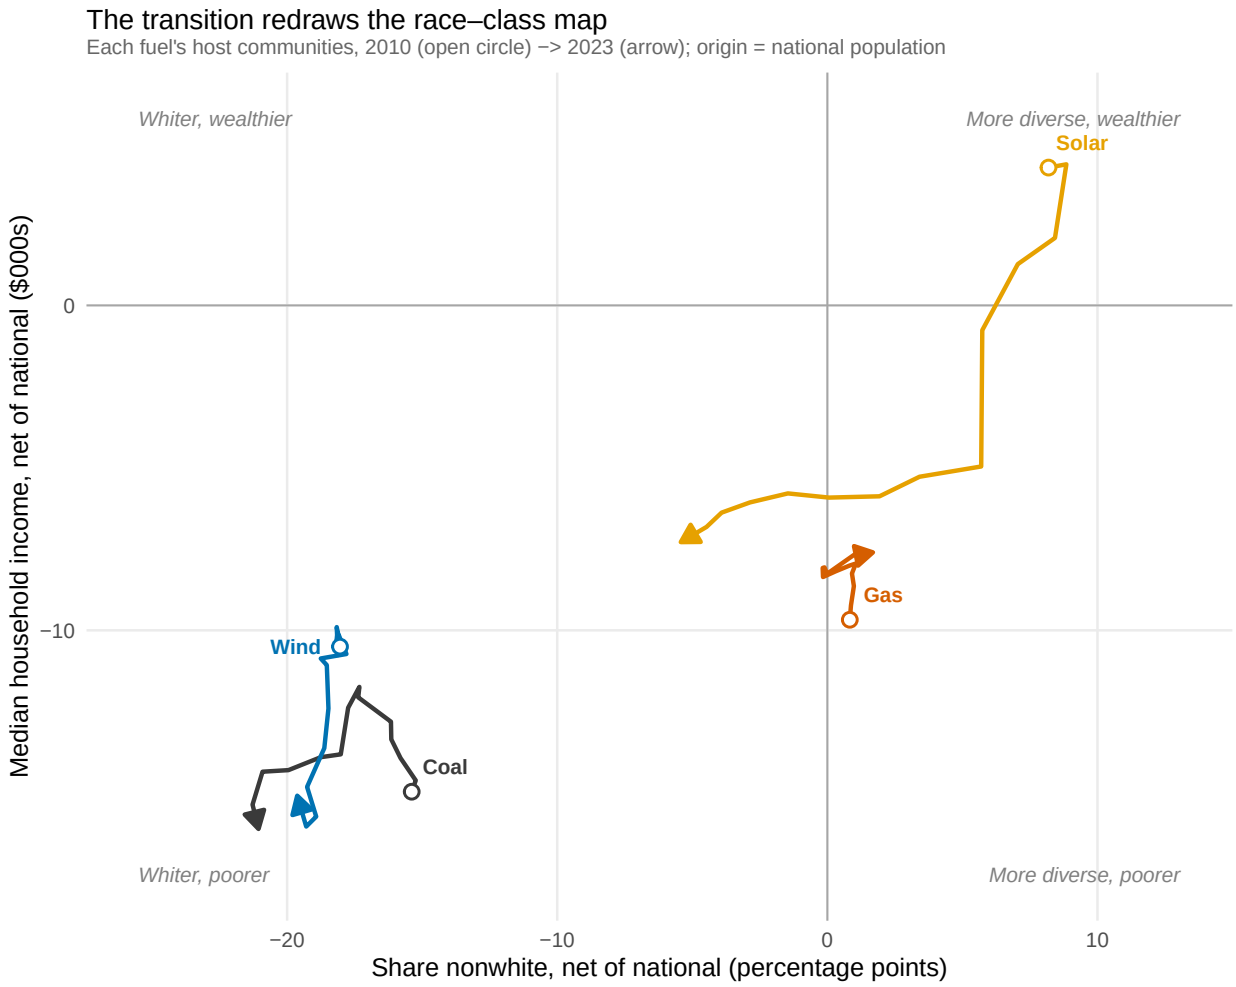

Figure 4.3: The race–class plane, 2010–2023: each transition-core fuel's mean host-community composition in share nonwhite and median household income, both net of the national population, traced from 2010 (open circle) to 2023 (arrowhead). The origin is the national benchmark; the crosshairs divide the quadrants. 3 km buffers.

4.5.2 Turnover versus neighborhood change (RQ2)

The trajectories of Figure 4.2 show the demographic map being redrawn; the decomposition asks how. Figure 4.4 partitions each fuel's change in mean host nonwhite share, over the primary 2010–2019 window, into the three channels the design distinguishes: the change among the communities of surviving plants, the contribution of entry, and the contribution of exit, which sum, by construction, to the total. Because the decomposition operates on raw composition rather than the detrended series, the within-survivor channel absorbs the national diversification common to every fuel, running a

similar modest positive across the fleet. The fuel-specific story is therefore told not by the within channel but by the turnover channels — and there it is told sharply.

Solar's change is an entry story almost in its entirety. New solar contributes some 10.1 points to the whitening of solar's mean — an entry channel that not only outweighs but reverses the 4.8-point more-nonwhite drift of its surviving communities, for a net move of 5.7 points toward the whiter side of the country. The decomposition thus locates the trajectory's signal precisely: solar's relative whitening is not its existing communities changing but the buildout itself, the thousands of new plants entering communities far whiter than the solar fleet they joined. Both figures show the shift, and the decomposition attributes it almost entirely to entry.

Coal's change, by contrast, is the work of exit. Where coal's surviving communities grew more nonwhite with the diversifying nation (a within channel of 3 points), the plants that retired sat in more nonwhite communities than those that remained, so their departure pulled the surviving mean back toward white by 2.4 points, very nearly cancelling the drift and leaving coal's mean almost unmoved. This is the demographic signature of departure: a fleet redrawing its own average by subtraction rather than by any change in who lives around it. The persisting fuels (hydro, nuclear, gas) turn over little and move with the national tide through the within channel alone. The contrast across the fleet, in other words, is not between fuels that change and fuels that do not, but between the channels through which the change arrives: entry for the technologies being built, exit for the technology being retired.

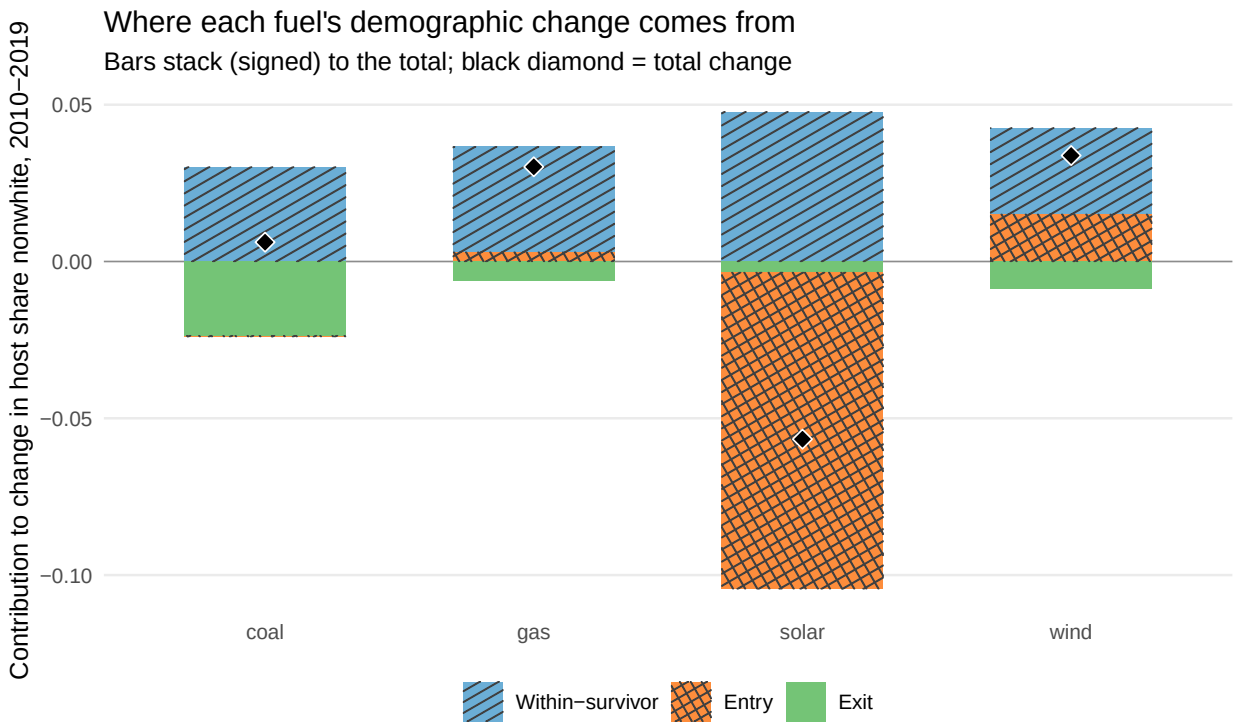

Figure 4.4: Melitz–Polanec decomposition of the change in host-community share nonwhite (3 km), 2010–2019, for the four transition-core fuels: the within-survivor, entry, and exit channels, signed and stacked to the total (black diamond). All eleven fuels and the 2010–2023 extension appear in Figure 4.11.

4.5.3 Where new capacity lands (RQ3)

The trajectory and its decomposition establish *that* the communities receiving new capacity differ from those that hosted the old, and by how much; the entrant-composition analysis asks *at what geographic scale* that difference is drawn. The question matters because two very different processes would leave the same fleet-wide footprint. New capacity might land, within any given place, in systematically whiter or more nonwhite neighborhoods than the incumbent fleet — a local sorting of technology onto community. Or the fleet might simply be relocating across the map — new plants built in whiter or more nonwhite *states* than the old, with no particular targeting of the communities within them. Figure 4.5 separates the two by estimating each fuel's entrant-versus-incumbent gap twice: once across the whole country, and once within states (net of state fixed effects). The distance between the two is the share of the entrant difference that is between-state relocation rather than within-place selection.

For the fuel that drives the transition's demographic signal, solar's whitening is relocation, almost in its entirety. New solar sits in communities some 12 points whiter than the incumbent solar fleet across the country as a whole — but that gap all but vanishes, to 2 points, once the comparison is made within states. The between-state component, 10 of those 12 points, accounts for nearly all of it. New solar is not, within California or Arizona or North Carolina, seeking out whiter neighborhoods than the solar already there; it is being built in whiter *states* than the early fleet was — the dispersal of the buildout out of its early Californian concentration and across the interior, the Plains, and the rural Southeast, which Figure 4.6 and Figure 4.13 record. The demographic map of solar is redrawn at the scale of the state, not the neighborhood. The map makes something else plain: the whitening is an average, not a uniform movement. A substantial share of the recent buildout lands in the rural Southeast, in communities far from white, so the fleet-wide shift toward whiter hosts coexists with a real and growing Black-rural tail, the same tail the distributional view leaves visible.

The pattern generalizes, with one qualification. Across the fleet the within-state entrant gaps are small — new plants, whatever the technology, tend to land in communities much like the incumbents they join once the state is held fixed. Gas is the partial exception: its entrant communities run more nonwhite than its incumbents' both between states and, modestly, within them (3 points within-state), the one fuel to show a trace of genuine within-place selection. A methodological caution belongs here. Because the entrant gap lives between states, and because the estimator clusters on state, the standard errors on these gaps are wide. What the entrant-composition analysis can suggest is the scale at which the redrawing is occurring. The transition moves capacity across the country, not onto particular communities within the places it already occupies. This fits the technoeconomic account of renewable siting: the endowments of resource and land that fix where a project is viable operate at the regional scale, and the composition of the immediate community matters only faintly, if at all.

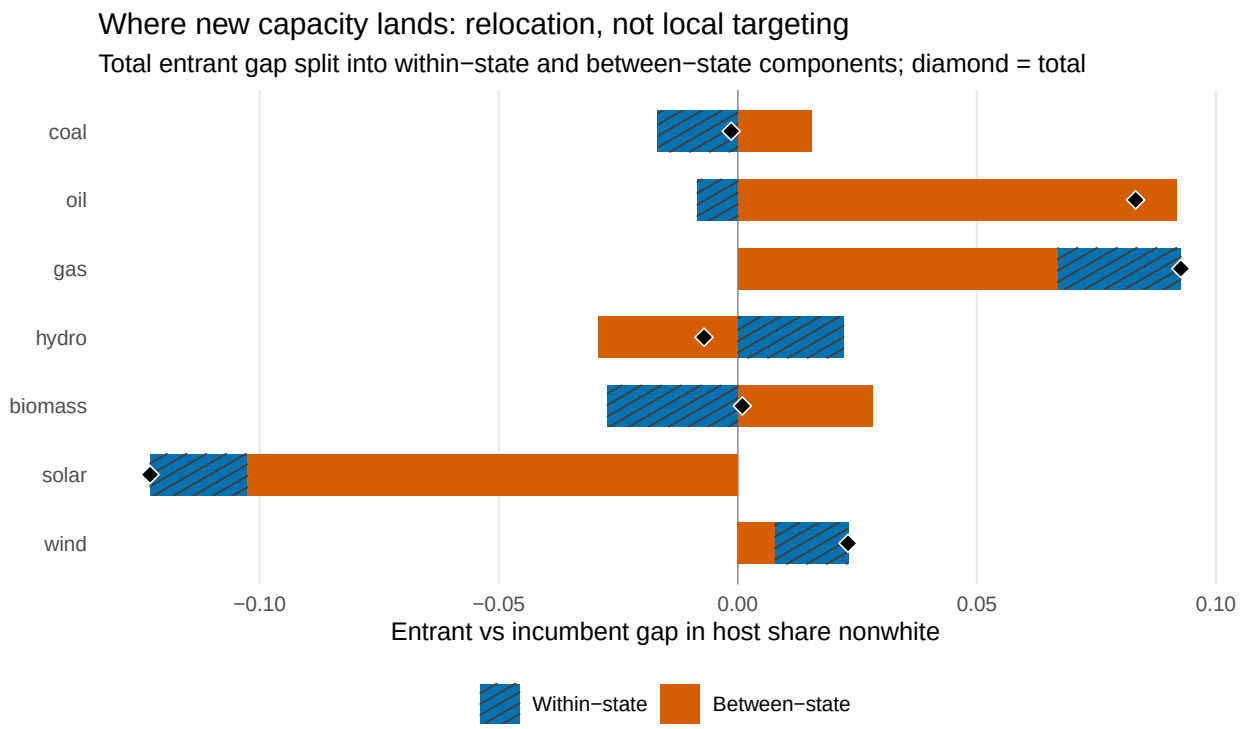

Figure 4.5: The total entrant-versus-incumbent gap in host-community share nonwhite (3 km), by fuel, split into a within-state component (net of state fixed effects) and a between-state component; the two bars sum to the total (black diamond). SEs are wide (see text); point estimates are read descriptively. All fuels and outcomes are in Figure 4.12.

The between-state relocation, in the case of solar, is visible in Figure 4.6, which plots the solar fleet by build era. The map shows the pre-2010 plants cluster tightly in California, while the transition-era cohorts (2018 and 2021) scatter across the Southeast, Texas, and the interior — the geographic face of the between-state shift, and direct evidence for the Californian origin the trajectory could only infer.

Figure 4.6 shows the relocation geographically: the pre-2010 incumbent fleet concentrates in California, and the 2010–2016 and 2017–2023 cohorts disperse into the South and Midwest.

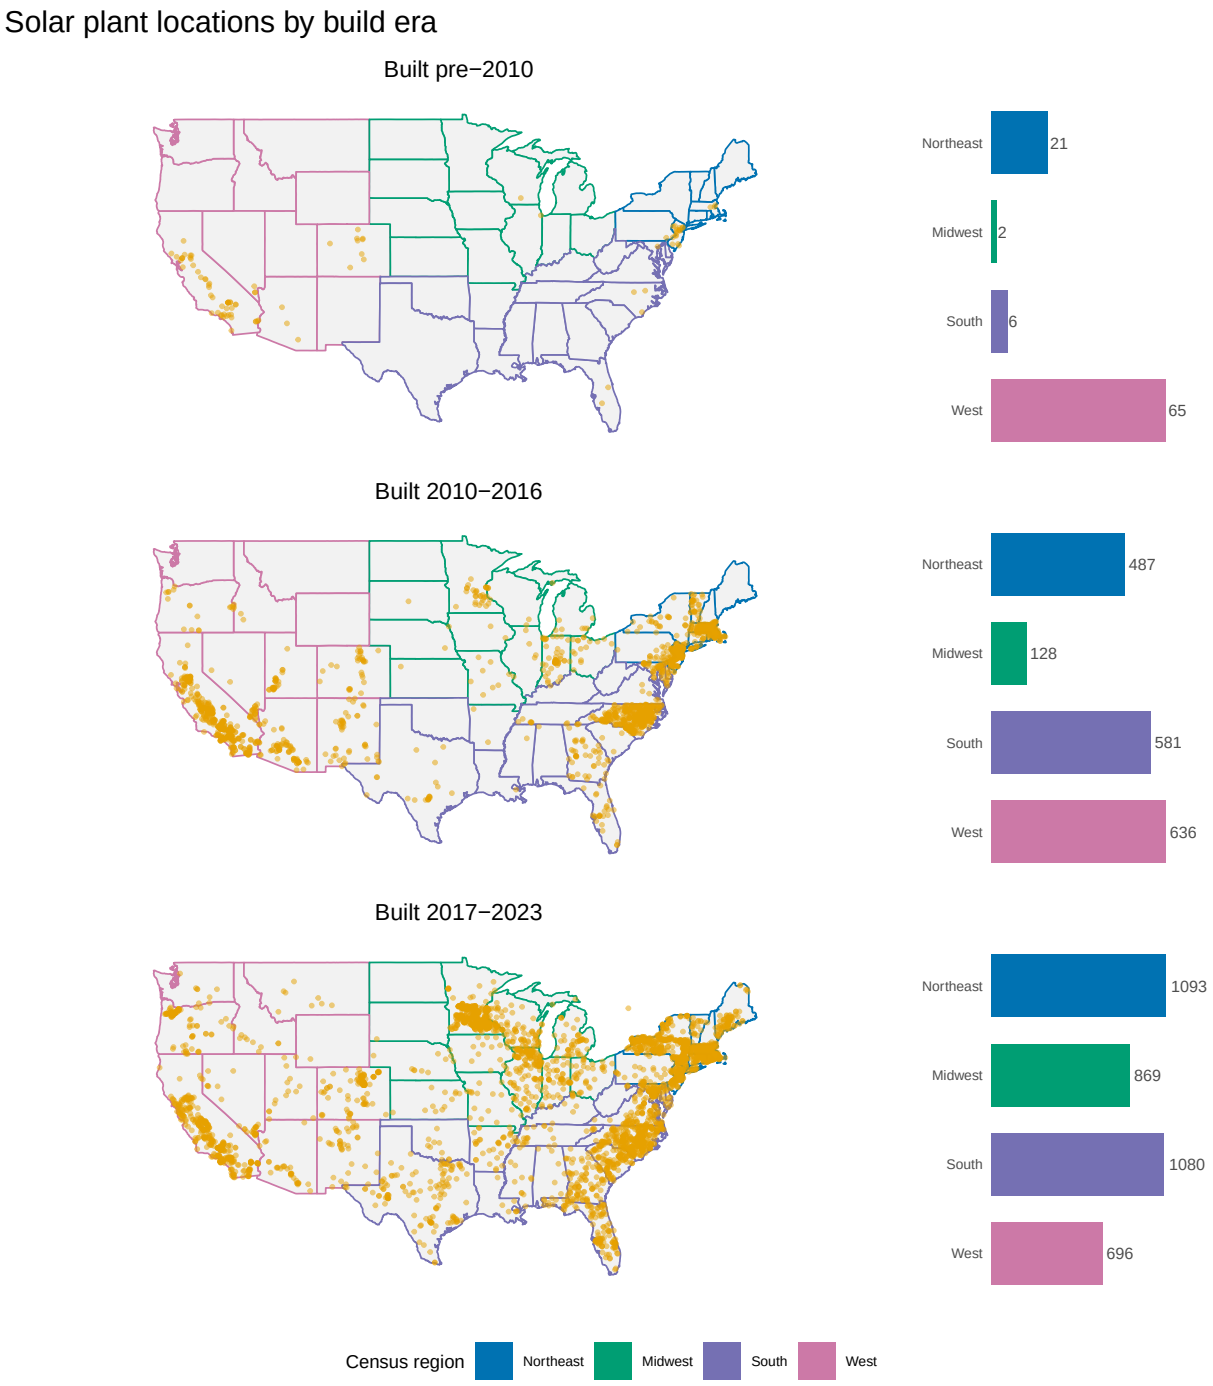

Figure 4.6: Utility-scale solar plant locations by build era (pinned coordinates), with plant counts by census region; state outlines are colored on the same region scale as a geographic key (Alaska, Hawaii, and DC fall outside the base map and appear only in the bars).

4.5.4 The Indigenous dimension

The theory set the Indigenous dimension apart from the rest, modelling proximity to tribal land as a threshold rather than a gradient and anticipating that it would behave unlike race — concentrated around particular technologies and locations, and shaped by a settler-colonial history the path-of-least-resistance frame cannot capture. The results bear that expectation out: the Indigenous dimension yields the analysis's single statistically distinguishable entrant effect.

The result has two parts, one related to the standing fleet and one related to new construction. Among plants already standing, the fuels likeliest to sit near tribal land are the old ones: geothermal, whose 16% of 2023 plants in high-Indigenous communities leads the fleet, followed by coal at 6% and hydro (geothermal's lead rests on a few dozen plants and carries a wide interval). This is the Western resource geography of geothermal steam fields, coal, and dams, inscribed into Native lands long before the transition began. Among plants newly built, solar is the one fuel whose capacity tilts, significantly, toward tribal-proximate communities: net of state and of rurality alike, an entrant solar plant is roughly 3.8 percentage points likelier than the incumbent solar fleet to sit in a high-Indigenous community (Figure 4.7). Unlike the racial signature of solar, a relocation *across* states that vanished *within* them, this one is a within-state, within-place selection the state fixed effects do not absorb. It survives where every other siting coefficient dissolves for a mechanical reason: tribal land is a within-state feature, so the between-state noise that swamps the racial estimates leaves it untouched.

National evidence gives the tilt a further reading. The dispossession that created the reservation system left tribes, unintentionally, with favorable wind and solar endowments, most of all on the lowest-income reservations. Yet reservation lands remain far less likely to host wind and solar projects than comparable lands just beyond their borders (by their estimates, 46 and 110 percent less likely, respectively), a gap worth more than US$19 billion in foregone lease and tax earnings through 2050 (Parker et al. 2024). Read against that evidence, the entrant tilt found here is an adjacency pattern. Because the outcome is proximity rather than tenure, a buildout that develops the land beside reservations while passing over the land within them would produce this signature: projects landing beside Native communities, ownership and revenue accruing elsewhere.

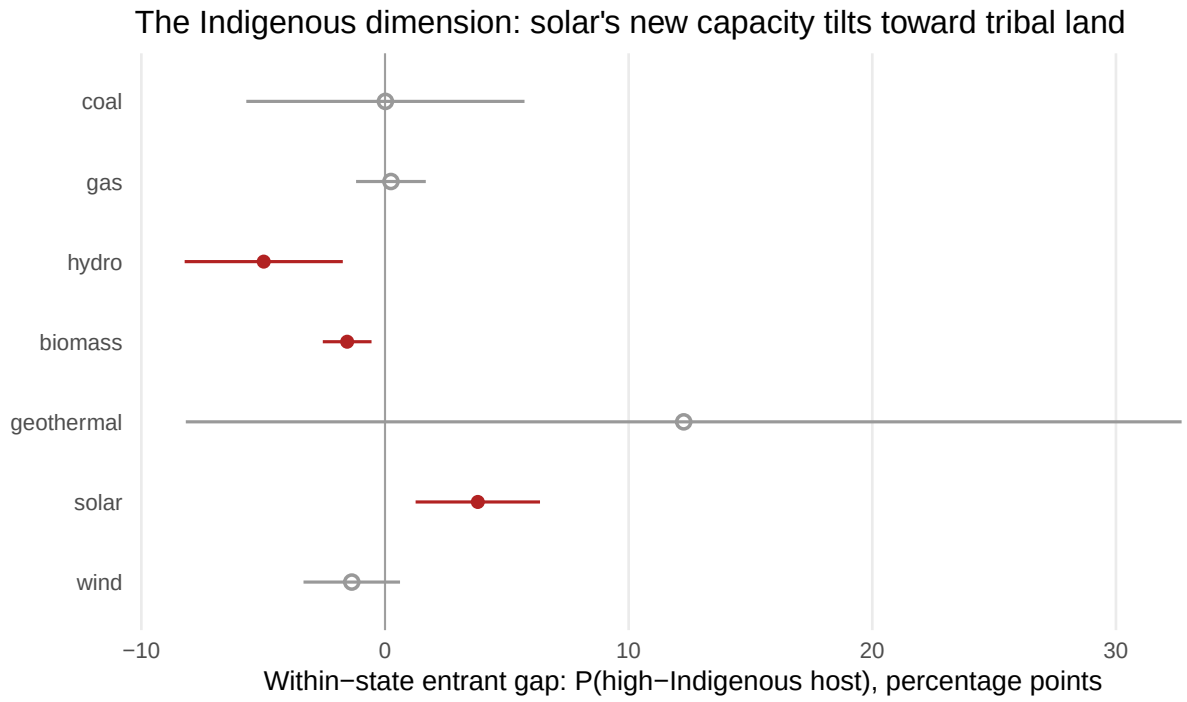

Figure 4.7: Within-state entrant-versus-incumbent gap in the probability of a high-Indigenous host community (>=5% Indigenous, 3 km buffer), by fuel, with 95% CIs (state-clustered); colored where the interval excludes zero.

4.6 Discussion

Read together, the trajectory, its decomposition, and the entrant composition tell a consistent story about how the spatial demography of American electricity is being redrawn as the system decarbonizes: fuel by fuel rather than uniformly; through the turnover of the fleet rather than the churn of neighborhood; and, where the renewables are concerned, at the scale of the state rather than the locale. The transition is not, on this evidence, sorting clean technology onto whiter communities and dirty onto more nonwhite ones within a shared local landscape. It is relocating the map: coal retiring from communities more nonwhite than the coal that remains, solar building out in whiter states than the early fleet occupied. The demographic consequences follow from where the retiring and the building fall.

Against the priors fixed in advance, the trajectory and the decomposition mostly conform. The fuels diverge rather than converge. Coal's host communities scarcely change, but as the nation

diversifies around them they sit ever further below the national nonwhite share. Solar's cross the national line outright, from more nonwhite than the nation to whiter. Gas's stay more nonwhite than the nation throughout, most distinctively in their Hispanic share. Wind's are as white as coal's, and nearly as low-income. The pattern is fuel-specific, as the theory anticipated, and the decomposition shows the same shape by channel: solar's change is entry almost in its entirety (an entry channel of 10.1 points), coal's runs through exit, and the persisting fuels drift only with the national tide. The exit channel, however, runs against its registered prior. The expectation, drawn from the retirement literature, was that the less-disadvantaged coal would retire first, steepening the disadvantage gradient across the surviving fleet (Richmond-Bryant et al. 2020). The data show the reverse: the coal that retired sat in more nonwhite communities than the coal that survives, so retirement, by subtraction alone, whitened the surviving fleet's average. The reversal is partly a matter of design: Richmond-Bryant compare the retiring fleet's burden against the nation's, weighted by the people who bear it, whereas this analysis compares retiring coal against the coal that survives, counting plants. The two questions can run in opposite directions across the same fleet. The decomposition adds the mechanism: the map moves because the fleet turns over, not because the communities around standing plants are themselves transformed.

The entrant composition is where the registered expectation runs into complication. The transformation prior for renewables held that the racial gradient of fossil siting would weaken rather than repeat. It holds, but not quite for the reason the prior supposed. Within states, new solar sits in communities essentially indistinguishable from the solar already there (a within-state gap of 2 points against a total of 12), so the fossil path of least resistance is indeed not reproduced in the renewable buildout. But the difference between new and old solar communities, so visible in the decomposition, does not thereby vanish; it *relocates*. It is a between-state phenomenon (Figure 4.6): new capacity built in whiter states, not whiter neighborhoods, with the local community's composition adding almost nothing once the state is held fixed. The technoeconomic account holds, but at a scale the prior did not specify. The resource and land endowments that decide where a solar project is viable (sun, open acreage, transmission) are facts of regional geography, not of the census tract, so the demographic signature of renewable siting is written at the resolution of the state, and the host community enters, if at all, only faintly. Gas is the one fuel to show a trace of the fossil pattern

within places, its entrants modestly more nonwhite than its incumbents even with the state fixed: the replication prior survives, weakly, and only for the emitting fuel the theory singled out.

Two qualifications bound the synthesis. The first is that solar's whitening is an average, and averages conceal. Beneath the fleet-wide shift toward whiter hosts runs a substantial and growing tail of solar in the Black-rural Southeast (Figure 4.13) that a flat "renewables go to white communities" reading would efface. The second is that race and class do not travel together across the fleet: the whitest fuels, coal and wind, host the lowest-income communities, while the more diverse ones sit closer to the national median, solar higher-income than the fossil fleet it joins. A composite index of 'disadvantage' would have flattened the crosscutting structure the separate treatment of race and class was designed to preserve. An intersectional treatment, examining race and class jointly rather than separately, is the natural next step (Raoul S. Liévanos 2015, 2019).

One further lens reframes the analysis. Every quantity reported so far counts *plants* — each generator one observation, weighted equally. Count *people* instead, weighting each plant by the population its buffer contains, and the fleet's host communities turn systematically more nonwhite (Figure 4.8). The reason is structural: more nonwhite communities are denser because the nonwhite population is disproportionately urban, a legacy of migration to cities and of the segregation that confined Black and immigrant households to central neighborhoods (Cushing et al. 2023; Raoul S. Liévanos 2015), so a plant sited in one covers more residents than a plant in the white rural interior, and person-weighting registers that difference. The shift is large and nearly universal: across the fleet, the population living near generation is more nonwhite than a plant-count average implies, and for some fuels the reweighting reverses the comparison to the nation outright. Solar, whiter than the nation counted by plant, is more nonwhite than it counted by person; wind, markedly whiter by plant, sits about at the national average by person; even coal's host communities, the whitest of any fuel by plant, are less so by the people who actually live in them. The plant-side account this chapter offers is the appropriate one for its question, where capacity is *sited*, but it is not the person-side account of who is *exposed*; the gap between them, visible here, is why the absolute-burden question requires a tract-side design.

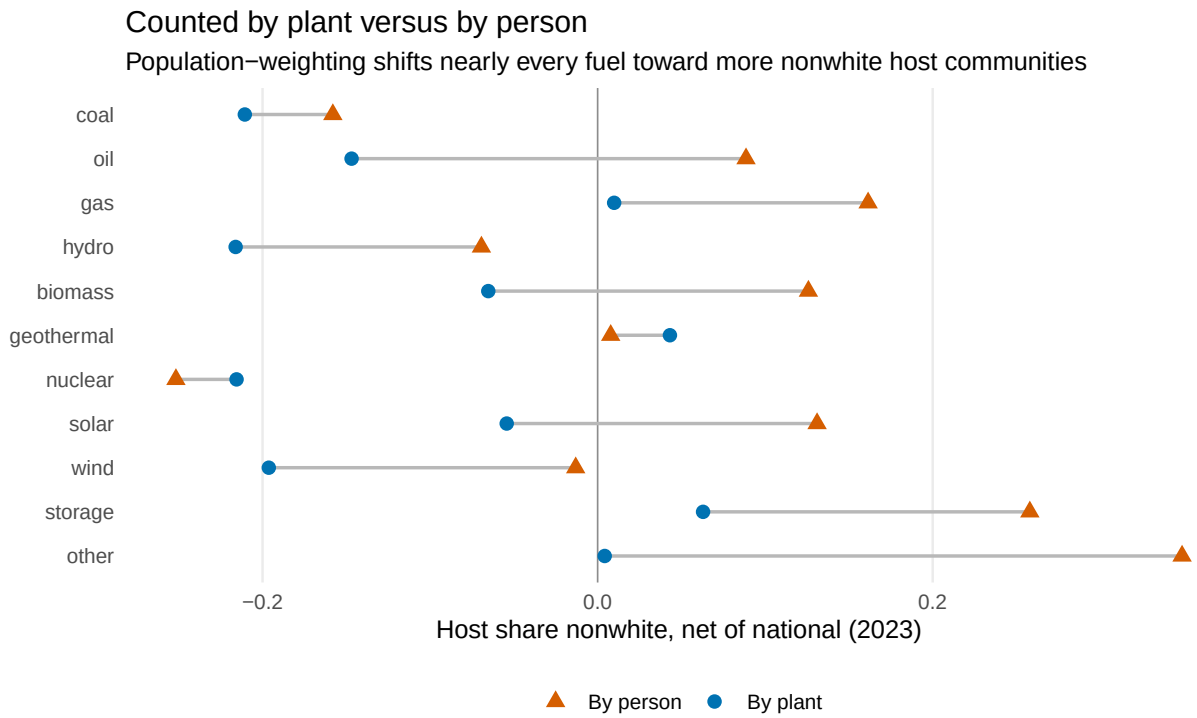

Figure 4.8: Host-community share nonwhite in 2023, net of the CONUS national population, by fuel, counted two ways: by plant (each generator weighted equally) and by person (each plant weighted by its buffer population).

Though the reading is descriptive, its policy implication is legible, and it follows from the twin findings that the map moves by turnover and at the scale of the state. A spatial demography redrawn by fleet turnover is one that siting and interconnection policy can reach (the queues, the permitting regimes, the state renewable standards that govern where capacity is built and retired), rather than one only housing and mobility policy could touch, as a story of neighborhood change would imply. That the renewable shift is between-state, moreover, locates the operative scale above the municipality and the county whose zoning fights dominate the popular account of clean-energy siting: the demographic composition of the buildout is being set, in the main, by which states build, which is to say by the state-level political economy of the transition rather than by the local contest over any single project. The coal side carries the transition-justice concern in its familiar form: whatever the aggregate benefits of retirement, the surviving fossil fleet hosts communities modestly more nonwhite and higher in poverty than the fleet as a whole, a residual concentration a just-transition policy would have to address.

A final observation concerns scale. Whatever the demographic gradient of renewable siting (faint, as this analysis finds it, and drawn between states rather than within them), the number of communities the transition draws near electricity generation is multiplying. Figure 4.9 counts the host communities of each fuel across the panel: where coal's footprint contracts with retirement and gas's holds roughly steady, solar's climbs steeply, from some 105 operating communities in 2010 to 5,558 by 2023, overtaking coal early, passing gas well before the decade is out, and ending it hosting more communities than coal, gas, and wind combined. The transition's most consequential distributional fact may be less that its new capacity lands in whiter communities on average than that it lands in so many *more* of them: the horizontal geography of renewable power multiplies the arenas of siting politics, and with them the communities that will contend, or fail to contend, over what is built nearby (Bridge et al. 2013). The composition of that vastly enlarged host population is the descriptive baseline this chapter offers; who comes to bear what, as the footprint keeps growing, is a question for a tract-side design.

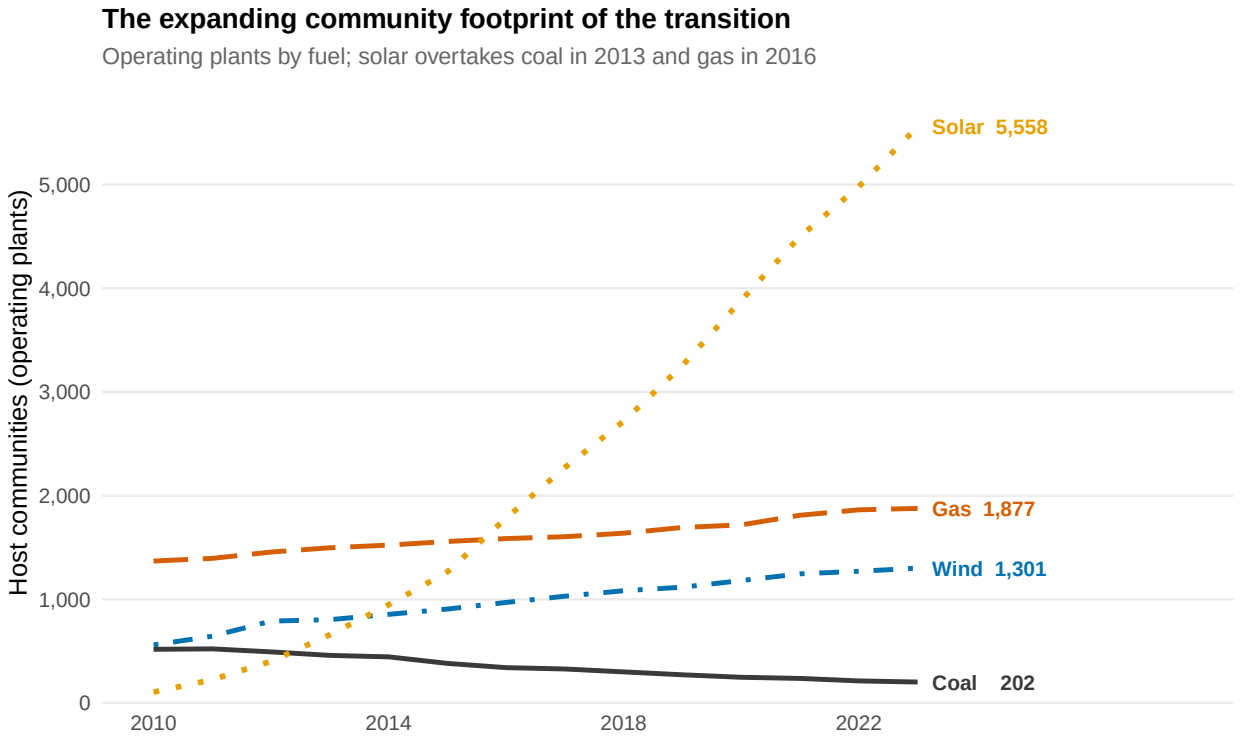

Figure 4.9: Number of host communities (operating plants) by energy type, 2010–2023, for the four transition-core fuels.

4.7 Limitations

Four limitations bound what the analysis can claim.

The first and most consequential is that the account is plant-side and relative, not person-side and absolute. As the plant-versus-person comparison showed (Figure 4.8), a fleet that whitens by plant count still hosts a disproportionately nonwhite population; the exposure question (how much pollution falls on whom) belongs to a tract-side design that takes the community, not the plant, as its unit (Richmond-Bryant et al. 2020).

The second is that the estimand is associational, for the reason the identification section set out: the survey's overlapping five-year windows smooth the demographic series too heavily to cut the clean temporal lag a causal siting-versus-sorting test would demand (Banzhaf, Ma, and Timmins 2019; Clay et al. 2025). The entrant analysis compares new and incumbent communities as observed, not a community against its earlier self, so its estimates are description, not effect. In Saha and Mohai's sense of the term, prior neighborhood composition predicting where hazards subsequently land (Saha and Mohai 2005), the entrant-composition analysis is therefore not a siting analysis, and it makes no claim to be one.

The third concerns the entrant composition. Its entrant gaps carry wide, state-clustered standard errors, so they are read descriptively rather than for significance (as the results noted). More substantively, the finer technoeconomic determinants of siting the theory names (resource quality, transmission access, land cost) are not modelled here beyond the single proxy of community density, so the between-state relocation is documented rather than explained: what draws the buildout to the states it favors is a question the resource and transmission data of a subsequent design would have to take up.

The fourth is that host-community composition moves for reasons of its own. The decomposition's within-survivor channel records that drift, but the models make no attempt to explain it: the housing and land markets, residential segregation, and racialized sorting that the post-siting literature identifies as the engines of neighborhood change (Been and Gupta 1997; Pastor, Sadd, and Hipp 2001; Mohai and Saha 2015) enter nowhere as covariates. Where incumbent communities diversify or whiten around a standing plant, the analysis registers the change without attributing it, and any account of why the buildout's host communities look as they do would have to model those forces

directly.

Two narrower cautions round out the account. The small-sample fuels (geothermal, storage, and the residual category) carry wide uncertainty and are read only as context for the substantial technologies, not as findings in their own right; and the energy-type classification, a capacity-weighted dominant fuel, does not separate combustion-turbine peakers from combined-cycle baseload within natural gas, so the gas result is entered for the fuel as a whole and a prime-mover split is left to robustness (Krieger, Casey, and Shonkoff 2016).

The demographic measures carry the same two caveats as the previous chapter's. The ACS's tract-level estimates are not screened for reliability by the coefficient-of-variation threshold the literature recommends (Folch et al. 2016; Raoul S. Liévanos 2018, 2019), and the documented geography of ACS unreliability, concentrated in the southern United States and in low-income communities and communities of color, includes the communities on which several of the chapter's findings turn, the Black-rural Southeast among them. And the decennial censuses of 2010 and 2020 are not used in the ACS's place: they carry no income or poverty measures, and two enumeration dates cannot supply the annual series a fourteen-year trajectory requires, so every demographic measure comes from the end-anchored five-year windows, with a decennial-based sensitivity analysis left to future work.

None of these is a defect to be corrected within the present design. Each marks the boundary between the question the standing fleet can answer (how the spatial demography of American generation is being redrawn, and by which channel) and the questions of absolute burden and causal mechanism, which would require a tract-side design.

4.8 Conclusion

This chapter set out to learn how the spatial demography of American electricity generation is being redrawn as the system decarbonizes. Where the nearest precedent fixed the fuel-and-demography question in a single cross-section (Cranmer et al. 2023), the analysis here followed the full fleet, renewables included, across the transition decade, decomposing each fuel's change into channels, splitting new entry into within- and between-state parts, and tracing energy-specific trajectories. The answer is one: the map is being redrawn fuel by fuel rather than uniformly, by the turnover of

the fleet rather than the churn of neighborhoods, and, for the renewables, at the scale of the state rather than the locale.

Located this way, the transition's spatial demography is one that policy can reach. Because it moves by turnover and between states, it answers to the interconnection queues, permitting regimes, and renewable standards of the states rather than to municipal zoning or to housing and mobility policy. The surviving fossil fleet's communities remain modestly more nonwhite and higher in poverty than the fleet as a whole, a residual for just-transition policy; and the sheer number of communities the buildout draws into hosting (solar's footprint alone up more than an order of magnitude) multiplies the arenas in which that politics will be contested.

The clean transition, on this evidence, is not erasing the geography of energy injustice but redrawing it — through which plants the fleet builds and retires, at the scale of the state, and across a host population that grows larger each year. Whether that redrawing narrows the fossil era's disparities or reinscribes them is the question the transition now poses; the descriptive baseline drawn here is a starting point for developing the answer.

4.9 Appendix

4.9.1 Variable inventory

Every variable the analysis uses is inventoried below by conceptual role — the identifiers and model terms, the energy-type predictor and the fleet-dynamics terms that structure the three analyses, and the community-demographic outcomes. Each row gives the variable, a plain-English label, its units, source, and a one-line construction. Only the 3 km buffer measures, the primary radius, are listed; their 5 km and 10 km analogues are identical in construction and retained as radius robustness.

4.9.1.1 Identifiers and model terms

Table 4.4: Codebook: identifiers and model terms.

Variable	Label	Units	Source	Operationalization

orispl	EIA plant identifier (ORIS ID)	integer	EIA-860	Direct field; primary key for the plant-year.
year	Observation year	2010–202	EIA-860	Annual panel year. Enters every model as a fixed effect, absorbing the secular national trend and the ACS overlapping-window smoothing.
state	State (two-letter)	factor	EIA-860	Plant's USPS state code. The siting regression's fixed-effect and cluster unit; the within-state vs between-state split turns on it.
latitude	Plant latitude (pinned)	degrees	EIA-860	Pinned to the median of the plant's valid continental-US coordinates across all panel years before buffering, so coordinate drift does not register as community change (App. Plant-location measurement).
longitude	Plant longitude (pinned)	degrees	EIA-860	As latitude; median-pinned before buffering.

4.9.1.2 Predictor and fleet dynamics

Table 4.5: Codebook: predictor and fleet dynamics.

Variable	Label	Units	Source	Operationalization
energy_type	Plant energy type (11-level; key predictor)	factor (11)	EIA-860 (duck_db primary_ energy_type)	Capacity-weighted dominant operable fuel per plant-year: the fuel category with the most operable nameplate capacity, ties broken alphabetically, NA when no operable generator. Eleven levels (coal, oil, gas, hydro, biomass, geothermal, nuclear, solar, wind, storage, other); coal is the reference.

entry_year	Build (operating-start) year	year	EIA-860	MIN(oldest_generator_year) across panel years; the fine date of a plant's construction.
entrant	Entrant indicator (siting analysis)	0/1	constructed	1 if entry_year □ [2010, 2023]; 0 = incumbent (built before the panel). The structural term interacted with energy_type in the siting regression.
presence_class	Decomposition role	factor	EIA-860 (R/diagnostics/05)	Presence across the decomposition's two endpoints: survivor (present both), entrant (later endpoint only), exiter (earlier only). Defines the entry/exit/within partition of the Melitz–Polanec decomposition.
retirement_year	Retirement year	year	EIA-860	First year a plant reports 'RE' (retired) status, unioned across data vintages; dates the exit margin.

4.9.1.3 Outcomes: community demographics

Table 4.6: Codebook: community-demographic outcomes.

Variable	Label	Units	Source	Operationalization
pct_nonwhite_3km	Share nonwhite, 3 km buffer (lead outcome)	fraction 0–1	ACS 5-yr × TIGER tracts	Population-weighted apportionment of tract shares to the 3 km buffer: weight = tract pop × (tract area inside ring ÷ tract area), end-anchored 5-year ACS window. 5 km and 10 km analogues retained as radius robustness.
pct_black_3km	Share non-Hispanic Black, 3 km buffer	fraction 0–1	ACS 5-yr × TIGER tracts	As pct_nonwhite_3km, for the non-Hispanic Black share (ACS B03002_004).
pct_hispanic_3km	Share Hispanic, 3 km buffer	fraction 0–1	ACS 5-yr × TIGER tracts	As pct_nonwhite_3km, for the Hispanic (any race) share (ACS B03002_012).

pct_asian_3km	Share non-Hispanic Asian, 3 km buffer	fraction 0–1	ACS 5-yr × TIGER tracts	As pct_nonwhite_3km, for the non-Hispanic Asian share (ACS B03002_006); small and right-skewed, read as supplementary.
pct_indigenous_3km	Share non-Hispanic AIAN, 3 km buffer	fraction 0–1	ACS 5-yr × TIGER tracts	As pct_nonwhite_3km, for the non-Hispanic American Indian / Alaska Native share (ACS B03002_005). Carried continuous but modelled via the threshold indicator below.
hi_indigenous_3km	High-Indigenous-proximity indicator, 3 km	0/1	constructed	1 if pct_indigenous_3km $\geq$ 0.05 (a $\geq$ 0.10 cutoff reported alongside). Indigenous proximity is a near-tribal-lands threshold, not a smooth gradient, so it enters as this binary; the siting model for it is a linear-probability specification.
pct_poverty_3km	Share below poverty, 3 km buffer	fraction 0–1	ACS 5-yr × TIGER tracts	Population-weighted apportionment of the tract poverty rate to the 3 km buffer (same weight as the race shares). 5 km / 10 km analogues as radius robustness.
median_hh_income_3km_real	Median household income, 3 km (real)	constant 2023 USD	ACS 5-yr × TIGER tracts, deflated	Population-weighted apportionment of tract median household income, deflated to constant 2023 dollars (R/join/04) so decade change registers real fortunes, not inflation.
ln_median_hh_income_3km	Log real median household income, 3 km	log(2023 USD)	constructed	log(median_hh_income_3km_real); the income form entered in the siting regression, its coefficient read as a proportional difference.

4.9.2 Plant-location measurement

We represent each plant by a single point and apportion community demographics to distance buffers around it. This appendix examines three measurement concerns that bear on that choice: the stability of the reported coordinate over time, the spatial extent of dispersed wind facilities, and the compactness of solar facilities.

4.9.2.1 Coordinate stability and pinning

EIA-860 reports one coordinate per plant-year, but that coordinate is not stable across years for a non-trivial share of plants — most acutely the spatially dispersed technologies (Table 4.7). Because plants do not move, year-to-year coordinate movement displaces the buffer and registers as *artifactual* community change; diagnostically, the spurious demographic jump scales monotonically with the size of the coordinate displacement. We therefore pin each plant to the median of its valid (continental-U.S.) coordinates across all panel years before buffering, so that within-plant change reflects the community rather than the geocode.

Table 4.7: EIA-860 single-point coordinate drift across years, by energy type (raw, pre-pinning).

Energy type	Plants	Drifters	% drift
wind	1404	353	25.1
oil	763	168	22.0
biomass	904	130	14.4
—	212	28	13.2
geothermal	74	5	6.8
gas	2272	144	6.3
other	63	3	4.8
hydro	1436	61	4.2
solar	5638	201	3.6
storage	295	10	3.4
coal	557	16	2.9
nuclear	65	0	0.0

4.9.2.2 Wind: turbine-array versus point buffers

A wind "plant" is in fact dozens of turbines spread over a large footprint (median bounding box $\approx$ 13.5 km across), so a point-buffer measures only a fraction of it. Using the U.S. Wind Turbine Database (Rand et al. 2020; Fujita et al. 2026), the analysis re-measures each matched facility as the union of buffers around all its turbines and compares it to the point-buffer (Table 4.8). The array captures several times more population, but the *aggregate* host-community composition is essentially unchanged: the median difference in community percent-nonwhite is near zero at every radius, with individual facilities reshuffling symmetrically about it. The point-buffer is therefore unbiased for the unweighted composition measures used in the primary analysis; the array measurement matters chiefly for population-weighted (exposure) framing.

Table 4.8: Wind host-community demographics: turbine-array buffers minus point buffers, by radius.

Radius (km)	Facilities	Median diff (pct nonwhite)	Mean abs diff	p90 abs diff	Pop ratio (array/point)
3	1235	0e+00	0.0208	0.0510	4.32
5	1236	0e+00	0.0233	0.0573	2.82
10	1237	1e-04	0.0275	0.0787	1.88

4.9.2.3 Solar facility footprint

Solar facilities, by contrast, are compact contiguous arrays. The U.S. Large-Scale Solar Photovoltaic Database records a median footprint whose area-equivalent radius is well under a kilometer (Table 4.9), far smaller than the smallest (3 km) buffer, so the centroid-buffer already encloses the facility and its immediate surroundings. No footprint correction is warranted for solar.

Table 4.9: U.S. large-scale solar facility footprint (USPVDB); area-equivalent radius.

Facilities	Median area (ha)	Median equiv-radius (m)	p90 equiv-radius (m)	Max equiv-radius (m)
6611	6	139	632	2446

4.9.3 All energy types: full trajectory reference

Figure 4.2 highlights the four transition-core fuels (coal, gas, solar, wind). For completeness, Figure 4.10 plots all eleven energy types on the same net-of-trend basis, so the seven fuels held back from the main figure (including the small-sample types carried with wide uncertainty) are visible. Fuel-years with fewer than thirty plants are dropped.

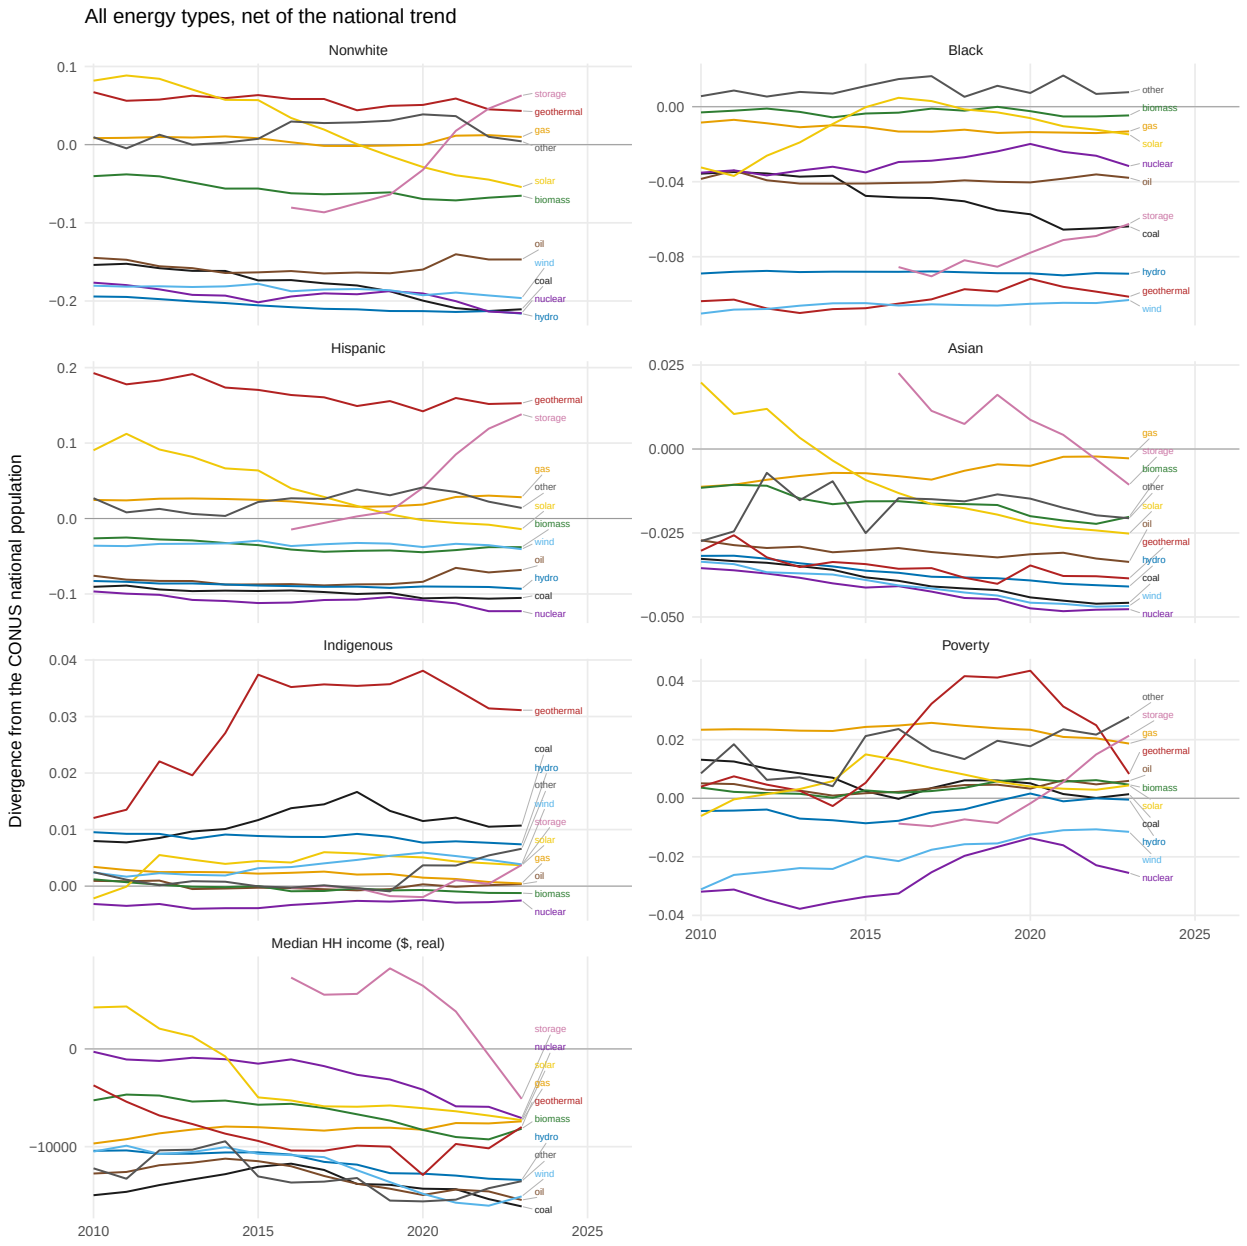

Figure 4.10: Host-community composition by energy type, net of the CONUS national trend, 2010–2023, all eleven energy types; zero is the national level. Income is real 2023 dollars; fuel-years with fewer than thirty plants are omitted.

4.9.4 Decomposition: all fuels and both windows

Figure 4.4 shows the four transition-core fuels for the primary 2010–2019 window. Figure 4.11 gives every energy type, for both the primary window and the 2010–2023 extension. Degenerate fuels (no plants surviving the full window — storage in particular) have no channel bars and are

shown by their total alone.

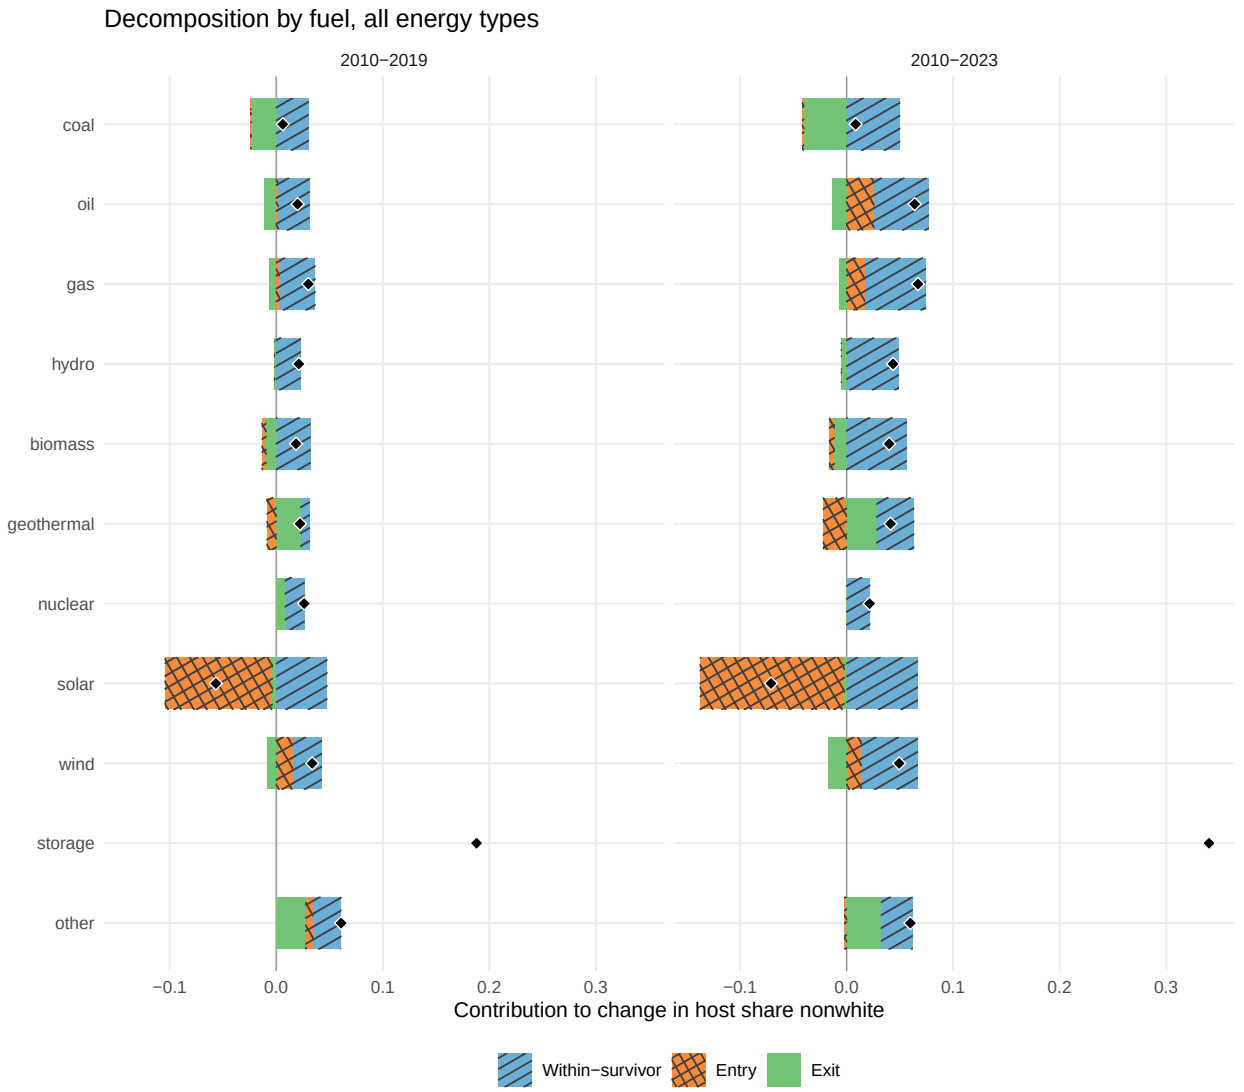

Figure 4.11: Melitz–Polanec decomposition of the change in host-community share nonwhite (3 km), all eleven energy types, for the primary 2010–2019 window and the 2010–2023 extension; channels signed, the black diamond the total. Fuels with no surviving plants across a window carry a total but no channel bars.

4.9.5 Entrant composition: all fuels and outcomes

Figure 4.5 shows the nonwhite gap for the substantial fuels. Figure 4.12 gives every fuel and every demographic outcome, each total entrant gap split into its within-state and between-state components. The high-Indigenous outcome is a linear-probability model; income is in log points. Nuclear (no

in-window entrants) is absent.

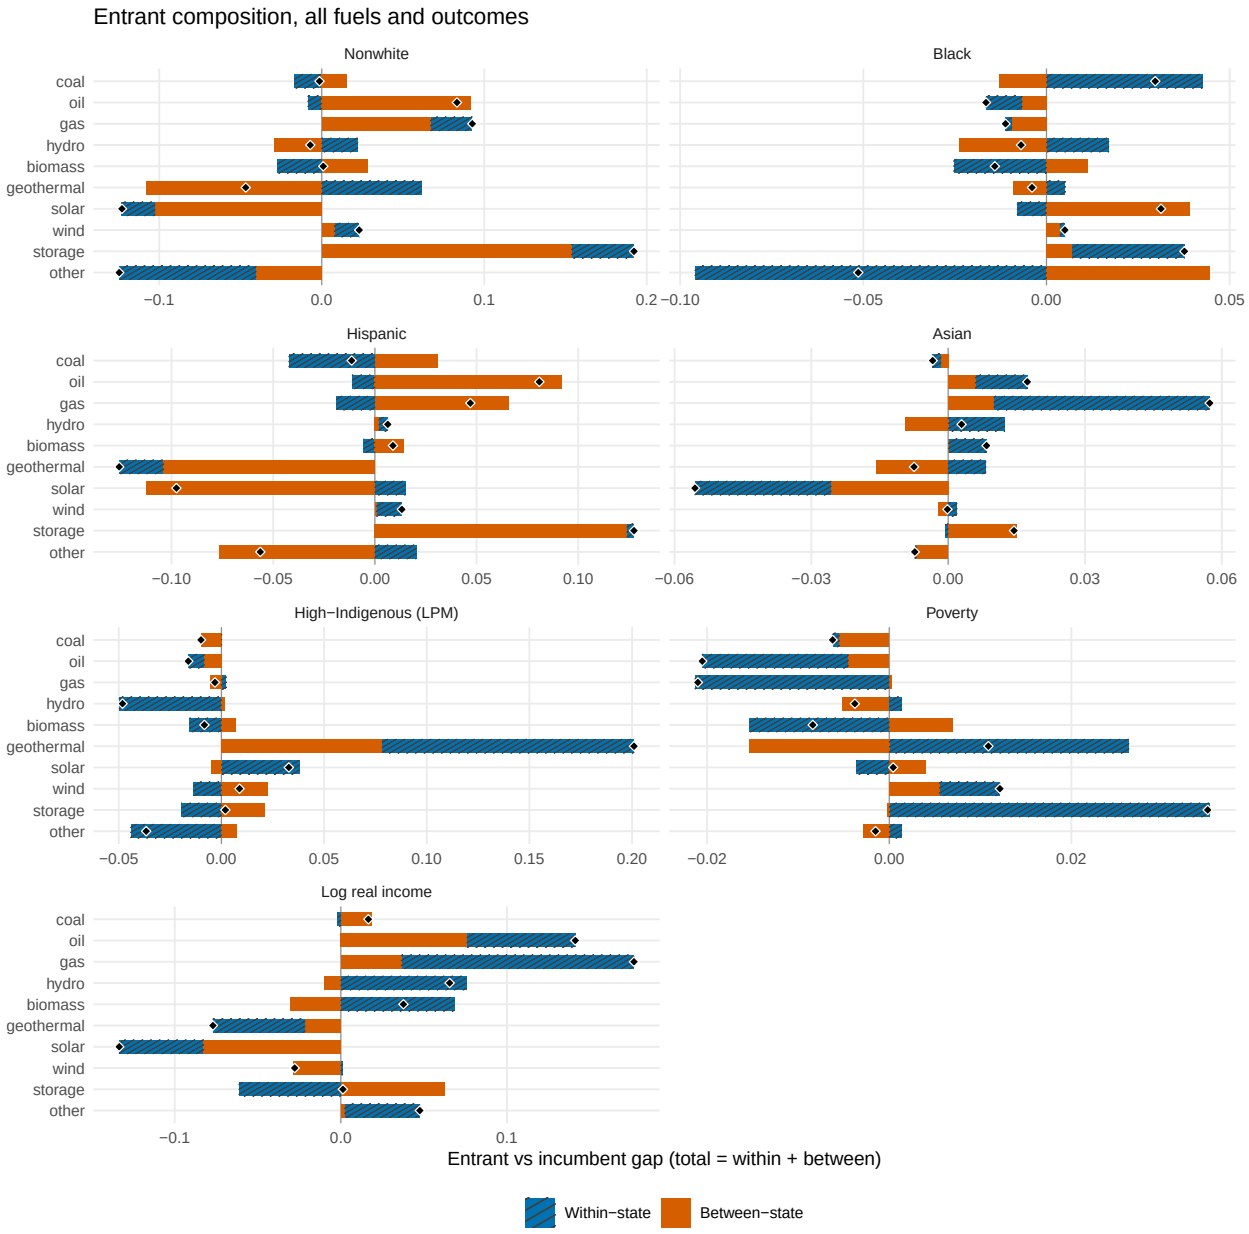

Figure 4.12: Entrant composition for all eleven energy types and all outcomes: each total entrant-versus-incumbent gap split into within-state and between-state components summing to the total (black diamond). Panels are free-scaled; the high-Indigenous panel is a linear-probability model and income is in log points. SEs are wide (see text).

4.9.6 Distributional context: solar host communities over time

The trajectory in Figure 4.2 reports solar's *mean* host-community composition, which falls across the decade. This appendix asks what the underlying distribution is doing: whether the whole mass of solar communities slides toward the national average, or a distinct lower-nonwhite (rural) mode emerges beneath the early Southwest concentration as the buildout disperses geographically (Figure 4.13).

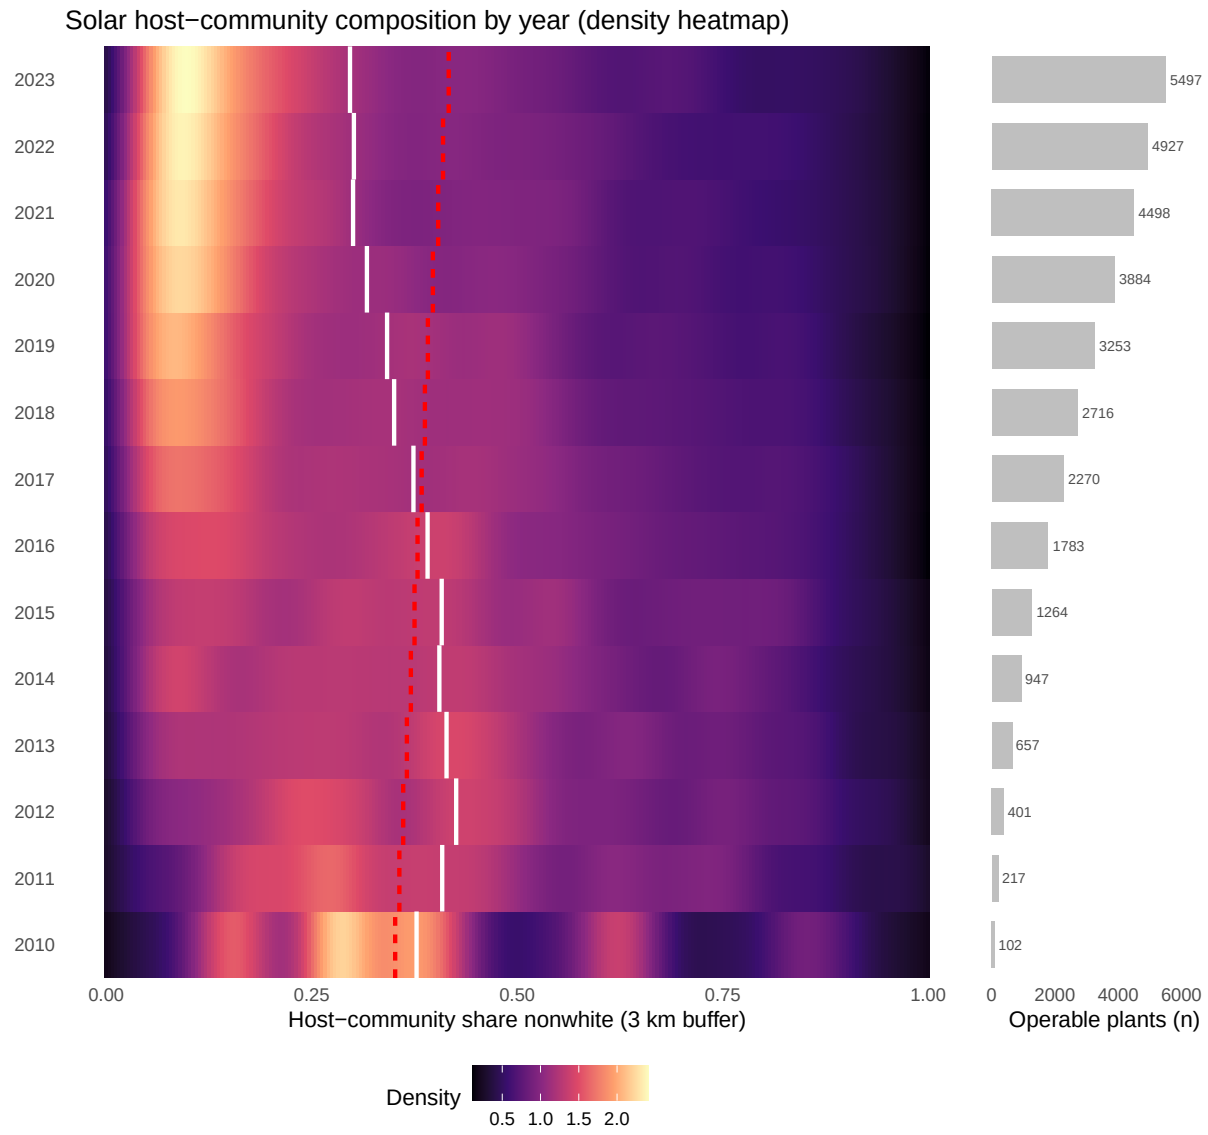

Figure 4.13: Per-year kernel density of solar host-community share nonwhite (3 km buffer), 2010–2023, as a heatmap (brighter = denser; left), with the operable-plant count each year (right). The solid white line is the annual median; the dashed red line the contemporaneous CONUS national share. Bandwidth 0.035 on a 256-point grid; a 30-plant floor binds on no year.

4.9.7 Robustness: buffer radius

The trajectory's headline, solar's crossing from more-nonwhite-than-nation to whiter-than-nation, is robust to the buffer radius (Figure 4.14). Widening the buffer from three to five to ten kilometres barely moves the series: the crossing holds in 2019 at every radius, and the endpoints shift by less

than a point, so the finding is not an artefact of the particular distance chosen. (The distinct question of *weighting*, counting plants versus counting people, is taken up in the main text at Figure 4.8, where population-weighting is shown to shift the fleet's host communities systematically more nonwhite.)

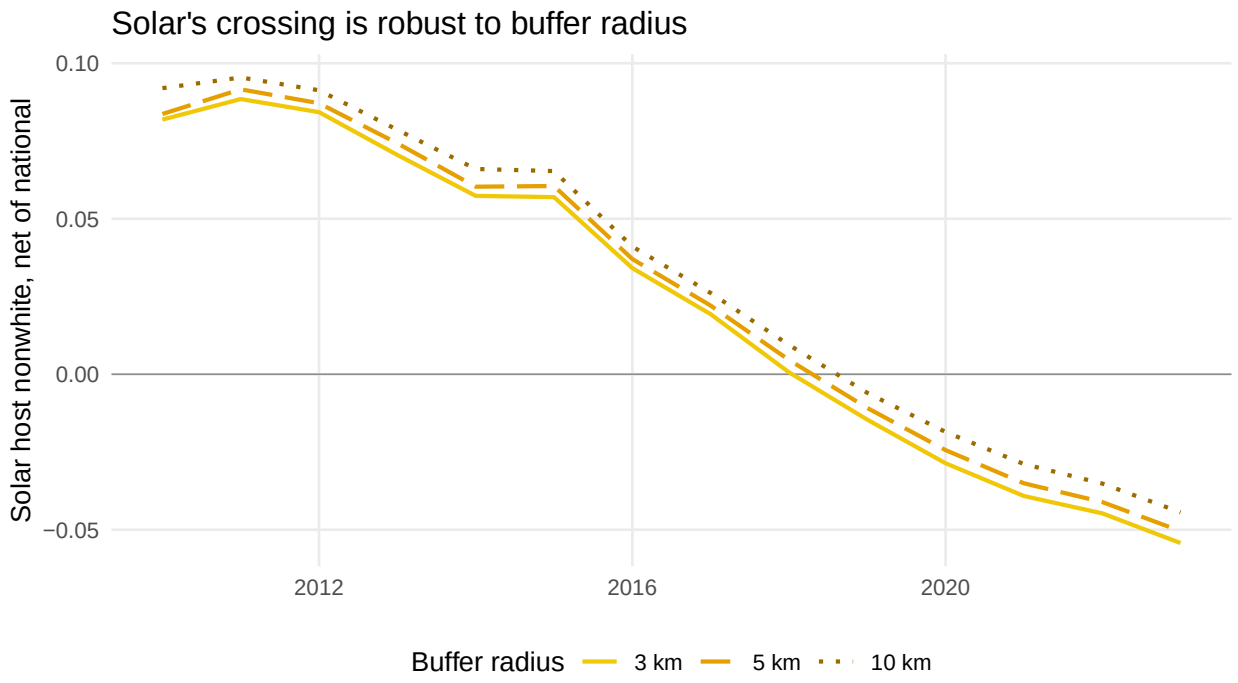

Figure 4.14: Solar's host-community share nonwhite, net of the CONUS national population, at 3, 5, and 10 km buffers.

4.9.8 Full results tables

The figures above report the entrant-composition and decomposition results graphically; the tables here give the underlying numbers. Table 4.10 and Table 4.11 report the entrant effects for the total (year fixed effects only) and within-state (adding state fixed effects) specifications; each cell is the per-fuel entrant-versus-incumbent gap, in percentage points, with CR1 (state-clustered) standard errors in parentheses. Their difference is the between-state component discussed in the text. The reading is descriptive: because the entrant gap is between-state while the estimator clusters on state, the standard errors are wide and no significance is asserted, the lone exception being solar on the Indigenous dimension. Table 4.12 gives the Melitz-Polanec channels.

Table 4.10: Entrant composition: total entrant-versus-incumbent gap (year fixed effects only), by fuel and outcome.

	Hi-Indig.	Ln inc.	Asian	Black	Hispanic	Nonwhite	Poverty
Coal	-1.0 (3.6)	1.7 (11.6)	-0.3 (0.3)	3.0 (7.4)	-1.1 (1.7)	-0.1 (6.8)	-0.6 (2.7)
Oil	-1.6 (0.7)	14.1 (6.8)	1.7 (0.8)	-1.7 (1.7)	8.1 (6.5)	8.3 (7.6)	-2.1 (1.7)
Gas	-0.3 (1.0)	17.6 (4.8)	5.7 (2.6)	-1.1 (1.2)	4.7 (1.5)	9.3 (3.2)	-2.1 (0.6)
Hydro	-4.8 (1.4)	6.5 (4.9)	0.3 (0.4)	-0.7 (1.2)	0.6 (2.4)	-0.7 (2.8)	-0.4 (1.5)
Biomass	-0.8 (0.5)	3.8 (3.6)	0.8 (0.8)	-1.4 (1.6)	0.9 (1.0)	0.1 (1.9)	-0.8 (0.7)
Geothermal	20.1 (9.8)	-7.7 (9.6)	-0.7 (0.5)	-0.4 (0.5)	-12.6 (6.0)	-4.7 (8.3)	1.1 (2.4)
Solar	3.3 (1.3)	-13.3 (6.3)	-5.6 (1.0)	3.1 (3.1)	-9.8 (2.9)	-12.3 (4.3)	0.0 (1.3)
Wind	0.9 (1.4)	-2.8 (2.0)	-0.0 (0.1)	0.5 (0.4)	1.3 (2.1)	2.3 (2.2)	1.2 (0.8)
Storage	0.2 (2.4)	0.1 (8.3)	1.4 (1.4)	3.8 (1.2)	12.7 (4.9)	19.2 (6.8)	3.5 (1.9)
Other	-3.7 (3.6)	4.8 (7.8)	-0.7 (2.6)	-5.1 (2.5)	-5.6 (7.8)	-12.5 (9.4)	-0.2 (2.1)

Note. Cells are per-fuel entrant marginal effects in percentage points (log points for income), CR1 state-clustered SEs in parentheses. Nuclear omitted (no in-window entrants). Descriptive.

Table 4.11: Entrant composition: within-state entrant-versus-incumbent gap (year and state fixed effects), by fuel and outcome; cells as in the previous table.

	Hi-Indig.	Ln inc.	Asian	Black	Hispanic	Nonwhite	Poverty
Coal	0.0 (2.9)	-0.2 (8.7)	-0.2 (0.3)	4.3 (6.3)	-4.2 (3.7)	-1.7 (7.4)	-0.1 (2.1)
Oil	-0.8 (0.9)	6.5 (4.6)	1.1 (0.8)	-1.0 (1.5)	-1.1 (2.6)	-0.9 (4.1)	-1.6 (1.0)
Gas	0.2 (0.7)	14.0 (4.4)	4.7 (2.2)	-0.2 (0.8)	-1.9 (1.1)	2.6 (1.6)	-2.1 (0.7)
Hydro	-5.0 (1.7)	7.6 (4.3)	1.2 (0.8)	1.7 (1.0)	0.4 (2.0)	2.2 (2.9)	0.1 (1.5)
Biomass	-1.6 (0.5)	6.8 (2.8)	0.8 (0.6)	-2.5 (1.2)	-0.5 (0.7)	-2.7 (1.5)	-1.5 (0.6)
Geothermal	12.3 (10.4)	-5.5 (9.3)	0.8 (1.4)	0.5 (0.3)	-2.2 (2.9)	6.1 (4.5)	2.6 (1.9)
Solar	3.8 (1.3)	-5.1 (3.3)	-3.0 (0.5)	-0.8 (1.5)	1.5 (2.2)	-2.0 (3.4)	-0.4 (0.9)
Wind	-1.4 (1.0)	0.1 (2.6)	0.2 (0.2)	0.1 (0.4)	1.2 (1.4)	1.5 (1.5)	0.7 (0.8)
Storage	-1.9 (2.2)	-6.1 (9.2)	-0.1 (1.1)	3.1 (1.0)	0.3 (3.2)	3.8 (5.1)	3.5 (2.4)
Other	-4.4 (3.6)	4.5 (7.6)	-0.0 (1.7)	-9.6 (4.3)	2.0 (5.8)	-8.4 (8.1)	0.1 (1.8)

Table 4.12: Melitz–Polanec decomposition of the change in host-community share nonwhite (3 km) by fuel and channel, for the 2010–2019 window and the 2010–2023 extension.

	n(surv)	n(entry)	n(exit)	Total	Within	Entry	Exit
2010-2019							
Coal	254	18	262	0.6	3.0	0.0	-2.4
Oil	389	75	168	2.0	2.9	0.2	-1.1
Gas	1136	556	224	3.0	3.4	0.3	-0.6
Hydro	1282	73	59	2.1	2.3	0.0	-0.2
Biomass	427	257	145	1.9	3.2	-0.4	-1.0
Geothermal	47	14	9	2.2	0.8	-0.9	2.3
Nuclear	57	0	7	2.6	1.8	0.0	0.8
Solar	92	3161	10	-5.7	4.8	-10.1	-0.3
Wind	508	607	51	3.4	2.8	1.5	-0.9
Storage	0	91	1	18.8	—	—	—
Other	28	15	7	6.1	2.7	0.7	2.8
2010-2023							
Coal	188	14	328	0.8	5.0	-0.1	-4.0
Oil	333	130	207	6.4	5.1	2.6	-1.3
Gas	1076	755	272	6.7	5.7	1.7	-0.7
Hydro	1229	84	99	4.4	4.9	0.0	-0.5
Biomass	351	226	217	4.0	5.6	-0.5	-1.2
Geothermal	46	17	10	4.1	3.5	-2.2	2.8
Nuclear	52	0	11	2.2	2.2	0.0	0.0
Solar	88	5409	14	-7.1	6.7	-13.6	-0.2
Wind	475	823	85	4.9	5.3	1.4	-1.7
Storage	0	261	1	34.0	—	—	—
Other	24	17	11	6.0	2.9	-0.2	3.2

Note. Percentage points; channels sum to the total. n(surv/entry/exit) are the survivor, entrant, and exiter counts; fuels with no survivors show total only.

4.9.9 Robustness: functional form

The trajectory and the decomposition operate on the mean share and so require the linear coding; the entrant-composition regression does not, and its headline (solar's entrants whiter than the incumbent fleet, gas's more nonwhite) is not an artefact of treating percent nonwhite as a continuous

share. Table 4.13 re-estimates the entrant gap under two alternative codings of the outcome: a majority-minority indicator (host share nonwhite above one-half) and a top-tertile indicator, both linear-probability models. The sign and the fuel-ranking are preserved throughout (solar negative, gas positive, wind and coal near zero under every coding), with the binary and tertile magnitudes running a little larger in absolute terms, as coarsened outcomes tend to, but never reversing.

Table 4.13: Entrant-versus-incumbent gap in host-community race (full specification, percentage points), for the four transition-core fuels, under three codings of the outcome: linear share nonwhite (primary), a majority-minority indicator (share > 0.5), and a top-tertile indicator; linear-probability models for the binaries.

	Majority-minority	Linear (primary)	Top tertile
coal	0.4	-0.1	-2.4
gas	15.0	9.3	15.8
solar	-14.2	-12.3	-18.5
wind	2.6	2.3	2.6

4.9.10 Robustness: Indigenous proximity threshold

The Indigenous measure enters the analysis as a high-proximity threshold rather than a smooth gradient (Section 4.9.1); Table 4.14 shows the choice of cutoff does not drive the result. The entrant gap for the high-Indigenous outcome holds its sign across the primary five-percent cutoff, a stricter ten-percent cutoff, and the continuous share: new solar sits modestly *more* often near substantially Native communities (a small counter-current beneath solar's overall whitening, consistent with utility-scale siting on the public lands of the interior West), while new coal and gas, the few that enter, sit slightly less often near them. The magnitudes are small and the continuous specification is the faintest of the three, but nothing flips with the threshold.

Table 4.14: Entrant-versus-incumbent gap in high-Indigenous proximity (full specification, percentage points), for the four transition-core fuels, at the primary >=5% cutoff, a >=10% cutoff, and the continuous non-Hispanic AIAN share.

	≥5% (primary)	≥10%	Continuous
coal	-1.0	-2.8	-1.1
gas	-0.3	-1.1	-0.3

solar	3.3	2.3	0.7
wind	0.9	0.7	0.3

4.9.11 Robustness: natural-gas prime mover

The energy-type classification is a capacity-weighted dominant fuel, so it cannot separate the combustion-turbine peakers long noted to sit near poor and minority communities (Krieger, Casey, and Shonkoff 2016) from the combined-cycle plants that carry gas's baseload. Splitting natural gas by prime mover shows that the weak gas replication the main analysis reported is, in fact, a peaker phenomenon. New peaking capacity sits in communities some 14 points more nonwhite than the incumbent fleet, and, decisively, it retains a within-state gap of 3.6 points, the local-targeting signature the fossil path of least resistance predicts. Combined-cycle entrants run the other way: 7 points *whiter* than incumbents, and 1.8 points whiter even within state. The decomposition tells the same story through the entry channel (Figure 4.15): over 2010–2023 the entry of new peakers pulls their host mean more nonwhite by 3.7 points, while new combined-cycle capacity pulls its mean whiter by 1.6. The lumped gas result of the main text is thus a blend of two opposed siting logics, and it is the peaker (the emitting, quick-start unit closest in kind to the fossil fleet) that reproduces the fossil pattern, on the emitting margin where the theory expected it.

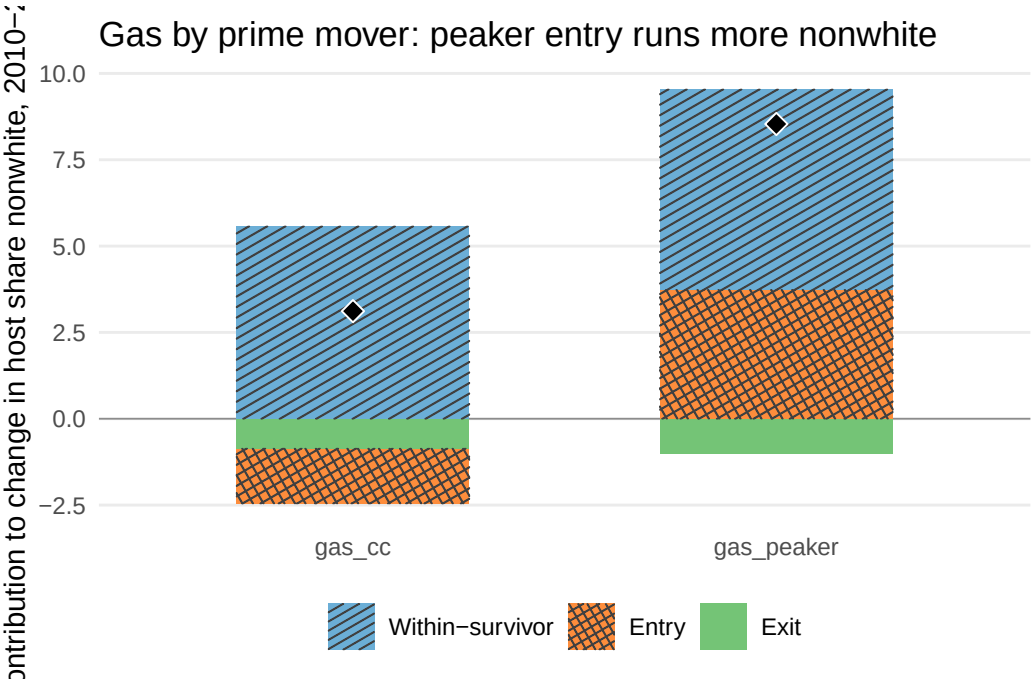

Figure 4.15: Melitz–Polanec decomposition of the change in host-community share nonwhite (3 km), 2010–2023, with natural gas split by prime mover into combined-cycle baseload (gas_cc) and simple-cycle peaking (gas_peaker); channels signed (percentage points), the black diamond the total.

5 Conclusion

Organizational Structure, Environmental Justice, and the Geography of the Green Transition in the U.S. Power Sector

5.1 The distributive question, answered

This dissertation opened with a distributive question: How do electricity generating facilities' organizational and community contexts shape both the burden of fossil-fuel pollution and the distribution of clean-energy benefits, and do those burdens and benefits fall equitably across lines of race, income, and geography? Collectively, the three chapters return an unexpected answer. In the fossil fleet, disparity does not emerge where the environmental justice (EJ) theoretical literature predicted, in who gets pollution-control equipment: neither organizational structure nor community demographics predicts adoption on either margin. Disparity appears to emerge downstream instead, in how plants are run once equipped. Beyond the fossil fleet, the transition is redrawing the spatial demography of energy generation fuel by fuel, and, by and large, via the turnover of the fleet rather than the churn of people through neighborhoods; for renewables specifically, that demographic transition is occurring at the scale of the state rather than the locale. In sum, and answering the question with which the introduction closed: the burden of the fossil capacity still in service concentrates, increasingly and measurably, on the more nonwhite communities beside the aging urban units the transition keeps alive, while the benefits of non-emitting generation land where the buildout lands, in the whiter states toward which solar is dispersing.

5.2 Findings in review

5.2.1 The Organization of Pollution Control

Chapter 2 asked who cleans up, and found the question misdirected. Organizational structure barely predicts whether a plant adopts NO_x control on either functional margin. The lone adoption signal is government ownership: public owners adopt the capital-intensive post-combustion option more readily, the pattern cheap public capital would produce, yet they are no cleaner per MWh for it. Organizational structure impacts emission outcomes downstream of adoption. Rate-regulated investor-owned utilities and cooperatives run markedly dirtier per MWh than merchant generators, and the penalty survives conditioning on the controls a plant carries, the implication being that the difference lies in how equipped plants are run. Chapter 2 settles on the reading that cost-of-service regulation pulls in two directions at once. Rate-base recovery, a regulatory mechanism allowing rate-regulated utilities to earn a legally authorized return on capital expenditures and infrastructure investments (incentivizing over-capitalization, i.e., 'gold-plating'), lowers the cost of abatement capital and rewards capital spending generally (Averch and Johnson 1962). But cost pass-through, another regulatory mechanism allowing rate-regulated utilities to pass operating costs through to ratepayers, blunts the operating-cost discipline a merchant faces directly. On the NO_x-per-MWh margin, the pass-through effect dominates the rate-base recovery financing effect. Furthermore, the cooperative result helps locate the mechanism more precisely: cooperatives recover costs like IOUs but hold no rate-base return, though they still carry the same penalty. Ownership dispersion adds a parallel channel, a modest but reliable per-owner penalty consistent with principal–agent slack. The 'super-polluter' expectation leaves no mark on intensity: plants owned by larger firms are no dirtier per MWh. The null is less surprising than it first appears. The disproportionality claim that animates the frame (Freudenburg 2006; Prechel and Istvan 2016) concerns the mass a large firm's generation commands, not the rate at which any one plant pollutes. But the per-MWh test asked a further question: whether centralized control over large pools of generation bred its own principal–agent slack, and a rate penalty to match. The answer was no.

5.2.2 Environmental Justice at the Operating Margin

Chapter 3 carried the environmental justice question to similar analytical machinery and received the same redirection. Adoption of NO_x emissions control technologies shows no demographic gradient: plants in more nonwhite or poorer communities are no less likely to carry control on either margin, combustion or post-combustion, and the capital-intensity differential (the likelihood of capital-intensive, more effective post-combustion control versus cheaper, less effective combustion modification) the design was built to detect is nil. Absent at adoption, the disparity appears instead in operation — in NO_x intensity per MWh — and it is recent. By 2021, though not in 2009 or 2018, plants in more nonwhite communities run measurably dirtier per MWh, conditional on the controls they carry and on how hard they run. The gradient is a racial one, not a poverty gradient in disguise: it holds with income and poverty entered separately, and it survives spatial correction. It appears only in the emission rate; on aggregate tonnage the association vanishes. The mechanism Chapter 3 deductively proposed is compositional. As the coal fleet retired, the dirtiest-per-MWh unit relocated from the large rural coal plant to the small, aging urban peaker. The peaker starts and stops with daily demand, and so spends a disproportionate share of its firing hours in start-up transients and low-load running, where flue gas falls below the temperature window post-combustion controls require. And it sits where the grid's energy demands dictate — inside the transmission-constrained metropolitan load pockets that must meet part of their peak demand locally — and therefore in the more nonwhite communities those pockets contain. In 2009 the highest-intensity decile sat in communities whiter than the fleet average; by 2021 it had crossed into more nonwhite ones. Chapter 3 identified an emerging, operational face of environmental injustice, distinct from the siting disparities the proximity literature has long documented (Bullard 1990; Mohai, Pellow, and Roberts 2009).

5.2.3 The Unequal Geography of the Green Transition

Chapter 4 expanded its scope to ask where the transition is actually unfolding, finding the spatial demography of U.S. electricity generation being redrawn fuel by fuel, by fleet turnover, and, for the renewables, between states. Solar's host communities cross the continental US grand mean outright, from more nonwhite than the nation to whiter, as the buildout disperses into whiter states. Within

any given state, new solar lands in communities essentially indistinguishable from the solar already there, so the whitening is relocation, not local targeting. Coal's relative whitening is exit-driven, and it runs against the registered prior: the coal that retired sat in more nonwhite communities than the coal that survives, so retirement whitened the surviving fleet's average by subtraction alone. Gas, predictably, is the lone trace of within-state fossil replication, its entrants modestly more nonwhite than its incumbents with the state held fixed, and sharpest among peakers. However, two caveats bound this synthesis: solar's whitening is an average beneath which a real and growing share of the buildout lands in the Black-rural Southeast, and the fleet is whiter counted by plant than by person, since weighting each plant by the population around it turns nearly every fuel's host communities more nonwhite.

Scored against the trichotomy fixed before estimation, the results land close to the registered expectations, with one instructive miss. Gas replicated the fossil siting gradient, as predicted: its entrants land in modestly more nonwhite communities than its incumbents even with the state held fixed, the lone within-state trace of the path of least resistance, and sharpest among peakers. Solar, expected to sit between replication and transformation, delivered transformation outright, though of a particular kind: the demographic association did not so much weaken as move its axis, from the racial composition of host communities to the state-scale geography of resource and land into which the buildout is dispersing. Coal moved through exit, as expected, but in the direction the prior did not anticipate: the coal that retired sat in more nonwhite communities than the coal that survives. The null, finally, holds in part but not in whole: much of where new capacity lands is fixed by resource potential and project viability before demography enters at all, yet demographic patterning survives that filter, in gas's within-state gradient and in the demographic signature of solar's between-state dispersal.

5.3 What the three chapters say together

5.3.1 From adoption to operation

Each tradition predicted its own gradient in who adopts control: organizational structure for the political economy literature, community disadvantage for environmental justice. Neither gradient

shows up at adoption. Both surface downstream, in how plants are operated. The operating margin, which neither tradition has yet made a central object, carries the signal in both cases: for organizational form in Chapter 2, for race in Chapter 3. The two-margin design, then, did more than test the traditions' hypotheses: its results argue for re-aiming their shared question at that margin. Who adopts control is, by the 2010s, largely settled by regulation and engineering. The live question, the one on which organization and inequality actually register, is: 'Who runs dirtier given the same controls?' The standard policy instruments, the subsidies and mandates that put equipment on site, are not built to answer it; even the NO_x trading programs (the NO_x Budget Program and its successors, the Clean Air Interstate Rule and the Cross-State Air Pollution Rule), the one instrument class aimed squarely at operations, reach operational emissions only in aggregate, for reasons the policy implications section takes up. And rising demand — data centers, electrification, growth itself — spares the highest-rate units on the other flank, deferring retirement, the only mechanism capable of improving the actual *composition* of the fleet.

The eco-structural premises the introduction set out fare unevenly, and instructively. The first, that production organizations are the proximate agents of harm, is vindicated in a form stronger than the literature usually claims: organizational identity registers not in what plants carry but in how the same equipment is run. The second, that ownership, financing, and regulation shape the decisions that determine harm, holds, though the operative mechanism proved to be cost pass-through rather than the cost of capital. The third, that harm is concentrated, held for mass but not for rate: the heaviest emitters by tonnage are not the dirtiest per megawatt-hour, and the super-polluter expectation left no mark on intensity, so the premise governs the fleet's total burden while the rate margin obeys a different logic.

5.3.2 The peaker at the junction

The three accounts converge on a single kind of plant: the small, aging, urban gas peaker. This kind of peaker is the unit form whose operating regime defeats the controls it carries, given the frequency at which its operating temperatures fluctuate outside the range in which NO_x control technologies are effective. The same unit carries Chapter 3's emergent racial gradient, sitting in the metropolitan load pockets whose demographics it inherits. And it supplies Chapter 4's lone within-state trace

of fossil replication, gas entrants landing modestly more nonwhite than gas incumbents, sharpest on the peaking margin. The peaker is what fleet turnover leaves behind. It is also what reliability requirements keep alive: paid to stand ready (Alliance et al. 2020), dispatched in the stressed hours when ozone chemistry can be most dangerous, and sited where the load is, which is to say where the people are, and disproportionately where the nonwhite and low-income people are. If this dissertation has an emblematic unit, it is the small, aging, urban gas peaker: the unit form in which organizational slack, racial disparity, and the transition's forever-unfinished geography meet.

5.3.3 Turnover as redistribution

Beneath all three chapters runs the same underlying process: the turnover of the electricity generating fleet. Each chapter treated compositional change as a feature to describe rather than a bias to purge. Coal's retirement drove the relocation of the dirty tail of high-rate emitters that produced Chapter 3's gradient, and whitened the surviving coal fleet by subtraction in Chapter 4. The entry of gas and renewables is meanwhile drawing the next map: gas reproducing a trace of the fossil gradient within states, solar redrawing it between them. Energy transitions, the historical record suggests, have so far been additions rather than substitutions (York and Bell 2019), and this one is following the pattern, driven forward by the global technoeconomic growth engine (data centers above all, stacking large new loads atop a base that is itself growing (Shehabi et al. 2024)), but also by decarbonization itself. Electrification, the backbone of decarbonization policy, raises electricity demand by design. All of these sources of rising demand slow retirement, and slowed retirement keeps the aging, high-rate units of Chapter 3's story online and running harder. Entry accelerates while exit stalls, in other words. The burden associated with electricity generation will not disappear with the coal fleet; it is emerging anew with the rise of natural gas as well as being reapportioned: concentrated per-MWh on the surviving and interstitial fossil units and the communities beside them, even as its aggregate mass declines, while the benefits of the buildout accrue on a geography whose demographic signature is being set state by state.

5.4 Policy implications

The findings carry three implications for policy. The first concerns where the existing instruments aim. The standing toolkit targets adoption and aggregate mass: performance standards put equipment on site, and the NO_x trading programs cap the total while reaching operations only in aggregate, their defining flexibility letting the highest-rate units, by definition the costliest to abate, comply by buying allowances rather than by running cleaner. The results here locate the disparity downstream of everything those instruments touch. Reaching it would take instruments aimed at the operating margin itself: unit-level rate standards, or dispatch rules that weigh a unit's emission rate alongside its bid. Chapter 2's merchant contrast is instructive here, though not as a brief for markets: the operating-cost exposure merchants face is associated with cleaner running per MWh, which is to say that operating discipline and organizational slack have real consequences for NO_x emissions, and that the presumption of public or rate-regulated ownership operating in the public interest does not extend automatically to this particular margin.

The second implication concerns entry. The transition multiplies the siting problem rather than shrinking it. A horizontal buildout of wind and solar involves orders of magnitude more host communities than the fossil system it replaces (Bridge et al. (2013); Smil (2010)). Deployment can stall precisely where the capacity to resist is greatest, in whiter and higher-income communities, so the social acceptance of facilities becomes a binding constraint on the transition's pace. The technoeconomic constraints that discipline renewable siting meanwhile bind the portable gas turbine far less, and friction on clean entry is therefore not neutral: it tilts the mix toward the fossil entrant that faces less of it.

The third implication is the broadest: taken together, the chapters argue that the transition itself, and not only its additions, must become the object of policy. Encouraging it requires regulation of both supply and demand in pursuit of three goals at once: encouraging the entry of renewables, discouraging the entry of new fossil capacity, and encouraging the retirement of old. The chapters give each limb an empirical face: entry is redrawing the map state by state, new gas is the one fossil entrant still replicating the old gradient within states, and retirement is the suppressed channel on which the whole distributive ledger turns.

5.5 Limitations and the road forward

Every estimate in this dissertation is associational: organizational structure, community demographics, and siting are not assigned, and the designs pool cross-sections rather than following plants through time. The outcome throughout is the emission *rate*: how cleanly power is made, on average, in NO_x per MWh. What any community actually breathes is a different quantity. For NO_x, though, that caveat is smaller than it seems: the harm arrives in high-intensity bursts, in the very hours when the high-rate units fire and ozone forms. Another important caveat is that the analysis is plant-side, characterizing the average neighborhoods of generators rather than the exposure of populations, though it does offer a brief plant-versus-person comparison to show exactly how much that choice conceals. A significant limitation of the analysis is that firm identity is the EIA reporting entity, not the ultimate corporate parent. The observed entities are mere shards of broader parent–subsidiary structures, whose liability-shielding forms go unobserved here (cf. Grant et al. (2010)).

Though the project's limitations are many, each boundary, appropriately understood, is a research program in and of itself. The tract-side companion design (taking the community rather than the plant as the unit, and asking about absolute burden and the causal ordering of siting versus sorting) is the natural sequel to both Chapters 3 and 4. The operating margin's mechanism could be supplemented with dispatch- and operator-level detail: startup frequency, ramping, fuel quality, and control utilization at the hour scale where both the emissions and the harm actually occur. And the corporate-hierarchy question (whether the within-firm differential treatment this design could not resolve sharpens when subsidiaries are traced to parents) merely waits on the corporate affiliation data with which to build it.

5.6 Coda

Whose air does the operating margin degrade, and where? That question, latent in Chapter 2's central finding, organized everything that followed. Chapter 3 answered it for the standing fleet: increasingly, and measurably by 2021, the air of more nonwhite communities, in the acute hours when the aging urban peaker fires. For the fleet to come, the answer is a map being redrawn by what retires and what is built, at the scale of the state, across a host population growing by thousands of

communities a year. Strictly speaking, there has never been an energy transition (York and Bell 2019). New sources of energy simply stack atop old ones, a paradox already recognized long ago, in which new capacity and new efficiency feeds demand rather than displacing what came before (York 2006). Every addition so far has rearranged the geography of environmental burden. The present buildout, on the evidence here, is doing the same, with retirement running as the suppressed channel: rising demand keeps the old units online beside the new, so the burden the transition was to retire is run overtime instead. It will keep doing so unless the imperative of transition itself becomes an object of policy — the mix, not additions or subtractions alone — a policy that does more than install the new technology, and instead creates the conditions under which the old can retire. That is also where the dissertation's opening frame comes to rest. If climate politics now lives implicitly in sectoral struggles over utilities and infrastructure (Bosworth 2026), then the questions fought out at the facility, who runs dirtier with the same equipment and who hosts what the transition builds and retires, are not adjacent to the climate fight; they are its present form.

References Cited

Aklin, Michaël, and Matto Mildenberger. 2020. "Prisoners of the Wrong Dilemma: Why Distributive Conflict, Not Collective Action, Characterizes the Politics of Climate Change." *Global Environmental Politics* 20 (4): 4–27. https://doi.org/10.1162/glep_a_00578.

Aklin, Michaël, and Johannes Urpelainen. 2013. "Political Competition, Path Dependence, and the Strategy of Sustainable Energy Transitions." *American Journal of Political Science* 57 (3): 643–58. https://doi.org/10.1111/ajps.12002.

Alliance, NYC Environmental Justice, New York Lawyers for the Public Interest, THE POINT CDC, UPROSE, and Clean Energy Group. 2020. "Dirty Energy, Big Money: How Private Companies Make Billions from Polluting Fossil Fuel Peaker Plants in New York City's Environmental Justice Communities—and How to Create a Cleaner, More Just Alternative." PEAK Coalition.

Allison, Paul D. 1999. "Comparing Logit and Probit Coefficients Across Groups." *Sociological Methods & Research*. https://doi.org/10.1177/0049124199028002003.

Anselin, Luc. 2009. "Spatial Regression." In *The SAGE Handbook of Spatial Analysis*.

Anselin, Luc, Anil K. Bera, Raymond Florax, and Mann J. Yoon. 1996. "Simple Diagnostic Tests for Spatial Dependence." *Regional Science and Urban Economics* 26 (1): 77–104. https://doi.org/10.1016/0166-0462(95)02111-6.

"Appeals Court Issues Split Ruling in Michigan Solar Permitting Suit: A State Appeals Panel Has Upheld Most of Michigan's Regulations Limiting Local Control over Renewable Energy Projects." 2026. *Associated Press*, May.

Arel-Bundock, Vincent. 2026. *Marginaleffects: Predictions, Comparisons, Slopes, Marginal Means, and Hypothesis Tests*. https://marginaleffects.com/.

Averch, Harvey, and Leland L. Johnson. 1962. "Behavior of the Firm Under Regulatory Constraint." *The American Economic Review* 52 (5): 1052–69. https://www.jstor.org/stable/1812181.

Bank, World. 2026. "Inflation, Consumer Prices (Annual %)."

Banzhaf, H. Spencer, Lala Ma, and Christopher Timmins. 2019. "Environmental Justice: Establishing Causal Relationships." *Annual Review of Resource Economics* 11 (1): 377–98. https://doi.org/10.1146/annurev-resource-100518-094131.

Banzhaf, H. Spencer, and Randall P. Walsh. 2008. "Do People Vote with Their Feet? An Empirical Test of Tiebout's Mechanism." *American Economic Review*. https://doi.org/10.1257/aer.98.3.843.

Been, Vicki. 1994. "Locally Undesirable Land Uses in Minority Neighborhoods: Disproportionate Siting or Market Dynamics?" *The Yale Law Journal* 103 (6): 1383. https://doi.org/10.2307/797089.

Been, Vicki, and Francis Gupta. 1997. "Coming to the Nuisance or Going to the Barrios? A Longitudinal Analysis of Environmental Justice Claims." *Ecology Law Quarterly* 24 (1): 1–56.

Berge, Laurent. 2026. *Fixest: Fast Fixed-Effects Estimations*. https://lrberge.github.io/fixest/.

Betz, Timm, Scott J. Cook, and Florian M. Hollenbach. 2018. "On the Use and Abuse of Spatial Instruments." *Political Analysis* 26 (4): 474–79. https://doi.org/10.1017/pan.2018.10.

Bivand, Roger. 2026. *Spdep: Spatial Dependence: Weighting Schemes, Statistics*. https://github.com/r-spatial/spdep/.

Borenstein, Severin, and James Bushnell. 2015. "The US Electricity Industry After 20 Years of Restructuring." *Annual Review of Economics* 7 (1): 437–63. https://doi.org/10.1146/annurev-economics-080614-115630.

Bosworth, Kai. 2026. "Where Is the U.S. Climate Movement?" *Verso*.

Boyce, James K., and Manuel Pastor. 2013. "Clearing the Air: Incorporating Air Quality and Environmental Justice into Climate Policy." *Climatic Change* 120 (4): 801–14. https://doi.org/10.1007/s10584-013-0832-2.

Breen, Richard, Kristian Bernt Karlson, and Anders Holm. 2018. "Interpreting and Understanding Logits, Probits, and Other Nonlinear Probability Models." *Annual Review of Sociology* 44 (1): 39–54. https://doi.org/10.1146/annurev-soc-073117-041429.

Breetz, H., M. Mildenberger, and L. Stokes. 2018. "The Political Logics of Clean Energy Transitions." *Business and Politics* 20 (4): 492–522. https://doi.org/10.1017/bap.2018.14.

Bridge, Gavin, Stefan Bouzarovski, Michael Bradshaw, and Nick Eyre. 2013. "Geographies of Energy Transition: Space, Place and the Low-Carbon Economy." *Energy Policy* 53: 331–40. https://doi.org/10.1016/j.enpol.2012.10.066.

Brulle, Robert J. 2018. "The Climate Lobby: A Sectoral Analysis of Lobbying Spending on Climate Change in the USA, 2000 to 2016." *Climatic Change* 149 (3-4): 289–303. https:

//doi.org/10.1007/s10584-018-2241-z.

Bullard, Robert D. 1990. *Dumping in Dixie: Race, Class, and Environmental Quality*. Boulder: Westview Press.

Bureau, US Census. 2026a. "American Community Survey 5-Year Data (2009-2024)." Government. *Census.gov*. https://www.census.gov/data/developers/data-sets/acs-5year.html; US Census Bureau.

———. 2026b. "TIGER/Line Shapefiles." Government. *Census.gov*. https://www.census.gov/geographies/mapping-files/time-series/geo/tiger-line-file.html; US Census Bureau.

Carley, Sanya, and David M. Konisky. 2020. "The Justice and Equity Implications of the Clean Energy Transition." *Nature Energy*. https://doi.org/10.1038/s41560-020-0641-6.

Carley, Sanya, David M. Konisky, Zoya Atiq, and Nick Land. 2020. "Energy Infrastructure, NIMBYism, and Public Opinion: A Systematic Literature Review of Three Decades of Empirical Survey Literature." *Environmental Research Letters* 15 (9): 093007. https://doi.org/10.1088/1748-9326/ab875d.

Casey, Joan A., Jason G. Su, Lucas R. F. Henneman, Corwin Zigler, and Al Et. 2020. "Improved Asthma Outcomes Observed in the Vicinity of Coal Power Plant Retirement, Retrofit and Conversion to Natural Gas." *Nature Energy*. https://doi.org/10.1038/s41560-020-0600-2.

Chakraborty, Jayajit, Juliana A. Maantay, and Jean D. Brender. 2011. "Disproportionate Proximity to Environmental Health Hazards: Methods, Models, and Measurement." *American Journal of Public Health*. https://doi.org/10.2105/ajph.2010.300109.

Christophers, Brett. 2022. *Rentier Capitalism: Who Owns the Economy, and Who Pays for It?*

London New York: Verso.

———. 2024. *The Price Is Wrong: Why Capitalism Won't Save the Planet*. London New York: Verso.

Clark, Lara P, Dylan B Millet, and Julian D Marshall. 2014. "National Patterns in Environmental Injustice and Inequality: Outdoor NO2 Air Pollution in the United States."

Clay, Karen, Danae Hernandez-Cortes, Akshaya Jha, Joshua A Lewis, Noah S Miller, and Edson R Severnini. 2025. "The Social Lifecycle Impacts of Power Plant Siting in the Historical United States."

Colgan, Jeff D., Jessica F. Green, and Thomas N. Hale. 2021. "Asset Revaluation and the Existential Politics of Climate Change." Journal {{Article}}. *International Organization* 75 (2): 586–610. https://doi.org/10.1017/s0020818320000296.

Conley, T. G. 1999. "GMM Estimation with Cross Sectional Dependence." *Journal of Econometrics* 92 (1): 1–45. https://doi.org/10.1016/S0304-4076(98)00084-0.

Cranmer, Zana, Laurel Steinfield, Javier Miranda, and Taryn Stohler. 2023. "Energy Distributive Injustices: Assessing the Demographics of Communities Surrounding Renewable and Fossil Fuel Power Plants in the United States." *Energy Research & Social Science* 100 (June): 103050. https://doi.org/10.1016/j.erss.2023.103050.

Crowder, Kyle, and Liam Downey. 2010. "Interneighborhood Migration, Race, and Environmental Hazards: Modeling Microlevel Processes of Environmental Inequality." *American Journal of Sociology*. https://doi.org/10.1086/649576.

Curran, Patrick J, and Stephen G West. 1996. "The Robustness of Test Statistics to Nonnormality

and Specification Error in Confirmatory Factor Analysis." *Psychological Methods*. https://doi.org/1082-989X/96/$3.00.

Currie, Janet, John Voorheis, and Reed Walker. 2023. "What Caused Racial Disparities in Particulate Exposure to Fall? New Evidence from the Clean Air Act and Satellite-Based Measures of Air Quality." *American Economic Review* 113 (1): 71–97. https://doi.org/10.1257/aer.20191957.

Cushing, Lara J., Shiwen Li, Benjamin B. Steiger, and Joan A. Casey. 2023. "Historical Red-Lining Is Associated with Fossil Fuel Power Plant Siting and Present-Day Inequalities in Air Pollutant Emissions." *Nature Energy* 8 (1): 52–61. https://doi.org/10.1038/s41560-022-01162-y.

Declet-Barreto, Juan, and Andrew A. Rosenberg. 2022. "Environmental Justice and Power Plant Emissions in the Regional Greenhouse Gas Initiative States." Edited by Diane Sicotte. *PLOS ONE* 17 (7): e0271026. https://doi.org/10.1371/journal.pone.0271026.

Devine-Wright, Patrick. 2009. "Rethinking NIMBYism: The Role of Place Attachment and Place Identity in Explaining Place-Protective Action." *Journal of Community & Applied Social Psychology*. https://doi.org/10.1002/casp.1004.

Donovan, Charles, and Christopher Corbishley. 2016. "The Cost of Capital and How It Afffects Climate Change Mitigation Investment." Briefing Paper No. 15. Imperial College UK.

Downey, Liam. 2005. "Assessing Environmental Inequality: How the Conclusions We Draw Vary According to the Definitions We Employ." Journal {{Article}}. *Sociological Spectrum* 25 (3): 349–69. https://doi.org/10.1080/027321790518870.

Düben, Christian. 2026. *Conleyreg: Estimations Using Conley Standard Errors*. https://github.com/cdueben/conleyreg.

Ebenstein, Avraham, Maoyong Fan, Michael Greenstone, Guojun He, and Maigeng Zhou. 2017. "New Evidence on the Impact of Sustained Exposure to Air Pollution on Life Expectancy from China's Huai River Policy." *Proceedings of the National Academy of Sciences* 114 (39): 10384–89. https://doi.org/10.1073/pnas.1616784114.

EIA. 2025. "Annual Electric Power Industry Report, Form EIA-860 - U.S. Energy Information Administration (EIA)." *U.S. Energy Information Administration*. https://www.eia.gov/electricity/data/eia860/; Department of Energy.

Eisenson, Matthew, Jacob Elkin, Ivonne Norman, Rebecca Coombs, Chadol Kim, Rex Koenig, Suzan Michalski, et al. 2025. "Opposition to Renewable Energy Facilities in the United States: June 2025 Edition." *Sabin Center for Climate Change Law*, June.

Elhorst, J. Paul. 2014. *Spatial Econometrics: From Cross-Sectional Data to Spatial Panels*. SpringerBriefs in Regional Science. Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-40340-8.

EPA. 2015. "Technical Support Document for the Cross-State Air Pollution Rule for the 2008 Ozone NAAQS." Technical {{Support Document}}. U.S. EPA, Office of Air and Radiation.

———. 2025. "Clean Air Markets Program Data (CAMPD) | US EPA." Government. *Clean Air Markets Program Data (CAMPD)*. https://campd.epa.gov; Environmental Protection Agency (EPA) Clean Air Markets Program (CAMP).

———. 2026. "Form EIA-923 Detailed Data with Previous Form Data (EIA-906/920) - U.S. Energy Information Administration (EIA)." Government. *Electricity*. https://www.eia.gov/electricity/data/eia923/index.php; Environmental Protection Agency (EPA) Energy Information Administration (EIA).

EPA eGRID. 2009. “Technical Support Document for the 8th Edition of eGRID with Year 2009 Data.” EPA eGRID.

———. 2018. “Technical Support Document for the 13th Edition of eGRID with Year 2018 Data.” EPA eGRID.

———. 2021. “16th Edition of eGRID with Year 2021 Data, Technical Support Document.” EPA eGRID.

Farrell, Justin, Paul Berne Burow, Kathryn McConnell, Jude Bayham, Kyle Whyte, and Gal Koss. 2021. “Effects of Land Dispossession and Forced Migration on Indigenous Peoples in North America.” *Science* 374 (6567): eabe4943. https://doi.org/10.1126/science.abe4943.

Fischer-Kowalski, Marina, and Helmut Haberl, eds. 2007. *Socioecological Transitions and Global Change: Trajectories of Social Metabolism and Land Use*. Advances in Ecological Economics. Cheltenham, UK ; Northampton, MA: Edward Elgar.

Folch, David C., Daniel Arribas-Bel, Julia Koschinsky, and Seth E. Spielman. 2016. “Spatial Variation in the Quality of American Community Survey Estimates.” *Demography* 53 (5): 1535–54. https://doi.org/10.1007/s13524-016-0499-1.

Fotheringham, A Stewart. 2009. “Geographically Weighted Regression.” In *The SAGE Handbook of Spatial Analysis*.

Fowlie, Meredith. 2010. “Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement.” *American Economic Review* 100 (3): 837–69. https://doi.org/10.1257/aer.100.3.837.

Fox, John, Sanford Weisberg, and Brad Price. 2026. *Car: Companion to Applied Regression*.

https://github.com/bprice2652/car_repo.

Franzese, Robert J., and Jude C. Hays. 2007. "Spatial Econometric Models of Cross-Sectional Interdependence in Political Science Panel and Time-Series-Cross-Section Data." *Political Analysis* 15 (2): 140–64.

Freudenburg, William R. 2006. "Environmental Degradation, Disproportionality, and the Double Diversion: Reaching Out, Reaching Ahead, and Reaching Beyond." *Rural Sociology* 71 (1): 3–32. https://doi.org/10.1526/003601106777789792.

Fujita, K. Sydny, Zachary H Ancona, Louisa A Kramer, Mary M Straka, Tandie E. Gautreau, Christopher P Garrity, Dana Robson, James (Jay) E. Diffendorfer, and Ben Hoen. 2026. "United States Large-Scale Solar Photovoltaic Database (Ver. 4.0, April 2026)." U.S. Geological Survey. https://doi.org/10.5066/P9IA3TUS.

Ge, Mengpin, Johannes Friedrich, and Leandro Vigna. 2026. "Where Do Emissions Come From? These Charts Explain Greenhouse Gas Emissions by Sector," April.

Gilbert, Angela, and Jayajit Chakraborty. 2011. "Using Geographically Weighted Regression for Environmental Justice Analysis: Cumulative Cancer Risks from Air Toxics in Florida." *Social Science Research*. https://doi.org/10.1016/j.ssresearch.2010.08.006.

Gonzalez, Destiny, and Alexandro Carlotto. 2026. "State Climate Policy Dashboard." Database. *State Climate Policy Dashboard*. https://www.climatepolicydashboard.org; Climate XChange.

Grant, Don, Kelly Bergstrand, and Katrina Running. 2014. "Effectiveness of US State Policies in Reducing CO2 Emissions from Power Plants." *Nature Climate Change* 4 (11): 977–82. https://doi.org/10.1038/nclimate2385.

Grant, Don, Andrew K. Jorgenson, and Wesley Longhofer. 2016. "How Organizational and Global Factors Condition the Effects of Energy Efficiency on CO2 Emission Rebounds Among the World's Power Plants." *Energy Policy* 94 (July): 89–93. https://doi.org/10.1016/j.enpol.2016.03.053.

Grant, Don, Andrew Jorgenson, and Wesley Longhofer. 2018. "Pathways to Carbon Pollution: The Interactive Effects of Global, Political, and Organizational Factors on Power Plants' CO2 Emissions." *Sociological Science* 5: 58–92. https://doi.org/10.15195/v5.a4.

Grant, Don, Wesley Longhofer, and Andrew Jorgenson. 2024. "Super Polluters." In *Routledge Handbook of Climate Change and Society*, 2nd ed. Routledge.

Grant, Don, Mary Nell Trautner, Liam Downey, and Lisa Thiebaud. 2010. "Bringing the Polluters Back In: Environmental Inequality and the Organization of Chemical Production." *American Sociological Review* 75 (4): 479–504. https://doi.org/10.1177/0003122410374822.

Green, Jessica, Jennifer Hadden, Thomas Hale, and Paasha Mahdavi. 2021. "Transition, Hedge, or Resist? Understanding Political and Economic Behavior Toward Decarbonization in the Oil and Gas Industry." *Review of International Political Economy*, July, 1–28. https://doi.org/10.1080/09692290.2021.1946708.

Greenberg, Pierce. 2017. "Disproportionality and Resource-Based Environmental Inequality: An Analysis of Neighborhood Proximity to Coal Impoundments in Appalachia." Journal {{Article}}. *Rural Sociology* 82 (1): 149–78. https://doi.org/10.1111/ruso.12119.

Grossmann, Matt, Marty P. Jordan, and Joshua Mccrain. 2021. "The Correlates of State Policy and the Structure of State Panel Data." Journal {{Article}}. *State Politics & Policy Quarterly* 21 (4): 430–50. https://doi.org/10.1017/spq.2021.17.

Hall, Peter, and David Soskice. 2001. “An Introduction to Varieties of Capitalism.” Book {{Section}}. In *Varieties of Capitalism: The Institutional Foundations of Comparative Advantage*.

Hamilton, Michael S. 1979. “Power Plant Siting: A Literature Review.” *Natural Resources Journal* 19: 75–95.

Hernandez-Cortes, Danae, Kyle C. Meng, and Paige Weber. 2023. “Decomposing Trends in US Air Pollution Disparities from Electricity.” *Environmental and Energy Policy and the Economy* 4 (January): 91–124. https://doi.org/10.1086/722674.

Hester, Jim, and Jennifer Bryan. 2024. *Glue: Interpreted String Literals*. https://glue.tidyverse.org/.

Hornier, Michael D. 2026. “79 Michigan Municipalities Take Fight for Local Zoning Control to State Supreme Court.” https://www.fosterswift.com/newsroom/news/michigan-supreme-court-zoning-control-renewable-energy-pa233-mpsc-appeal.

Hothorn, Torsten, Achim Zeileis, Richard W. Farebrother, and Clint Cummins. 2022. *Lmtest: Testing Linear Regression Models*. https://CRAN.R-project.org/package=lmtest.

Iannone, Richard, Joe Cheng, Barret Schloerke, Shannon Haughton, Ellis Hughes, Alexandra Lauer, Romain François, JooYoung Seo, Ken Brevoort, and Olivier Roy. 2026. *Gt: Easily Create Presentation-Ready Display Tables*. https://gt.rstudio.com.

InfluenceMap. 2024. “Carbon Majors: 2024 Data Update.” https://carbonmajors.org/briefing/Carbon-Majors-2024-Data-Update-35466.

Johnson, Boh L. 1995. “Resource Dependence Theory: A Political Economy Model of Organizations.”

Johnson, Leland L. 1973. "Behavior of the Firm Under Regulatory Constraint: A Reassessment." *The American Economic Review* 63 (2): 90–97. https://www.jstor.org/stable/1817057.

Karlson, Kristian Bernt, Anders Holm, and Richard Breen. 2012. "Comparing Regression Coefficients Between Same-sample Nested Models Using Logit and Probit." *Sociological Methodology*. https://doi.org/10.1177/0081175012444861.

Kaufman, Alexander C. 2020. "New York City's Hottest New Energy Fight." *Grist*.

Khanchel, Imen, Naima Lassoued, and Cyrine Khiari. 2025. "Are Pollution Control Bonds and Public Ownership Really Blessing for Utility Firms?" *Journal of Public Affairs* 25 (1): e70004. https://doi.org/10.1002/pa.70004.

Kim, Hae-Young. 2013. "Statistical Notes for Clinical Researchers: Assessing Normal Distribution (2) Using Skewness and Kurtosis." *Restorative Dentistry & Endodontics* 38 (1): 52. https://doi.org/10.5395/rde.2013.38.1.52.

Kim, Sung Eun, and Johannes Urpelainen. 2017. "The Polarization of American Environmental Policy: A Regression Discontinuity Analysis of Senate and House Votes, 1971–2013." *Review of Policy Research* 34 (4): 456–84. https://doi.org/10.1111/ropr.12238.

Klein, Naomi. 2014. *This Changes Everything : Capitalism Vs. The Climate*. Book. First Simon & Schuster hardcover edition. New York: New York : Simon & Schuster.

Knuth, Sarah, Ingrid Behrsin, Anthony Levenda, and James McCarthy. 2022. "New Political Ecologies of Renewable Energy." *Environment and Planning E: Nature and Space* 5 (3): 997–1013. https://doi.org/10.1177/25148486221108164.

Komsta, Lukasz, and Frederick Novomestky. 2022. *Moments: Moments, Cumulants, Skewness,*

Kurtosis and Related Tests. https://www.r-project.org.

Konisky, David M. 2009. "Inequities in Enforcement? Environmental Justice and Government Performance." *Journal of Policy Analysis and Management* 28 (1): 102–21. https://doi.org/10.1002/pam.20404.

Konisky, David M., and Christopher Reenock. 2013. "Compliance Bias and Environmental (In)Justice." *The Journal of Politics*. https://doi.org/10.1017/S0022381613000170.

Konisky, David M., and Tyler S. Schario. 2010. "Examining Environmental Justice in Facility-Level Regulatory Enforcement." *Social Science Quarterly* 91 (3): 835–55. https://doi.org/10.1111/j.1540-6237.2010.00722.x.

Krajewski, Allison K., Claire M. Holesovsky, David M. Lehmann, Anne Barkley, James Brown, Peter Byrley, Catheryne Chiang, et al. 2026. "Integrated Science Assessment (ISA) for Oxides of Nitrogen." Integrated {{Science Assessment}} EPA/600/R-25/232. US Environmental Protection Agency (EPA).

Krieger, Elena M., Joan A. Casey, and Seth B. C. Shonkoff. 2016. "A Framework for Siting and Dispatch of Emerging Energy Resources to Realize Environmental and Health Benefits: Case Study on Peaker Power Plant Displacement." *Energy Policy*. https://doi.org/10.1016/j.enpol.2016.05.049.

Kuha, Jouni, and Colin Mills. 2020. "On Group Comparisons With Logistic Regression Models." *Sociological Methods & Research* 49 (2): 498–525. https://doi.org/10.1177/0049124117747306.

Lachapelle, E., R. Macneil, and M. Paterson. 2017. "The Political Economy of Decarbonisation: From Green Energy "Race' to Green "Division of Labour'." Journal {{Article}}. *New Political Economy* 22 (3): 311–27. https://doi.org/10.1080/13563467.2017.1240669.

Lee, Barrett A., Sean F. Reardon, Glenn Firebaugh, Chad R. Farrell, Stephen A. Matthews, and David O'Sullivan. 2008. "Beyond the Census Tract: Patterns and Determinants of Racial Segregation at Multiple Geographic Scales." *American Sociological Review* 73 (5): 766–91. https://doi.org/10.1177/000312240807300504.

LeSage, James. 2009. "Motivating and Interpreting Spatial Econometric Models." In *Introduction to Spatial Econometrics*.

LeSage, James, and R. Pace. 2014. "The Biggest Myth in Spatial Econometrics." *Econometrics* 2 (4): 217–49. https://doi.org/10.3390/econometrics2040217.

Lewis-Beck, Michael, Tim Futing Liao, and Alan Bryman. 2026. "Middle-Range Theory." *The Sage Encyclopedia of Social Science Research Methods*, August. https://doi.org/10.4135/9781412950589.n560.

Liévanos, R. S., P. Greenberg, and R. Wishart. 2018. "In the Shadow of Production: Coal Waste Accumulation and Environmental Inequality Formation in Eastern Kentucky." *Social Science Research* 71 (March): 37–55. https://doi.org/10.1016/j.ssresearch.2018.01.003.

Liévanos, Raoul S. 2015. "Race, Deprivation, and Immigrant Isolation: The Spatial Demography of Air-Toxic Clusters in the Continental United States." *Social Science Research* 54 (November): 50–67. https://doi.org/10.1016/j.ssresearch.2015.06.014.

———. 2018. "Retooling CalEnviroScreen: Cumulative Pollution Burden and Race-Based Environmental Health Vulnerabilities in California." *International Journal of Environmental Research and Public Health* 15 (4): 762. https://doi.org/10.3390/ijerph15040762.

———. 2019. "Air-Toxic Clusters Revisited: Intersectional Environmental Inequalities and Indigenous Deprivation in the U.S. Environmental Protection Agency Regions." *Race and*

Social Problems 11 (2): 161–84. https://doi.org/10.1007/s12552-019-09260-5.

Malm, Andreas. 2016. *Fossil Capital: The Rise of Steam-Power and the Roots of Global Warming*. London ; New York: Verso.

McAdam, Doug, and Hilary Boudet. 2012. *Putting Social Movements in Their Place : Explaining Opposition to Energy Projects in the United States, 2000-2005*. Book. Cambridge Studies in Contentious Politics. Cambridge: Cambridge University Press.

McCoy, Ember D. 2017. "Which Came First, Coal-Fired Power Plants or Communities of Color?" PhD thesis, University of Michigan.

Meadows, Jim. 2024. "Some Midwest States Look to Counter Local Opposition to Wind and Solar Farm Projects : NPR." *All Things Considered*. Urbana, IL: National Public Radio.

Meckling, J. 2015. "Oppose, Support, or Hedge? Distributional Effects, Regulatory Pressure, and Business Strategy in Environmental Politics." *Global Environmental Politics* 15 (2): 19–+. https://doi.org/10.1162/GLEP_a_00296.

Melitz, Marc J., and Sašo Polanec. 2015. "Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit." *The RAND Journal of Economics* 46 (2): 362–75. https://doi.org/10.1111/1756-2171.12088.

Mildenberger, Matto. 2020. *Carbon Captured: How Business and Labor Control Climate Politics*. American and Comparative Environmental Policy. Cambridge, Massachusetts ; London, England: The MIT Press.

Miller, Gary J. 2005. "The Political Evolution of Principal-Agent Models." *Annual Review of Political Science* 8 (1): 203–25. https://doi.org/10.1146/annurev.polisci.8.082103.104840.

Milman, Oliver. 2022. "'It's Got Nasty': The Battle to Build the US's Biggest Solar Power Farm." *The Guardian*, October.

Milman, Oliver, and Andrew Witherspoon. 2024. "The US Says It Needs up to 22m Acres for the Solar Energy Transition – Here's What That Looks Like." *The Guardian*, January.

Mohai, Paul, and Bunyan Bryant. 1992. "Race, Poverty & the Distribution of Environmental Hazards: Reviewing the Evidence." *Race, Poverty & the Environment* 2: 3–27. https://www.jstor.org/stable/41555195.

Mohai, Paul, David Pellow, and J. Timmons Roberts. 2009. "Environmental Justice." *Annual Review of Environment and Resources* 34 (1): 405–30. https://doi.org/10.1146/annurev-environ-082508-094348.

Mohai, Paul, and Robin Saha. 2006. "Reassessing Racial and Socioeconomic Disparities in Environmental Justice Research." *Demography*. https://doi.org/10.1353/dem.2006.0017.

———. 2007. "Racial Inequality in the Distribution of Hazardous Waste: A National-Level Reassessment." *Social Problems*. https://doi.org/10.1525/sp.2007.54.3.343.

———. 2015. "Which Came First, People or Pollution? Assessing the Disparate Siting and Post-Siting Demographic Change Hypotheses of Environmental Injustice." *Environmental Research Letters* 10 (11): 115008. https://doi.org/10.1088/1748-9326/10/11/115008.

"Monthly Energy Review." 2026. U.S. Energy Information Administration.

Mood, C. 2010. "Logistic Regression: Why We Cannot Do What We Think We Can Do, and What We Can Do About It." *European Sociological Review* 26 (1): 67–82. https://doi.org/10.1093/esr/jcp006.

Mueller, J. Tom, and Matthew M. Brooks. 2020. "Burdened by Renewable Energy? A Multi-Scalar Analysis of Distributional Justice and Wind Energy in the United States." *Energy Research & Social Science* 63 (May): 101406. https://doi.org/10.1016/j.erss.2019.101406.

Mühleisen, Hannes, and Mark Raasveldt. 2026. *Duckdb: DBI Package for the DuckDB Database Management System*. https://r.duckdb.org/.

Müller, Kirill. 2025. *Here: A Simpler Way to Find Your Files*. https://here.r-lib.org/.

Müller, Kirill, and Hadley Wickham. 2026. *Tibble: Simple Data Frames*. https://tibble.tidyverse.org/.

NCSL. 2026. "State Partisan Composition." *National Conference of State Legislators*. https://www.ncsl.org/about-state-legislatures/state-partisan-composition.

Newell, Peter, Philipp Pattberg, and Heike Schroeder. 2012. "Multiactor Governance and the Environment." *Annual Review of Environment and Resources* 37 (1): 365–87. https://doi.org/10.1146/annurev-environ-020911-094659.

Nunez, Yanelli, Jaime Benavides, Jenni A. Shearston, Elena M. Krieger, Misbath Daouda, Lucas R. F. Henneman, Erin E. McDuffie, Jeff Goldsmith, Joan A. Casey, and Marianthi-Anna Kioumourtzoglou. 2024. "An Environmental Justice Analysis of Air Pollution Emissions in the United States from 1970 to 2010." *Nature Communications* 15 (1): 268. https://doi.org/10.1038/s41467-023-43492-9.

O'Shaughnessy, Eric, Ryan Wiser, Ben Hoen, Joseph Rand, and Salma Elmallah. 2023. "Drivers and Energy Justice Implications of Renewable Energy Project Siting in the United States." *Journal of Environmental Policy & Planning* 25 (3): 258–72. https://doi.org/10.1080/1523908X.2022.2099365.

Oates, Wallace E., and Paul R. Portney. 2003. "The Political Economy of Environmental Policy." In *Handbook of Environmental Economics*, 1:325–54. Elsevier. https://doi.org/10.1016/S1574-0099(03)01013-1.

Parker, Dominic P., Sarah Johnston, Bryan Leonard, Daniel Stewart, and Justin B. Winikoff. 2024. "Economic Potential of Wind and Solar in American Indian Communities." *Nature Energy* 9 (11): 1360–68. https://doi.org/10.1038/s41560-024-01617-4.

Pastor, Manuel, Jim Sadd, and John Hipp. 2001. "Which Came First? Toxic Facilities, Minority Move-In, and Environmental Justice." *Journal of Urban Affairs*. https://doi.org/10.1111/0735-2166.00072.

Paterson, Matthew, and Xavier P-Laberge. 2018. "Political Economies of Climate Change." *WIREs Climate Change* 9 (2): e506. https://doi.org/10.1002/wcc.506.

Pebesma, Edzer. 2026. *Sf: Simple Features for r*. https://r-spatial.github.io/sf/.

Pedersen, Thomas Lin. 2025. *Patchwork: The Composer of Plots*. https://patchwork.data-imaginist.com.

Pellow, David N. 2000. "Environmental Inequality Formation: Toward a Theory of Environmental Injustice." *American Behavioral Scientist* 43 (4): 581–601. https://doi.org/10.1177/0002764200043004004.

———. 2016. "Toward a Critical Environmental Justice Studies: Black Lives Matter as an Environmental Justice Challenge." *Du Bois Review: Social Science Research on Race* 13 (2): 221–36. https://doi.org/10.1017/S1742058X1600014X.

Perrow, Charles. 1997. "Organizing for Environmental Destruction." *Organization & Environment*.

https://doi.org/10.1177/0921810697101009.

Polanyi, Karl. 2010. *The Great Transformation: The Political and Economic Origins of Our Time*. 2. Beacon paperback ed., [reprinted]. Boston, Mass: Beacon Press.

Prechel, Harland. 2023. "Organizational Political Economy, Corporate Power, and the Great Acceleration of Environmental Pollution in the United States." In *Handbook of Inequality and the Environment*, 288–307. Cheltenham, UK: Edward Elgar Publishing. https://doi.org/10.4337/9781800881136.00029.

Prechel, Harland, Amber Blazek, and Ernesto F. L. Amaral. 2025. "Toward Theory Consolidation: Stratification, Organizational, and Political-Legal Effects on Greenhouse Gas Emissions." *Energy Research & Social Science* 128 (October): 104330. https://doi.org/10.1016/j.erss.2025.104330.

Prechel, Harland, and Alesha Istvan. 2016. "Editors' Pick: Disproportionality of Corporations' Environmental Pollution in the Electrical Energy Industry." *Sociological Perspectives* 59 (3): 505–27. https://doi.org/10.1177/0731121416629991.

Prechel, Harland, and George Touche. 2014. "The Effects of Organizational Characteristics and State Environmental Policies on Sulfur-Dioxide Pollution in U.S. Electrical Energy Corporations." *Social Science Quarterly* 95 (1): 76–96. https://doi.org/10.1111/ssqu.12052.

Pronobis, Marek. 2020. "Chapter 4 - Reduction of Nitrogen Oxide Emissions." In *Environmentally Oriented Modernization of Power Boilers*, 79–133. Elsevier. https://doi.org/10.1016/B978-0-12-819921-3.00004-2.

Przeworski, Adam. 1991. *Capitalism and Social Democracy*. Studies in Marxism and Social Theory. Cambridge London New York Paris: Cambridge university press Éd. de la Maison des sciences

de l'homme.

Pulido, Laura, Steve Sidawi, and Robert O. Vos. 1996. "An Archaeology of Environmental Racism in Los Angeles." *Urban Geography* 15: 419–39.

R Special Interest Group on Databases (R-SIG-DB), Hadley Wickham, and Kirill Müller. 2024. *DBI: R Database Interface*. https://dbi.r-dbi.org.

Raimi, Daniel, Sanya Carley, and David Konisky. 2022. "Mapping County-Level Vulnerability to the Energy Transition in US Fossil Fuel Communities." *Scientific Reports*. https://doi.org/10.1038/s41598-022-19927-6.

Ramanan, Abbe. 2026. "Phase Out Peakers." *Clean Energy Group*. https://www.cleane-group.org/initiatives/phase-out-peakers/.

Ramanan, Abbe, Eva Morgan, Matthew Shorraw, and Eleanor Breslin. n.d. "How Data Centers Are Reshaping Pennsylvania's Energy Landscape."

Rand, Joseph T., Louisa A. Kramer, Christopher P. Garrity, Ben D. Hoen, Jay E. Diffendorfer, Hannah E. Hunt, and Michael Spears. 2020. "A Continuously Updated, Geospatially Rectified Database of Utility-Scale Wind Turbines in the United States." *Scientific Data* 7 (1): 15. https://doi.org/10.1038/s41597-020-0353-6.

Richards, John R. 2000. "EPA Air Pollution Training Institute (APTI) Course 418, Control of Nitrogen Oxides Emissions." EPA Contract No.: 68D99022. U.S. EPA, Air Pollution Training Institute.

Richardson, Neal, Ian Cook, Nic Crane, Dewey Dunnington, Romain François, Jonathan Keane, Bryce Mecum, et al. 2026. *Arrow: Integration to Apache 'Arrow'*. https://github.com/apache/

arrow/.

Richmond-Bryant, Jennifer, Ihab Mikati, Adam F. Benson, Thomas J. Luben, and Jason D. Sacks. 2020. "Disparities in Distribution of Particulate Matter Emissions from US Coal-Fired Power Plants by Race and Poverty Status After Accounting for Reductions in Operations Between 2015 and 2017." *American Journal of Public Health* 110 (5): 655–61. https://doi.org/10.2105/AJPH.2019.305558.

Ringquist, Evan J. 2005. "Assessing Evidence of Environmental Inequities: A Meta-Analysis." *Journal of Policy Analysis and Management*. https://doi.org/10.1002/pam.20088.

Saha, Robin, and Paul Mohai. 2005. "Historical Context and Hazardous Waste Facility Siting: Understanding Temporal Patterns in Michigan." *Social Problems* 52 (4): 618–48. https://doi.org/10.1525/sp.2005.52.4.618.

Sareen, Siddharth, Shayan Shokrgozar, James McCarthy, and Steven Wolf. 2025. "The Dawn of Solar Photovoltaics: Emergent Political Economies at the Solar–Agri–Land Nexus." *Sustainability Science* 20 (4): 1271–76. https://doi.org/10.1007/s11625-025-01721-8.

Schroeder, Jonathan P. 2007. "Target-Density Weighting Interpolation and Uncertainty Evaluation for Temporal Analysis of Census Data." *Geographical Analysis* 39 (3): 311–35. https://doi.org/10.1111/j.1538-4632.2007.00706.x.

Seto, Karen C., Steven J. Davis, Ronald B. Mitchell, Eleanor C. Stokes, Gregory Unruh, and Diana Ürge-Vorsatz. 2016. "Carbon Lock-In: Types, Causes, and Policy Implications." Journal {{Article}}. *Annual Review of Environment and Resources* 41 (1): 425–52. https://doi.org/10.1146/annurev-environ-110615-085934.

Shehabi, Arman, Alex Newkirk, Sarah J. Smith, Alex Hubbard, Nuoa Lei, Md Abu Bakar Siddik,

Billie Holecek, Jonathan Koomey, Eric Masanet, and Dale Sartor. 2024. “2024 United States Data Center Energy Usage Report,” December. https://doi.org/10.71468/P1WC7Q.

Slowikowski, Kamil. 2026. *Ggrepel: Automatically Position Non-Overlapping Text Labels with Ggplot2*. https://ggrepel.slowkow.com/.

Smil, Vaclav. 2010. “Power Density Primer: Understanding the Spatial Dimension of the Unfolding Transition to Renewable Electricity Generation (Parts 1-5).” Report.

Sorrels, John L, David D Randall, Karen S Schaffner, and Carrie Richardson Fry. 2019. “Chapter 2: Selective Catalytic Reduction.” In *EPA Air Pollution Control Cost Manual, Section 4: NOx Controls.*, 1–93. U.S. EPA, Office of Air Quality Planning and Standards.

Stokes, Danielle. 2021. “Renewable Energy Federalism.” *Minnesota Law Review* 106 (March).

Stokes, Leah C., Emma Franzblau, Jessica R. Lovering, and Chris Miljanich. 2023. “Prevalence and Predictors of Wind Energy Opposition in North America.” *Proceedings of the National Academy of Sciences* 120 (40). https://doi.org/10.1073/pnas.2302313120.

Tessum, Christopher W., David A. Paolella, Sarah E. Chambliss, Joshua S. Apte, Jason D. Hill, and Julian D. Marshall. 2021. “$PM_{2.5}$ Polluters Disproportionately and Systemically Affect People of Color in the United States.” *Science Advances* 7 (18): eabf4491. https://doi.org/10.1126/sciadv.abf4491.

Tobler, W. R. 1970. “A Computer Movie Simulating Urban Growth in the Detroit Region.” *Economic Geography* 46 (June): 234. https://doi.org/10.2307/143141.

US EPA, OAR. 2016. “Nonattainment Areas for Criteria Pollutants (Green Book).” Collections and {{Lists}}. https://www.epa.gov/green-book.

Wainwright, Oliver. 2023. “How Solar Farms Took over the California Desert: ‘An Oasis Has Become a Dead Sea’.” *The Guardian*, May.

Walker, Kyle. 2025. *Tigris: Load Census TIGER/Line Shapefiles*. https://github.com/walkerke/tigris.

Wallerstein, Immanuel. 2004. *World-Systems Analysis: An Introduction*. Duke University Press. https://doi.org/10.2307/j.ctv11smzx1.

Walsh, Edward. 1981. “Resource Mobilization and Citizen Protest in Communities Around Three Mile Island.” Journal {{Article}}. *Social Problems* 29 (October 1981).

Wamsted, Daniel. 2023. “Private Equity in PJM: New Risks for Limited Partners, Private Capital,” Institute for Energy Economics and Financial Analysis.

Wickham, Hadley, Winston Chang, Lionel Henry, Thomas Lin Pedersen, Kohske Takahashi, Claus Wilke, Kara Woo, Hiroaki Yutani, Dewey Dunnington, and Teun van den Brand. 2026. *Ggplot2: Create Elegant Data Visualisations Using the Grammar of Graphics*. https://ggplot2.tidyverse.org.

Wickham, Hadley, Romain François, Lionel Henry, Kirill Müller, and Davis Vaughan. 2026. *Dplyr: A Grammar of Data Manipulation*. https://dplyr.tidyverse.org.

Wickham, Hadley, and Lionel Henry. 2026. *Purrr: Functional Programming Tools*. https://purrr.tidyverse.org/.

Wickham, Hadley, Davis Vaughan, and Maximilian Girlich. 2025. *Tidyr: Tidy Messy Data*. https://tidyr.tidyverse.org.

Wilke, Claus O. 2025. *Ggridges: Ridgeline Plots in Ggplot2*. https://wilkelab.org/ggridges/.

Wüstenhagen, Rolf, Maarten Wolsink, and Mary Jean Bürer. 2007. "Social Acceptance of Renewable Energy Innovation: An Introduction to the Concept." *Energy Policy*. https://doi.org/10.1016/j.enpol.2006.12.001.

Xie, Yihui. 2025. *Knitr: A General-Purpose Package for Dynamic Report Generation in r*. https://yihui.org/knitr/.

Yergin, Daniel. 2012. *The Quest: Energy, Security, and the Remaking of the Modern World*. Penguin.

———. 2021. *The New Map: Energy, Climate and the Clash of Nations*. New York: Penguin Books, an imprint of Penguin Random House LLC.

York, Richard. 2006. "Ecological Paradoxes: William Stanley Jevons and the Paperless Office." *Research in Human Ecology* 13.

York, Richard, and Shannon Elizabeth Bell. 2019. "Energy Transitions or Additions?" Journal {{Article}}. *Energy Research & Social Science* 51: 40–43. https://doi.org/10.1016/j.erss.2019.01.008.

Zeileis, Achim, and Thomas Lumley. 2024. *Sandwich: Robust Covariance Matrix Estimators*. https://sandwich.R-Forge.R-project.org/.

Zhu, Hao. 2024. *kableExtra: Construct Complex Table with Kable and Pipe Syntax*. http://haozhu233.github.io/kableExtra/.