THE UNIVERSITY OF CHICAGO

SHORT CIRCUITED: COSTLY TRANSITIONS UNDER THE CLEAN AIR ACT

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
ANDREW NELSON MCKINLEY

CHICAGO, ILLINOIS
AUGUST 2025

This dissertation would not have been possible without the countless people who have shaped me into who I am today. While this work is dedicated to you all, I would like to highlight a handful of individuals.

To Colleen Hongisberg, thank you for all your support and guidance over the years, this journey would not have been possible without your near decade-spanning efforts. To Robert Jackson Jr., for taking a chance on me and continuing to support my journey throughout the years. To Patrick Corrigan and Ed Rock, your mentorship and votes of confidence provided the push I needed to set out on this journey. To Hans Christensen, Christian Leuz, Philip Berger, Bobby Bartlett, Joshua Macey, and Hajin Kim, thank you for being selfless with your time and support.

To Adam Garrison, Patrick Gilsenan, and Christos Hountalas, thank you for your patience and companionship across decades and cities. To Lauren Mostrom, for your love and support.

To Meemaw, may I never lose your spirit.

ABSTRACT

Tightening future environmental regulations can perversely accelerate fossil-fuel extraction and emissions by lowering prices today, a phenomenon known as the Green Paradox. I introduce a novel mechanism amplifying the Green Paradox, the "depreciation channel," whereby firms use soon-to-be-stranded capital more intensively ahead of stricter rules. Exploiting exogenous variation from a shock to legal procedure in implementing U.S. power plant regulations, I demonstrate this channel and quantify its impact. I estimate a 6.5% rise in annual CO_2 emissions and air pollution, with an associated social cost of over $22 billion. To mitigate these damages, I argue that the EPA should incorporate these costs into its final target and invoke its emergency authority under the Clean Air Act to curb short-term increases in air pollution.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I thank the members of my Ph.D. committee, Christian Leuz, Hans Christensen, Philip Berger, and Colleen Honigsberg, for their guidance and support over the years. I am also grateful to Ryan Bubb, Eugene Fama, Jacob Goldin, Daniel Hemel, Ryan Kellogg, Hajin Kim, Joshua Macey, Lauren Mostrom, Maximilian Muhn, David Weisbach, Anthony Welsch, and participants at ESIFCLIM, the American Law and Economics Association 2024 Annual Meeting, Chicago Booth and NYU School of Law for their helpful comments and suggestions. All errors are my own.

CHAPTER 1

SHORT CIRCUITED: COSTLY TRANSITIONS UNDER THE CLEAN AIR ACT

1.1 Introduction

Environmental regulations, whether cap-and-trade or command-and-control, are viewed as essential to addressing two pressing policy issues: climate change and air pollution. In the U.S., policies targeting greenhouse gas (GHG) emissions have been inconsistent, progressing in fits and starts. Some have argued that the delay and unpredictability create perverse incentives for firms to increase short-run emissions, potentially exacerbating climate damages rather than curtailing them (e.g., Sinn 2008; **gowrisankaran2022policy**).

This effect, known as the Green Paradox, extends the logic of Hotelling 1931 by introducing an anticipatory response: as firms expect future GHG regulation to lower the future value of extraction, they accelerate fossil fuel production today. This intertemporal shift increases supply, depresses prices (Gerlagh 2011; Hoel 2012; Lemoine 2017), raises consumption, and front-loads emissions, diminishing the effectiveness of regulation.

Yet the empirical relevance of the Green Paradox to U.S. power plant emissions regulations is unclear. One reason is institutional: vertically integrated utilities often recover fuel costs through automatic pass-through mechanisms such as fuel adjustment clauses (Macey 2020). Cost-of-service regulation further shields these firms from market fluctuations, insulating them from the price incentives that drive the Green Paradox. On the other hand, the price channel may be particularly powerful given that fuel accounts for most of the power sector's marginal costs (Fabrizio, Rose, and Wolfram 2007). Further, the impact of price can vary with market structure, where contractual resale restrictions make for sticky prices and market power may limit pass-through (Preonas 2024).

This paper argues that the Green Paradox operates through more than just prices. I isolate an additional channel rooted in firms' incentives to accelerate the consumption of long-lived, soon-to-be-stranded capital. Power plants, particularly coal-fired units, are durable assets–the average U.S. coal plant is over fifty years old–yet changes in output consume some of the plant's useful life, making economic depreciation a marginal cost. When policy changes make future operations more costly, firms have an incentive to increase output today to extract residual value. Or as the CEO of DTE Energy remarked in response to the Clean Power Plan: "Now we have an aging coal fleet....I always like to characterize them as old, being from the Motor City, old cars, old Thunderbirds. And you can replace a fender, replace a transmission. At some point, you're going to have to replace the whole car."[1] In essence, regulation barring dirty generation tomorrow reduces the asset's value to zero, prompting firms to consume the asset's productive capacity while they still can. This dynamic, the "depreciation channel", reinforces the price channel, increasing short-run emissions beyond what prices alone would predict.

I begin by providing reduced-form evidence on the depreciation channel and its contribution to the Green Paradox using plausibly exogenous variation from an unexpected change in American judicial procedure, the issuance of a stay by the Supreme Court during litigation over the Clean Power Plan (CPP). Enacted in 2015, the CPP aimed to cut power plant sector emissions to 32% below 2005 levels by 2030. Specifically, the regulation sought to minimize coal and shift generation towards renewables and natural gas. In a historic first, however, the Supreme Court intervened before any lower court ruling, halting implementation and abruptly altering expectations that the rule would take effect. This unexpected intervention enables me to overcome two identification challenges that have plagued empirical investigations of the Green Paradox.

First, the passage of legislation is unsuitable for quasi-natural experiments because the public nature of the legislative process allows firms to adjust their behavior in anticipation of future regulation. In contrast, the stay was issued without public consultation and before

---

[1] DTE Energy Company, Barclays CEO Energy-Power Conference 2015.

oral arguments, when the public typically receives information about potential rulings. These factors create a clean event date when firms update their expectations about the viability of the CPP. As the CEO of DTE Energy noted during an earnings call the day after the stay, “I must say that was a surprise to me.”[2] Moreover, the timing of legislation, even if sharply observed, can coincide with shifts in economic fundamentals, prompting lawmakers to act. Here, the timing of the Supreme Court’s intervention was driven by changes in its philosophical makeup, not market conditions or electoral pressure.

I analyze the depreciation channel in two steps. First, I examine changes in out-of-merit generation. The merit order ranks power plants by marginal cost, from lowest to highest; deviations from this order–generation by higher-cost plants when lower-cost options are available—are defined as out-of-merit. I reconstruct the merit order by balancing authority using plant-level data on monthly fuel costs, allowing plants to move up and down the merit order as the Green Paradox’s price channel alters marginal costs. Crucially, because this analysis directly controls for fuel prices, any remaining changes in out-of-merit generation must reflect shifts in non-fuel price marginal costs.

Using a within–balancing authority design that controls for preexisting incentives for inefficient generation and fuel price changes, I find a 7.5% decline in out-of-merit generation following the stay. This suggests that price effects alone cannot explain the observed behavioral response.

Next, I show that the change in out-of-merit generation is accompanied by increased economic depreciation of coal power plants. Changes in power plant output are called “ramping”, and each time a plant ramps, it reduces the life expectancy of the plant and leads to increased maintenance costs (Kumar et al. 2012). Each generator can be thought of as consisting of a stock of depreciation, a unit of which must be consumed when output changes, leading to changes in the firm’s marginal cost of generation and, thus, its future output decisions. Under standard Hotelling 1931 logic, as the future marginal value of remaining depreciation

---

[2] DTE Energy Company, Q4 2015 Earnings Call, Feb 10, 2016.

falls, the firm should increase its consumption of depreciation today. Using hourly data on power plant operations, I find that ramping falls by 1.1% or 13.6% relative to the pre-stay mean. In sum, prices are unlikely to explain the full change in generator behavior, and generators respond by increasing consumption of their capital assets.

The narrow focus on prices has another drawback: we lack direct measures of Paradox's impact on the outcomes of interest—GHG emissions and pollution. Treating the paradox as a simple subsidy may misstate its impact on pollution due to market structure, the omission of the depreciation channel, and the differing emissions profiles of the competing uses for fossil fuels. This paper addresses that gap by providing the first direct estimates of the Green Paradox's impact on particulate matter (PM2.5) and CO_2 emissions, using satellite and administrative data and identification from the CPP stay.

Most of the Clean Power Plan's economic benefits stemmed from reduced particulate matter (PM2.5) pollution, a byproduct of coal-fired generation with severe health effects, including cancer, stroke, heart attacks, and Alzheimer's disease (Ebenstein et al. 2017). Henneman et al. 2023 finds that PM2.5 emissions from coal power plants led to approximately 460,000 deaths between 1999 and 2020. Globally, its impact on life expectancy is comparable to tobacco use and significantly worse than alcohol use, unsafe water, and malnutrition (Air Quality Life Index 2023). To measure the impact of the Green Paradox on PM2.5 emissions, I use daily satellite measurements of PM2.5 concentrations for each zip code in the U.S. I then compare zip codes with a coal-fired power plant to those without, while controlling for the spatial dispersion of pollutants. Following the stay, PM2.5 concentrations fell by approximately 6.5% in zip codes containing a coal power plant.

Next, I examine GHG emissions using a difference-in-differences approach. This analysis compares coal-fired power plants, which were strictly regulated under the Clean Power Plan, with natural gas plants that were largely exempt from these pollution reductions. The results indicate that coal-fired plants reduced CO_2 emissions by 7.7% in response to the stay of the CPP, after controlling for weather, state-level macroeconomic factors, and unobserved differences among generating units.[3] Robustness checks confirm that the results are not driven by specific states or changes in state policy, and that the estimated coefficients are far outside the range of placebo tests.

Using the EPA's current social cost of carbon ($190 per ton) and my PM2.5 estimates, I calculate annual social costs of roughly $14 billion. Including co-pollutants (NH_3, NO_x, SO_2, and VOCs) raises this estimate to $22.5 billion per year.

These findings carry implications for policymakers. As a first step, the EPA could include the effects of the Green Paradox in their cost-benefit analysis, as they recently did for the Phase 3 Heavy Duty Trucks rule (U.S. Environmental Protection Agency 2024a). While the Green Paradox was not considered in the analyses for the Clean Power Plan or recent power plant performance standards, the Trucks rule demonstrates that the EPA could consider these costs. Furthermore, the EPA could use such estimates to jointly determine the phase-in period and reduce final targets to offset emissions accelerated by the Green Paradox (Revesz and Kong 2011). This approach could help maintain a stable total stock of CO_2 emissions, even if other paradox-mitigation measures are not pursued.

Finally, Congress vested the EPA with expansive powers under Section 303 of the Clean Air Act to take action to reduce air pollution, even absent existing standards (Glicksman and Adashek 2024). Under Section 303, the EPA could seek to enjoin all coal power plants from engaging in "unreasonable" changes in utilization in a single proceeding. The use of an ex post reasonableness standard is likely to preserve incentives for power plants to stabilize the grid in response to unexpected shocks, but should enable the EPA to prevent the most egregious behavior. When combined with tighter standards, the EPA should be able to greatly reduce the Paradox's importance.

Beyond analyzing the effects of the Clean Power Plan, this paper contributes to the

---

[3] The treated group includes oil and biomass plants; however, since coal plants constitute the majority of non-natural gas units, "treated" and "coal power plants" are used interchangeably.

extensive literature on the Green Paradox (e.g., Lemoine 2017; Norman and Schlenker 2024; Grafton et al. 2014; Okullo, Reynès, and Hofkes 2021; Di Maria, Lange, and Van der Werf 2014) and anticipatory market effects (Hall 1971; Judd 1985).[4] This paper extends the literature in several ways.

First, this paper introduces a novel aspect of the Green Paradox by extending it from resource extraction to capital investments. Specifically, coal-fired power plants, facing the threat of early retirement due to regulation, accelerate the economic depreciation of their assets to maximize their utilization before regulatory constraints take hold. This effect amplifies the price channel previously documented in the literature. Relatedly, recognizing that the Green Paradox's impact is often underestimated when measured solely through price channels, this paper provides new evidence on its social costs. By leveraging a well-defined exogenous shock with a specific event date, the study offers causal, reduced-form evidence of the Green Paradox, finding that it carries social costs in excess of \$22 billion per year. This approach enhances the understanding of the paradox compared to earlier studies that relied on time-series price data or structural models.

Second, this paper also contributes to the legal literature on optimal policy for legal transitions (Kaplow 1986; Levmore 1999; Nash and Revesz 2007; Shavell 2008). Shavell 2008 is the closest to this paper as it explicitly analyzes the impact of exempting existing durable assets from new regulation. That paper argues that with durable assets, exemptions from stricter standards can minimize social costs. This paper introduces an additional layer to the cost-benefit analysis in the optimal transitions literature by showing that polluting assets scheduled for future shutdowns may adjust their behavior to impose additional social costs in the short term. An analysis of the optimal transition rule should, therefore, account for these costs.

The remainder of this paper proceeds as follows. Section II lays out the theoretical

---

[4] Others focus on the long-run effects of energy substitutes (Grafton et al. 2014 and Okullo, Reynès, and Hofkes 2021). Parton and Dundas 2020 study housing permits and building regulations. This is not an exhaustive listing of Green Paradox papers, but a partial accounting.

intuition guiding the empirical approach. Section III explains the institutional details of my shock: the Clean Power Plan's stay. Next, Section IV introduces the data used in this paper. Section V lays out the research designs and discusses the results. Section VI calculates the estimated damages and explores several policy proposals to mitigate the paradox. Section VII concludes.

1.2 Theoretical Underpinning

The hypothesis underlying my analysis hinges on the concept of intertemporal substitution, wherein proposed emissions regulations alter the planning framework of firms by impacting the regulatory treatment of existing pollution sources. The legal implications of such regulations, particularly through the lens of "grandfathering," have been extensively examined in the legal literature (e.g., Kaplow 1986; Levmore 1999; Stavins 2006; Nash and Revesz 2007; Shavell 2008). Grandfathering exempts existing power plants from compliance with new regulatory standards,[5] a policy that Kaplow 1986 critiques as typically unwarranted. According to Kaplow, grandfathering functions as a form of insurance for operators of existing facilities, inadvertently introducing quasi-moral hazard. This regulatory insurance diminishes incentives for firms to proactively respond to anticipated regulatory changes, potentially leading to overinvestment in assets subject to future changes in regulation.

For example, consider a coal power plant owner deciding between installing scrubber A or scrubber B to reduce SO_2 emissions. Scrubber A meets the current regulatory standard and is slightly less expensive than scrubber B. Suppose it is known that the EPA will soon require all power plants to install scrubber B. If the new rule mandates the installation of scrubber B for all plants, the owner would likely choose scrubber B now to avoid the costs of first installing scrubber A and later replacing it with scrubber B once the rule takes effect. However, if the plant were exempted from this rule (i.e., "grandfathered"), the firm would

---

[5] However, substantial modifications to existing plants may subject them to the new standards.

opt for scrubber A, the cheaper option. This grandfathering provision thus incentivizes the firm to install a less effective scrubber, ultimately leading to higher SO_2 emissions in the meantime.

Shavell 2008 refines this analysis by examining the role of grandfathering for durable assets—those with long useful lives. He argues that grandfathering may be optimal when the compliance costs associated with new regulations exceed the expected benefits. For existing plants, compliance often yields only marginal improvements (e.g., upgrading from scrubber A to scrubber B), while the full cost of compliance must be incurred. In contrast, new plants benefit from both the full regulatory improvements and the associated costs. Given this cost-benefit imbalance, there are scenarios where grandfathering may indeed be appropriate.

In the utility sector, some degree of grandfathering is almost invariably required. This is because energy supply and demand must constantly be in balance, and there can be large swings in demand. Ensuring reliability requires sufficient generating capacity, and adding capacity, which requires regulatory approval and construction, can take many years. The recent New Source Performance Standards for GHG emissions on power plants provide one example (U.S. Environmental Protection Agency 2024b). That regulation requires legally binding commitments from coal power plant owners specifying the date they will retire their plants, starting seven years from the final rule. That deadline was extended from the proposed rule to ensure the electrical grid's reliability.

Given that some degree of grandfathering is necessary, what constitutes the optimal grandfathering policy? The Green Paradox cautions against a straightforward comparison of compliance costs with the benefits of achieving final pollution targets, as it is potentially misleading. Introducing a new rule may alter market behavior well before the grandfathering period concludes, potentially resulting in increased interim emissions. For electricity generators—the focus of this paper—this effect operates through two channels.

The first way the Green Paradox can increase emissions is what I will call the "price channel." This channel is the focus of much of the Green Paradox literature (e.g., Lemoine 2017). To understand this, imagine a new regulation that proposes to ban coal-fired electricity generation in ten years instead of the twenty years allowed under the current rule. Let us also assume that it is always profitable for companies to mine coal during this period. Given these circumstances, coal companies will want to sell all their coal within the ten-year window, rather than spreading sales over twenty years. If they do not, they risk missing out on profits, meaning they have not profit maximized.

When the deadline for coal use is moved up, lower coal prices reduce the marginal cost of coal-fired power plants, increasing the amount of electricity generated by coal. Additionally, standard economic reasoning suggests that prices and extraction rates adjust immediately. If the prices and extraction rates do not immediately adjust, then traders could make a profit by selling coal today and buying it back in the future; this is known as the no-arbitrage condition. As a result, coal usage goes up, despite the regulator's intention to reduce emissions by ending coal use sooner (Gerlagh 2011; Hoel 2012). Panel A of Figure B.1 illustrates this concept visually for coal extraction, where the announcement of the new policy is signified by the dark green dotted line.

The second way that the Green Paradox can increase emissions, which is novel to this paper, is what I will call the "depreciation channel." Conventionally, coal power plants have been thought of as "sunk costs" from the firm's perspective. Their construction costs have already been paid, so they should not impact the firm's forward-looking, profit-maximizing production decisions.

However, coal power plants are not exactly like sunk costs. This is perhaps easiest to see when considering the example of firms that produce electricity using a portfolio of energy resources, such as coal and natural gas. In making operating decisions about how much of each resource to use, firms weigh not only standard marginal costs but also the wear and tear on the capital stock itself. Each time a power plant generates electricity, it uses up a bit of this wear and tear, which effectively becomes a marginal cost that enters firms' production decisions. When regulation is expected to remain stable, electricity generators plan an optimal path to consume each type of power plant's stock of wear and tear over time. However, if regulation targets a specific resource, such as coal, the future marginal value of wear and tear on a coal plant decreases relative to that on a natural gas plant. This shift prompts firms to substitute coal-fired plants for cleaner natural gas plants, which are conserved for use after the regulation goes into effect.[6]

Imagine that, under normal circumstances, a power plant operator would use the wear and tear on their coal plant over twenty years to maximize profit. But now, the new regulation from above requires the plant to close in ten years. In this case, the operator would want to use up all of the remaining wear and tear before the closure deadline; otherwise, any remaining wear and tear becomes worthless once the plant shuts down. To consume this wear and tear faster, the operator could run the plant in ways that wear it down more quickly. Running a power plant typically causes wear and tear, but the effect accelerates when plants significantly adjust their power output, a process known as "ramping." Through this depreciation channel, the Green Paradox creates an incentive, beyond the price channel, for the plant operator to use the coal plant more intensively and ramp more frequently while it is still permitted.

Both channels operate simultaneously and reinforce one another. Through the price channel, a lower coal price reduces the marginal cost of coal-fired electricity, encouraging higher utilization of coal plants. Additionally, the power plant owner, aiming to use up all remaining wear and tear before the regulation takes effect, will further increase the plant's operating time. Therefore, following the regulation's announcement, we would expect to see

---

[6] The depreciation channel is not limited to firms with both coal and natural gas power plants; it also affects firms operating only coal power plants, which may adjust their bidding behavior in capacity markets. When the wear and tear on the plant becomes worthless in the near future, generators may be more inclined to turn the plant on and off frequently to meet demand, leading to more frequent bidding.

a rise in coal plant utilization, as illustrated in Panel B of Figure B.1. The extent to which generators can exploit these incentives—the magnitude of the Green Paradox—depends on their ability to shift production away from cleaner plants.

Testing for the Green Paradox in real-world data presents a challenge, as information about new policies often gradually escapes to the market. Major regulations are frequently informally leaked to gain feedback and may also undergo notice-and-comment periods or legislative debates. These disclosures make it difficult to pinpoint exactly when market participants adjust their expectations. However, we can observe the Green Paradox when a regulation is unexpectedly invalidated. Following such a shock, we would expect the Green Paradox effects to reverse as market agents return to their original trajectories, illustrated by the curves to the right of the dotted red line in both panels of Figure B.1. This paper adopts this approach by examining the unexpected stay of the Clean Power Plan, which is discussed in detail in the next section.

1.3 Institutional Background

This section introduces the Clean Power Plan and outlines how the regulation mandated emissions reductions in the power generation sector. Following this, I provide an overview of the litigation path of the CPP. Finally, I discuss the extent to which the issuance of the stay represents an unanticipated shock to market participants' expectations. Ultimately, the unprecedented nature of the stay makes for a clean event date.

1.3.1 History of the Clean Power Plan

Following the failure to pass cap-and-trade legislation in 2010,[7] the Obama Administration introduced an alternative approach on June 2, 2014, with the announcement of the Clean

---

[7] The dissolution of this effort is the focus of Lemoine 2017.

Power Plan by the EPA.[8] The regulation aimed to reduce GHG emissions from power plants by at least 32% by 2030. To achieve this, the EPA utilized federalism by setting individual emissions targets for 47 states, allowing each state to choose its policy mix to meet its assigned target.[9]

States were assigned targets using a four-step procedure (Plumer 2015). First, the EPA tallied the number of fossil fuel power plants in the state and their associated GHG emissions. Next, the EPA assessed how much power plants could cut using feasible technology at reasonable costs. Reductions largely came from operating coal plants more efficiently, increasing the operating time of natural gas plants, and increasing the renewable energy share. Third, using their assessment of reasonable emissions savings, the EPA applied those reductions to each power plant. Finally, the EPA summed up total emissions curtailment by state and subtracted them from their starting emissions to get a final target for all 47 states. As a result, states faced different percentage reductions based on their energy profiles, a feature I leverage in my empirical strategy.

Ultimately, these targets would be the final goal for 2030, but the first intermediate targets would become binding in 2022. A lengthy implementation period was necessary because, as one utility CEO stated "...if you had to comply...you'll be shutting down plants before you'd get [the] infrastructure in to replace them...who can site a transmission line in under 5 years? No one in this room. No one ever–anywhere can do that."[10]

Even though the CPP's binding targets were still years away, the logistics of electricity generation required immediate action. Electricity supply must continuously meet demand, and constructing new power plants can take years. Therefore, utilities needed to adjust their operating plans and capital expenditures well in advance of the CPP's requirements, which necessitated significant investments in natural gas and renewable generation. As

---

[8] The rule was finalized on August 3, 2015, and then published in the federal register on October 23, 2015. For a general overview of the rule and the relevant legal issues, see Hammond and Pierce Jr 2016.

[9] Vermont and Washington D.C. were exempt from the Clean Power Plan as they did not have large fossil fuel or electric plants. Hawaii and Alaska were also excluded due to their unique electrical grids.

[10] Fortis Inc. - Analyst/Investor Day (Oct. 6, 2015).

the CEO of Fortis put it succinctly on an earnings call in July of 2015 "...we are already taking significant actions related to our coal fleet ahead of the pending [C]lean [P]ower [P]lan announcement."[11] The Supreme Court's issuance of a stay suspended enforcement of the Clean Power Plan indefinitely, and with it, generators' incentives to increase the utilization of coal power plants. Whether the stay was anticipated by markets is a crucial point I address below.

1.3.2 Litigation Path of the Clean Power Plan

The Clean Power Plan enjoyed some industry support but was controversial, with the EPA receiving over 4.3 million comment letters regarding the proposed rule (U.S. Environmental Protection Agency 2015a; "Recent Regulation: The Clean Power Plan" 2016). Industry participants launched several unsuccessful and unusual legal challenges before the rule was finalized, each promptly dismissed by the courts.[12] These preemptive legal maneuvers allowed the EPA to anticipate and refine its legal arguments for the rule's final version. Finally, after the rule was published, the EPA faced further lawsuits aiming to block implementation of the CPP.

On February 9, 2016, the case *West Virginia v. EPA*,[13] took an unprecedented turn when the Supreme Court voted 5-4 to issue a stay against enforcing the Clean Power Plan. This marked the first time the Court had stayed a regulation before a lower court ruled on its merits (Liptak and Davenport 2016). Opponents of the regulation argued "the stay may well be a fatal blow..." (Wolf 2016). Figure B.2 provides a timeline of the legal developments of the CPP.

The issuance of the stay is a compelling event date for several reasons. First, the unprecedented nature of the Supreme Court's stay signaled to market participants that the CPP

---

[11] Fortis Inc., Q2 2015 Earnings Call, Jul 31, 2015.

[12] In re Murray Energy Corp., 788 F.3d 330, 333 (D.C. Cir. 2015), In re West Virginia, D.C. Cir., No. 15-1277, In re Peabody Energy Corp., D.C. Cir., No. 15-1284.

[13] 597 U.S. ___ (2022).

was unlikely to survive judicial scrutiny. In considering a stay, courts weigh, among other things, the probability of the party prevailing on the merits of the dispute.[14] By granting the stay, the Court indicated that the EPA was likely to lose, and the unexpected manner of the ruling highlighted the Court's confidence in that outcome. This undercuts the rationale for using the later release of the decision in *West Virginia v. EPA* as the event date since much of the relevant information would have been priced into the market years earlier.

The unexpected nature of the Supreme Court's stay provides advantages relative to using the date the rule was first announced. In the United States, administrative agencies often leak rules before formally proposing them, allowing for preliminary feedback through informal channels. Once proposed, a rule enters a formal comment period for additional public input.[15] The final rule is only published in the Federal Register and becomes effective no sooner than 30 days thereafter. This multi-step process creates numerous opportunities for information leaks, complicating efforts to pinpoint when expectations change. Consequently, event study methodologies around the rule's enactment are challenging due to the uncertainty about the timing of expectation shifts.

In contrast, the unexpected judicial intervention of the Supreme Court's stay offers a clear event date, as it lacks a comparable public process. Additionally, the stay was issued without oral arguments—a stage where judges typically provide hints about their stance—further preventing any information leakage. Analyzing the market response to this sudden change in expectations, particularly in commodity and electricity markets, offers cleaner causal identification of the Green Paradox than currently exists in the literature.

Moreover, the timing of the Supreme Court's stay, near the end of the Obama administration, reduced perceptions of the rule's survival chances. Historically, since World War II, the incumbent party of a two-term president has won the subsequent election only once, in-

---

[14] A party seeking a stay must show that: (1) they will likely prevail on the merits; (2) they will suffer irreparable injury if the stay is denied; (3) the other parties will not be substantially harmed; and (4) the public interest will be served by granting the stay. Long v. Robinson, 432 F.2d 977 (4th Circuit, 1970).

[15] About Regulations.gov, Regulations.gov, `https://www.regulations.gov/learn` (last visited Nov. 4, 2024).

creasing the Republicans' odds of winning office. If the Supreme Court invalidated the CPP, the Obama administration would not have had time to propose a new rule incorporating the Court's decision, nor would the Democrats likely have been able to push through a legislative fix. However, this does not imply that market participants expected the Republican candidate to simply repeal the Clean Power Plan. Had the Supreme Court upheld the CPP, the rule would likely have remained binding regardless of the election outcome. The downfall of the Affordable Clean Energy Rule—the Trump-era replacement for the CPP detailed in the Appendix—demonstrates the stickiness of administrative rules, as this attempt to alter the CPP was invalidated by the Court.

One potential drawback to using the stay as an event date is the timing of Justice Antonin Scalia's death on February 13, 2016, which left a vacancy on the Supreme Court (Liptak 2016). With Scalia's passing, the President would need to nominate a successor, who would then require Senate confirmation. At that time, the Senate was controlled by Republicans, while the White House was held by Democrats, creating a scenario where any nominee would need to appeal to both sides. If market participants believed that Scalia's death increased the odds of the CPP's survival, this could potentially bias the results toward zero.

However, this concern is mitigated for several reasons. First, the Republican-controlled Senate ultimately refused to fill the vacancy before the 2016 election, increasing the likelihood that the Clean Power Plan would not survive. This decision is included in the event period, meaning any initial market reaction to Scalia's death would be offset by the later Senate decision, which should be observable in the event time plots as a delayed effect. Ultimately, the impact of Scalia's passing is an empirical question. Figure B.18 and Table C.7 in the Appendix show that abnormal returns after his death align with expectations of a reduced likelihood of the Clean Power Plan's survival.[16] On balance, the stay provides an unexpected

---

[16] Bakke et al. 2023 provides some evidence supporting this hypothesis, noting negative returns for Republican-leaning firms upon news of Scalia's passing. However, their measure is broader than just coal mining firms and uses close-to-close returns, which might capture additional factors beyond the Clean Power Plan itself.

shock to expectations with minimal downsides and serves as the event date for the remainder of this paper.

1.4 Data

First, I collect data from the Public Utility Data Liberation Project (PUDL) to calculate the merit order and the realized utilization of power plants. Using the methodology from Cicala 2022, I construct the merit order and determine out-of-merit generation for each hour. Hourly demand by balancing authority is obtained from FERC Form 714, which tracks electricity demand every hour during the event period. Calculating hourly generation presents additional complexities.

Hourly data on fossil fuel power plant operations and emissions is sourced from the Clean Air Markets Program Database's (CAMPD) Continuous Emissions Monitoring System (CEMS), which covers 96% of power plant emissions.[17] CAMPD stems from federal emissions trading programs like the Cross-State Air Pollution Rule and the Acid Rain Program and applies to all plants with a generating capacity of at least 25 megawatts. In addition to emissions data, CEMS provides hourly gross generation and the percentage of the hour that a fossil fuel plant operates. I calculate a conversion ratio for gross generation using monthly net generation from the U.S. Energy Information Administration (EIA) Form 923 to estimate hourly net generation, accounting for energy consumption and losses in the generation process. Net generation is the amount of output that can be used to meet demand, it represents the amount of electricity generated after meeting the power plant's own uses and losses from equipment.

Additional data from EIA Forms 923 and 860 in PUDL provides fuel usage and fuel costs, enabling the calculation of heat rates and marginal costs for each generator.[18] Figure B.3

---

[17] United States Environmental Protection Agency (EPA)."Clean Air Markets Program Data." Office of Atmospheric Protection, Clean Air Markets Division. https://campd.epa.gov/.

[18] All power plants connected to the grid with a production capacity of more than one megawatt are required to file Form 923 annually. As of 2013, plants powered by non-coal fossil fuels must report fuel

plots the total monthly value of fuel purchased for coal-, gas-, and oil-fired power plants as reported in Schedule 2. While coal usage has fallen over time, it remained a significant source of electricity generation during the 2015 to 2016 period.

For non-fossil fuel plants, I use various data sources to allocate monthly generation to the hourly level. Solar generation is estimated by using hourly solar irradiance data—measured by satellite at the panel coordinates—to allocate monthly net generation based on each hour's share of total monthly irradiance. Wind and hydroelectric generation are estimated using corresponding data, such as streamflow readings from the closest downstream meter for hydro plants and measured wind speeds for wind plants. Nuclear power generation is distributed evenly across hours, using daily output percentages provided by the U.S. Nuclear Regulatory Commission (NRC).

Figure B.4, inspired by Cicala 2022, presents electricity generation data. The U.S. produces over 6 TWh (terawatt-hours) of electricity per day. Panel A shows seasonal fluctuations in electricity usage, with demand peaking in the summer. Panel B illustrates average hourly electricity usage across a week, revealing consistent weekly patterns. To address the seasonality shown in Panel A, I include flexible weather controls, specifically average temperature and precipitation. For the pattern in Panel B, I use unit of analysis × day-of-week × hour fixed effects, where the unit of analysis is either the emissions unit, the power plant, or the utility. The fixed effect can be thought of as comparing the change in behavior at a given power plant at noon on Wednesdays over time.

Third, CO_2 emissions data are sourced from CEMS. Facilities are required to install sensors that record hourly emissions of sulfur dioxide, carbon dioxide, and nitrogen compounds. These sensors undergo regular calibration and audits by the EPA and state regulators. Several studies confirm the accuracy of these emissions reports relative to satellite or aerial measurements (Cusworth et al. 2021; Cusworth et al. 2023; Velazco et al. 2011). Even if

---

purchases if they have a capacity of at least 200MW. Nearly all natural gas power plants exceed this threshold (`https://www.eia.gov/todayinenergy/detail.php?id=38312`). The threshold for coal plants is 50MW (`https://www.eia.gov/electricity/monthly/pdf/technotes.pdf`).

reporting behavior changes, the empirical strategy remains robust so long as changes do not coincide with the stay's issuance. I also use power plant characteristics from the Emissions & Generation Resource Integrated Database (eGRID) and restrict the analysis to plants that did not change fuel sources during the event period.

A quick plot of the raw CEMS data in the twelve weeks around the stay provides initial evidence consistent with the Green Paradox and the depreciation channel. Panel A of Figure B.5 plots the mean hours spent ramping in each hour by coal power plants, the treated group, in yellow, and natural gas power plants, the control, in green. We can see that coal power plants ramp twice as frequently before the stay, but the difference between the two groups collapses shortly after the event. Panel B of Figure B.5 plots the mean hourly CO_2 emissions for treated and control plants for the twelve weeks on either side of the stay. We similarly see a marked decrease in the level of emissions from coal-fired power plants.

Fourth, I obtain data on PM2.5 pollution from Wei, Xing, Shtein, et al. 2022, which provides daily PM2.5 concentrations at the zip code level using a combination of monitoring stations, satellite data, meteorological information, and chemical transport models. The data are initially measured on a 1 km by 1 km grid, with concentrations averaged for each zip code based on the grid's centroid location, producing a measure of PM2.5 in micrograms per cubic meter ($\mu g/m^3$) for each zip code. I winsorize the data at the top and bottom 1% to mitigate potential errors. Weather data are sourced from the National Oceanic and Atmospheric Administration (NOAA) at the zip code-day level.

Diesel vehicle traffic is a significant source of PM2.5 emissions. To account for this alternative source, I include daily traffic counts from the Department of Transportation's Travel Monitoring Analysis System (TMAS).[19] TMAS data are collected from monitoring stations along the nation's highways, recording vehicle counts at the lane-direction-station level, and are aggregated to the zip code-day level using Census shapefiles. For the approximately 1.6%

---

[19] Note that Maine did not report traffic data for February and May of 2016; all other states reported during the event window. The data can be accessed here: https://www.fhwa.dot.gov/policyinformation/tables/tmasdata/.

of traffic observations falling outside a zip code, I assign them to the nearest zip code.

In an additional test in the Appendix, I use coal extraction data from the EIA, which is reported at the state-week level. I aggregate this data across all states to produce a national-week total. Also in the Appendix, I analyze the Green Paradox using market prices. Return data are sourced from CRSP, and factor returns are accessed via WRDS. To ensure the robustness of my results, I screen for confounding events, such as earnings announcements, using Factiva.

Table C.1 presents summary statistics for the samples used in the main specifications (i.e., out-of-merit generation, PM2.5 concentrations, and CO_2 emissions). As expected, treated and control facilities reveal several differences, with treated observations producing more CO_2 than the control units. We can also see that, consistent with ramping being costly to generators (Kumar et al. 2012; Reguant 2014; Borrero, Gowrisankaran, and Langer 2024), it is unlikely to occur in any given hour.

1.5 Research Design and Results

This section introduces the research design and presents the results for each test used to assess the presence and magnitude of the Green Paradox. First, it examines the paradox's impact on out-of-merit generation to demonstrate the existence of the depreciation channel. Next, it analyzes the impact of the paradox on PM2.5 pollution, which accounted for most of the benefits of the Clean Power Plan. Then, using hourly power plant data, it investigates the effect on CO_2 emissions, offering a means to validate the PM2.5 findings under the assumption that the production function for both pollutants is relatively stable. The section concludes with a series of robustness checks, confirming the primary results.

1.5.1 The Presence of the Depreciation Channel

A key contribution of this paper is extending the analysis of the Green Paradox beyond the price channel to examine how it influences firms' asset utilization. In the case of electricity generation, power plants are long-lived assets with substantial useful life years remaining. When regulations accelerate a plant's retirement, owners face an incentive to increase their utilization of affected plants, since any remaining wear and tear will be valueless upon the regulation's implementation date. More short-run production at polluting plants leads to increased emissions during the regulation's implementation period—a dynamic introduced above as the depreciation channel.

Demonstrating the presence of the depreciation channel requires two steps. The first requires an analysis showing that changes in fuel prices cannot fully explain the change in generator behavior. To do so, I look at changes in out-of-merit generation after recalculating marginal costs each month to account for changes in fuel prices. If there are non-input price factors shifting behavior, they should show up as changes in out-of-merit generation before and after the stay. Second, it must be shown that coal power plants are shifting their behavior to consume less economic depreciation after the stay. I do this by comparing changes in ramping between treated and control plants after the event.

Changes in Out-of-Merit Dispatch

A measure that separates the depreciation from the price channel requires that changes in the measure of generator behavior cannot be attributed to price changes. Ideally, we would observe an unexpected shift in expectations of future environmental regulation that leaves prices unchanged but increases incentives for affected power plants to consume their remaining wear and tear. I approximate this ideal test by examining deviations from merit order dispatch and comparing the extent of these deviations before and after the stay of the CPP.

The merit order represents the lowest-cost supply curve, ranking power plants from lowest to highest marginal cost of generation, with electricity prices determined by the marginal cost of the last plant required to meet demand. Panel A in Figure B.6 illustrates this concept graphically: the vertical red line shows demand, while the marginal plant, shaded in grey, sets the price.

Calculating the merit order for power plants is straightforward, as it depends on each plant's marginal cost, which is determined by its efficiency (known as its heat rate) and fuel prices. This relationship allows me to separate the price channel from the depreciation channel. Since marginal costs are calculated based on the actual prices paid for fuel, the price channel's effect of reducing coal prices would, in turn, lower the marginal cost of coal-fired generation under the Green Paradox. With lower marginal costs, coal generators would move up in the merit order, so part of their increased generation aligns with the merit order once the paradox takes effect.

However, the second mechanism, the depreciation channel, does not enter into the merit order calculation because the remaining wear and tear is not captured by the plant's heat rate or fuel prices. The depreciation channel does lower the generator's marginal cost, so it should lead to more generation by the plant beyond what one would expect due to price. Because the merit order accounts for the price effect, but not the depreciation channel, this increased generation appears as larger deviations from the merit order, which is called out-of-merit generation. By comparing realized dispatch with the merit order, I control for price effects, ensuring that any increase in out-of-merit generation reflects the depreciation channel.

Panel B of Figure B.6 provides an example of out-of-merit generation, where a higher-cost unit (shown in yellow on the left) produces electricity despite not being the lowest-cost option. In this scenario, the generator may incur a loss, raising the question of why it would operate at all. One reason is that power plants take time to start up and shut down, and during this transition period, they continue generating electricity. During such times, it is more profitable for the plant owner to sell the electricity rather than let it go to waste. Vertically integrated utilities have an additional incentive, as they can pass losses onto customers, making them less sensitive to cost differences between generation sources. These incentives for out-of-merit generation are always present, regardless of the Green Paradox. In other words, these non-depreciation incentives are relatively constant over short periods of time.

To test for the depreciation channel, I assume that these other factors remain unchanged apart from the shock to expectations from the staying of the CPP. I then examine differences in out-of-merit generation before and after the stay, a method known as an event study. By comparing out-of-merit generation in this way, the event study effectively filters out stable factors, leaving temporary changes—such as those driven by the depreciation channel—to explain any observed differences.

Before delving into the analysis, I'll briefly explain the construction of the merit order. First, I calculate heat rates and fuel prices at the generating unit-month level using data from the EIA on monthly net generation and fuel consumed.[20] Since fuel prices are reported monthly, the merit order is updated accordingly each month. Within each balancing authority, generators are ranked by marginal cost, and cumulative net generation is summed until it meets demand. I derive hourly renewable generation following the methodology outlined in the Data section, assigning it a marginal cost of zero. For nuclear energy, I assume that plants are required to consistently supply the grid due to their role in providing continuous electricity. The remaining demand after subtracting out non-fossil fuel sources is the amount of demand that must be met by fossil fuels, which dispatch according to marginal cost.

To measure inefficiency, I sum the megawatts of out-of-merit generation for each balancing authority. Balancing authorities are responsible for ensuring that electricity supply and demand are balanced within specific areas and can direct plants to adjust operations to maintain service (U.S. Department of Energy 2023). Both the CEMS data on plants and

---

[20] All heat rates are adjusted for steam generation.

the EIA Form 714 data on balancing authorities provide hourly reports, enabling analysis at the balancing authority-hour level. The merit order sample spans from 2014 through 2017, allowing deviations to be evaluated within a multi-year context. Panel A of Figure B.7 plots the mean out-of-merit generation for each day in the sample. While the data are noisy, there is an extended level of out-of-merit generation between late 2015 and early 2016 while the Clean Power Plan was in place.

Panel B of Figure B.7 displays the total out-of-merit generation per hour, adjusted for hour × day-of-week and day-of-year × balancing authority fixed effects to eliminate seasonal influences. Additionally, it accounts for population-weighted daily average temperature and precipitation. To improve visualization, the total out-of-merit generation is smoothed.[21] The vertical red line marks February 9, 2016—the date the stay was issued—occurring just after a local maximum in out-of-merit generation. The dotted yellow line represents the mean out-of-merit net generation before the stay, while the dotted grey line shows the mean after the stay.[22]

The Clean Power Plan was introduced in June 2014, during which time out-of-merit generation remained relatively stable, slightly above zero (Panel B of Figure B.7). However, after the final rule was announced in late 2015, there was a noticeable rise in out-of-merit generation. This increase likely reflects that market expectations shifted towards implementation of the CPP and that the depreciation channel began to drive additional coal plant utilization as operators anticipated future regulation. Following the stay, there is a clear drop in out-of-merit generation, with a difference in means of approximately 1,000-megawatt hours (MWh), supporting the idea that the stay tempered the acceleration in coal plant depreciation.

In the pre-stay period, the estimates are generally close to zero; however, after the stay, there is a clear decline. Table C.2 presents the average effect using a single post-stay indi-

---

[21] Smoothing is done via restricted maximum likelihood.

[22] Anecdotally, the Rocky Mountain Institute has reported similar findings, indicating that coal plants continued operating out of merit until after the stay was issued, `https://utilitytransitionhub.rmi.org/economic-dispatch-operations/`.

cator, with an observed impact of approximately 7.5%. This effect size aligns with findings from the PM2.5 and GHG emissions tests below, providing further evidence for the existence of the Green Paradox. To illustrate the dynamics of the change in out-of-merit generation, Figure B.8 plots the results of regressing the specification from column (4) in Table C.2 on event-time dummy variables.

Overall, the out-of-merit generation results indicate that changes in coal prices alone cannot fully explain the shifting behaviors of power plant owners. The next section tests for the presence of the depreciation channel by directly measuring changes in power plant operations that increase consumption of wear-and-tear.

Changes in Depreciation-Inducing Behavior

Given the evidence above that prices alone cannot explain shifts in generators' behavior, this section directly tests for operational changes that cause economic depreciation of power plant assets. One behavior that increases wear-and-tear is called "ramping." Ramping is when power plants change the amount of electricity they generate rather than running at constant capacity. The effect on economic depreciation is not dissimilar to the difference between city and highway miles on an automobile. When driving a car in the city, there is often stop-and-go traffic, which produces additional economic depreciation on the vehicle. This effect is similar to ramping at power plants. Conversely, when driving a car on the highway, the wear-and-tear per mile is less than that of city driving. This reduced depreciation is similar to running a power plant near max capacity.

Ramping is also connected with out-of-merit generation in two ways. First, when a power plant starts up or shuts down, it will not be operating as efficiently as it would be at full capacity. Reduced efficiency means that generators require a higher price per unit of output, and that higher price per unit, at least temporarily, is recorded as out-of-merit generation. Second, the depreciation channel reduces the private marginal cost of operating the power plant. However, this private marginal cost is not represented in the merit order calculation, which includes only heat rate and fuel costs; so, to the extent that the depreciation channel increases utilization above and beyond the price channel, it shows up as out-of-merit generation. The out-of-merit results from the previous section include both effects.

I operationalize the depreciation channel test using the proxy for ramping following Borrero, Gowrisankaran, and Langer 2024. $Ramp\_MinMax$, is an indicator variable set to one when a power plant deviates by at least 10% from its minimum or maximum generation capacity. The ideal experiment would involve altering future pollution restrictions on a randomly selected subset of power plants, then comparing changes in ramping behavior to plants not subject to new regulations. I approximate this approach by leveraging the effectively random timing of the Clean Power Plan stay as a measure of changing regulatory burdens. Since the CPP primarily targeted coal plants with significant regulations, while imposing only minor changes on natural gas plants, I use coal plants as the treated group and natural gas plants as the control.[23] By comparing the differences in outcomes between these two plant types before and after the stay, I can estimate the magnitude of the Green Paradox.

Specifically, I use a difference-in-differences design of the following form to compare differences between the treated and the control group after the stay:

$$\text{Ramp\_MinMax}_{i,t} = \beta_1 \text{Treated}_i \times \text{Post}_t + \beta_2 \text{Treated}_i \times \text{X}_{i,t} + \beta_3 \text{X}_{i,t} + \gamma + \alpha \tag{1.1}$$

In this model, $Ramp\_MinMax_{i,t}$ represents the ramping proxy for generating unit i at hour t. The variable $Treated_i$ is an indicator for whether a unit is treated (i.e., a fossil-

---

[23] Although natural gas plants were included in CPP regulations, they played a critical role in the CPP's strategy to reduce GHG emissions, unlike coal plants, which were largely slated for retirement. Natural gas plants are markedly less polluting than coal, so increasing generation from natural gas would lead to large reductions in PM2.5 and CO_2 emissions. The CPP sought to incentivize increased utilization and capacity from natural gas.

fuel-powered plant that does not use natural gas, nearly all of which are coal-fired). $X_{i,t}$ includes population-weighted average temperature and precipitation at the state level, with weights based on each zip code's total population within the state. I include these weather controls because they likely impact electricity demand as seen in panel A of Figure B.4, an effect that should not be misattributed to the Green Paradox. Additionally, I interact these weather controls with the treatment variable to account for differences in coal and natural gas power plants' responses to weather changes. The parameter γ represents state fixed effects, which control for state-level policies and broader shocks (e.g., economic slowdowns). The parameter α includes unit of analysis × day-of-week × hour (i.e., generator, power plant, or utility) fixed effects to account for the cyclical patterns shown in panel B of Figure B.4. The sample period spans from 12 weeks before to 12 weeks after the stay.

Table C.3 shows the results. Across the columns, we can see a pretty stable estimate of around a 1.1% reduction in the probability of ramping in any given hour for treated power plants. Interestingly, Columns (7) and (8) highlight that the effect is most concentrated among the top decile of most polluting plants (signified by Top Decile) and in states with above median reductions under the Clean Power Plan. Put differently, the changes in consumption of economic depreciation are strongest for coal power plants where the regulation was more binding, consistent with what would be expected if the regulation was driving the behavior. To assess the plausibility of the parallel trends assumption I provide an event plot in Figure B.9 by interacting the treated variable with event time dummies. The pre-period estimates are all insignificant except for the Christmas-New Year's week at $t = -6$, and we can see a reduction in total ramping before reverting to a similar gap by week 12.

1.5.2 Changes in Power Plant Emissions

This section provides reduced-form estimates of the changes in air pollution and carbon emissions from the Green Paradox. It first looks at PM2.5 emissions, which made up the bulk of the benefits from reducing coal utilization. Next, it measures changes in CO_2 production by regulated power plants.

Changes in PM2.5 Emissions

The EPA estimated the Clean Power Plan's benefits at $26-$45 billion annually, including the prevention of thousands of deaths, hundreds of thousands of asthma attacks, and nearly half a million missed school or work days (U.S. Environmental Protection Agency 2015a). These benefits stem primarily from anticipated reductions in PM2.5 pollution rather than greenhouse gases like CO_2. PM2.5 particles, measuring less than 2.5 micrometers, can penetrate deep into the lungs, where they cause significant health issues.[24] PM2.5 pollution originates from both primary sources, which directly emit these particles, and secondary sources, which release gases that transform into PM2.5 in the atmosphere. Coal power plants contribute to both types of particulate pollution.

Ideally, I would directly measure both types of PM2.5 and randomly impose additional regulation on certain power plants; however, this approach is impractical. While direct PM2.5 measurements from smokestacks would capture primary emissions, they would miss secondary emissions formed through atmospheric chemical reactions. To account for both sources, I use daily PM2.5 concentration data at the zip code level, derived from satellite measurements, and introduce a novel technique to predict the path of emissions over time, enabling me to attribute them to their source.

Accurately tracking secondary PM2.5 precursors and the movement of primary PM2.5 particles across regions is essential for correctly attributing pollution to individual zip codes. Using an approach that compares zip codes with coal power plants (the treated group) to those without (the control group) could lead to biased estimates due to the spread of PM2.5 particles by wind and other factors. For instance, if a coal plant in Chicago emits PM2.5,

---

[24] Particulate Matter (PM) Pollution, Centers for Disease Control and Prevention, `https://www.cdc.gov/air/particulate_matter.html` (last visited Nov. 4, 2024).

prevailing west-to-east winds might carry some of these pollutants to Gary, Indiana, where there are no local emission sources. In this case, designating Chicago as the treatment unit and Gary as the control could produce biased results: pollution from Chicago drifting into Gary would blur the difference between them, thereby diminishing the observed effect. This is called a spillover effect, and a sufficiently robust research design should control for them.

To account for spillover effects, areas like Gary should be treated as indirectly affected units. Identifying these spillovers can be challenging. One approach is to define treated units as all zip codes within a specified radius of a coal power plant. However, this method can still introduce bias. For instance, if Milwaukee falls within this radius but is unaffected by west-to-east winds due to its location north of Chicago, the estimated treatment effect may again be biased toward zero. To control for spillovers, I adopt an approach similar to Hernandez-Cortes and Meng 2023, using NOAA's Hybrid Single-Particle Lagrangian Integrated Trajectory (HYSPLIT) model (Stein et al. 2015) to identify spillover zip codes.[25] This atmospheric dispersion model effectively balances computational efficiency with accuracy in tracking long-distance particle dispersion and has been widely applied in studies of PM2.5 emissions from coal power plants (e.g., Yerramilli et al. 2012; **HENNEMAN2019271**). HYSPLIT models the atmospheric trajectory of a particle released from a specific location and height, accounting for weather conditions (such as rain, which can bring particles back to Earth) and exposure to solar radiation, which can cause particle decay.

I use HYSPLIT as follows: For each U.S. coal power plant reporting to the EIA, I enter its coordinates and smokestack height into the HYSPLIT model. I then simulate 24-hour forward particle trajectories with four daily emissions releases, tracking particle locations at hourly intervals. This simulation is repeated daily throughout the sample period. A zip code is classified as treated if it contains at least one coal power plant, while it is designated as a spillover unit if HYSPLIT simulations show it was exposed to emissions from a coal power plant in the previous 24 hours.

---

25. I implement HYSPLIT via the splitR package in R.

Figure B.10 shows HYSPLIT's output for Illinois's largest coal plant, the Prairie State Energy Campus, during January 2016. The trajectories are color-coded by week, and for each day of the simulation, a particle is released every six hours, with each trajectory traced hourly over 24 hours. Each dot in Figure B.10 represents the hourly position of an emission particle. The wide dispersion of paths across weeks underscores the importance of accurately tracking PM2.5 travel patterns and the limitations of classifying zip codes as spillovers fixed over time. For instance, Gary, Indiana qualified as a spillover unit on only a single day in January 2016.

Since I cannot randomly assign treatment to coal power plant zip codes, I approximate this approach by leveraging the effectively random timing of the stay of the CPP. Given the unexpected nature of the stay, it likely acted as a shock to market participants. By comparing outcomes in zip codes with coal power plants, which were affected by the regulation, to those without, we can estimate the Green Paradox effect as the change in differences between these two groups before and after the stay. This method is known as a difference-in-differences design.

Specifically, to identify the effect of the stay on particulate emissions, I run a difference-in-differences model of the following specification:

$$\text{Log PM } 2.5_{i,t} = \beta_1 \text{Treated}_i \times \text{Post}_t + \beta_2 \text{Spillover}_{i,t} + \beta_3 \text{Treated}_i \times \text{X}_{i,t} + \beta_4 \text{X}_{i,t} + \gamma + \alpha \quad (1.2)$$

where $Log\ PM2.5_{i,t}$ is the log concentration of PM2.5 in micrograms per cubic meter, for zip code i at time t. The variable $Treated_i$ is an indicator set to one if a zip code contains a coal plant, while $Spillover_{i,t}$ is an indicator set to one if a trajectory entered a zip code's geographic boundaries in the past 24 hours. The vector $X_{i,t}$ includes time-varying controls, such as average temperature and precipitation for each zip code, which I interact with the treated variable. I control for these weather variables for two reasons: first, to account for demand fluctuations unrelated to the Green Paradox, thereby isolating the effect of interest; and second, because weather—especially precipitation—can significantly influence atmospheric PM2.5 levels, effects that should be removed from the estimated effect. In one specification, I also include the log of daily traffic volume, as diesel engines are a major source of PM2.5. Additionally, α denotes zip code or zip code × week fixed effects, a topic I return to below.

Figure B.11 displays the sample average PM2.5 concentration for each ZCTA in the dataset.[26] The figure reveals higher PM2.5 concentrations in industrial areas, such as the Rust Belt and the Eastern Seaboard, compared to more rural regions like Montana. The darkened area in California represents the Central Valley, where particulate matter often becomes trapped between mountain ranges. These patterns highlight the importance of controlling for both observed and unobserved zip code-specific factors through zip code fixed effects. Additionally, seasonal variations—such as plant closures on holidays or prescribed burns—are more accurately captured with zip code × week fixed effects.

Finally, it's important to note that natural gas plants also produce PM2.5 pollutants, although at a significantly lower rate than coal plants. Since it is not possible to separately measure PM2.5 emissions by source within a zip code, any increase in operating time by natural gas plants relative to coal plants could bias my results toward zero. Consequently, my estimate of the Green Paradox effect may be conservative.

The results of the particulate matter tests are in Table C.4. The positive and significant coefficients in columns (2), (3), and (4) suggest that the spillover group assignment was effective. However, when zip code × week-of-year fixed effects are included, the coefficient becomes statistically insignificant, indicating that spillover effects are consistent within a given week. Traffic volume is also an important control variable, with a 50% increase in traffic equating to a 2% rise in PM2.5 pollution. Precipitation is also a key control, as it tends to remove particles from the atmosphere through a process known as scavenging (Ikeuchi, Murakami, and Watanabe 2015).

---

26. ZCTAs are used in Figure B.11 instead of zip codes due to the lack of shapefiles at the zip code level.

Column (5) of Table C.4 is the preferred specification, as it incorporates policy differences, local economic changes such as reduced production during the Christmas-New Year period, and other unobserved shocks like crop burning or wildfires through zip code × week fixed effects. According to this specification, I estimate that the stay caused a roughly 6.5% reduction in PM2.5 concentrations in zip codes containing coal power plants. This reduction amounts to approximately 0.40 $\mu g/m^3$ at the sample mean. Column (6) presents the results after controlling for the log volume of traffic measured by TMAS. Although this adjustment reduces the sample size considerably, the effects in the treated zip codes become more pronounced and remain marginally significant. In summary, the Clean Power Plan appears to have increased short-run PM2.5 emissions, resulting in meaningful economic and health losses.

Finally, to assess the validity of the parallel trends assumption, I aggregate the data to the zip code-week level and interact the treated variable with weekly indicators in event time. These results require a caveat. The untreated group is changing throughout the sample period because I allow the spillover treatment variable to turn on and off each day. Therefore, the control group will shift slightly between periods, and drops one interaction term due to collinearity, but it should still provide some information on whether parallel trends are reasonable. Figure B.12 presents the results. There does not appear to be a trend in the pre-period, and we see a slight delay before PM2.5 concentrations fall. As shown below, this delay is consistent with the power plant-level results below.

Changes in CO_2 Emissions

Next, I estimate the effect of the Green Paradox on direct measures of power plant operations. As in the depreciation channel tests, the ideal experiment would involve altering future pollution restrictions on a randomly selected subset of power plants, then comparing changes in utilization and CO_2 emissions to plants not subject to new regulations. I approximate this approach by leveraging the effectively random timing of the Clean Power Plan stay as a measure of changing regulatory burdens. Since the CPP primarily targeted coal plants with significant regulations, while imposing only minor changes on natural gas plants, I use coal plants as the treated group and natural gas plants as the control.[27] By comparing the differences in outcomes between these two plant types before and after the stay, I can estimate the magnitude of the Green Paradox.

Specifically, I use the same difference-in-differences design as utilized in the test of changes in ramping behavior. This analysis focuses on the primary goal of the Clean Power Plan: reducing coal-fired electricity generation and associated greenhouse gas emissions. According to the Green Paradox, we should expect coal plant utilization and emissions to decrease following the issuance of the stay. My objective is to estimate the net effect of this phenomenon—the metric most relevant to policymakers who must account for any potential substitution between coal and natural gas—rather than simply tracking changes in coal electricity generation. Accurate identification relies on the assumption that, absent the CPP stay, coal and natural gas plants would have followed parallel trends over the event window.

One potential threat to the assumption of parallel trends is the differing economics of natural gas and coal-fired power plants. While visual inspection of the event studies provides evidence supporting this assumption, there is also a conceptual basis for believing that parallel trends are plausible. Although natural gas and coal-fired power plants have different emissions levels, this study focuses on changes in CO_2 emissions rather than absolute levels. Operational differences, such as the ability of these plants to start up and shut down quickly, are accounted for by fixed effects in the model, which control for predictable changes in demand and plant characteristics. Moreover, if firms are more willing to incur start-up costs in anticipation of future regulation, this change is part of the causal effect the study

---

[27] Although natural gas plants were included in CPP regulations, they played a critical role in the CPP's strategy to reduce GHG emissions, unlike coal plants, which were largely slated for retirement. Natural gas plants are markedly less polluting than coal, so increasing generation from natural gas would lead to large reductions in PM2.5 and CO_2 emissions. The CPP sought to incentivize increased utilization and capacity from natural gas.

aims to isolate. Given that the CPP would have led to the early retirement of coal plants, theory suggests we should observe increased utilization of these plants as the value of future depreciation diminishes, thus drawing more of their physical depreciation into the current period.

Finally, I exploit two measures of cross-sectional variation. First, I use the EPA's state-specific emissions reduction targets, as detailed in Section III, to create an interaction term called "Most Treated." This variable equals one for power plants in states with emissions reduction targets above the median. Thus, this test exploits variation in the extent to which the CPP was likely to bind the operations of the power plants in a given state. Therefore, treated plants in these states should have been slightly more affected than their peers.

Next, I examine whether the Green Paradox's incentives vary by plant pollution levels. To explore this, I interact "Top Decile," a variable indicating whether a power plant was among the top 10% most polluting before the CPP, with the treated and post terms. This interaction captures how the regulation's impact may be more pronounced for the dirtiest power plants.

The result, looking at changes in metric tons of CO_2, is in Table C.5. Treated power plants reduce their emissions by about 7.7% on average (Column 5). The top decile of generating units (Column 7), on average, reduces emissions by much more than the treated units (15%). There is also a larger effect in the states where the CPP was more binding, with reductions of over 9% (Column 8).

To visually investigate the plausibility of the parallel trends assumption, I also interact the treated indicator with dummy variables for each week to look for visual evidence of a violation of the parallel trends assumption. Figure B.13 plots the results. There is a slight delay before coal power plant emissions stabilize at a lower level, and there does not appear to be a violation of the parallel trends assumption.

To summarize, the GHG emissions tests find that power plants increased their emissions under the Clean Power Plan. Encouragingly, the magnitude is similar to the estimates in the PM2.5 tests. Given a constant pollution production function, this suggests that any potential spillovers between treated and control plants are likely small in magnitude. Encouragingly, these results are consistent with several robustness checks.

1.5.3 Robustness Checks

I include a battery of robustness checks to rule out alternative explanations. These checks involve examining media sources for state-level policy changes, performing a leave-one-state-out analysis, and conducting placebo tests with random assignments to the treatment group. The results remain robust across these additional tests.

First, I address the possibility of concurrent events affecting coal around the time of the stay. To mitigate this concern, I searched Factiva and UtilityDive for news events impacting coal during the event window. I identify three instances of state-level regulatory changes affecting utilities.

New Jersey and Oregon approved state-level laws to incentivize green energy (Walton 2016; Shallenberger 2016). These laws were intended to reduce coal use and, therefore, can be viewed as state-level Clean Power Plans. Since these laws were passed after the stay, their inclusion in the above tests should bias the results towards the null. Specifically, coal plants in New Jersey and Oregon did not have the same incentives to reduce usage as plants in other states after the stay. Consequently, if these state laws were perceived as binding, coal power plants in these states would have exhibited less change in their operating behavior compared to those without such laws.

Additionally, in March 2016, Ohio passed a law guaranteeing income for eight years for coal and nuclear plants operated by FirstEnergy and AEP (Shallenberger and Bade 2016). Like New Jersey and Oregon, the Ohio law could bias the results toward the null. However, this law overlaps with only a few days of the post-period. From the perspective of the parallel trends assumption, coal subsidies should reinforce the assumption's plausibility by making coal more competitive, thereby partially offsetting the secular rise in natural gas due to its cost advantages.

To validate my findings, I estimate models excluding observations from New Jersey, Ohio, and Oregon. Table C.6 presents the results.[28] The inferences are unchanged. Additionally, I perform a leave-one-out analysis, dropping one state at a time and estimating the treatment effects for both the intensive and extensive margins of power plant adjustments. The resulting coefficients, plotted in Figure B.14, are all negative and closely clustered, indicating that no single state and its policies are driving the results.

A second worry is that from 2015 until the spring of 2016, the Chinese economy experienced a slowdown, which led to falling energy prices (Irwin 2018). To minimize the impact of this macroeconomic shock, I use short event windows and control for its impacts using fixed effects where possible. Given that the stay was issued towards the end of this slowdown, any effects from the economic downturn would likely bias the results against finding evidence of the Green Paradox. Specifically, the slowdown in the latter half of 2015 would reduce demand for electricity (and its inputs) in the pre-period, while the subsequent upswing in the spring of 2016 would increase demand, leading to higher utilization of all power plants in the post-period. However, this macroeconomic shock poses a threat to the identification strategy only if it differentially affects treated (coal-fired) plants compared to control (natural-gas-fired) plants. Economy-wide shocks should impact both types of plants similarly after controlling for their specific characteristics. Further, the out-of-merit results suggest that the paradox was in effect beyond the price impacts that moves in the Chinese economy may have produced.

A related concern is that the results reflect the secular decline in coal, rather than the CPP. Although there has been a shift from coal to natural gas power plants, this shift does

---

[28] I do not drop Maryland following Hughes v. Talen, as it affected bidding for a single plant, and the impact of its holding was uncertain. The decision also happens after the end of the weekly tests.

not necessarily undermine the empirical approach. New power plants involve significant capital investment, and their adoption tends to be uneven over time (See Figure B.3). To account for long-term shifts in generation capacity, I focus on variation within short event windows. Moreover, by using plant × day-of-week × hour fixed effects in the analysis, I compare operational changes within the same power plant at specific hours and days of the week over time. This approach helps control for long-term differences between treated and control groups.

Concerns about both long-term trends and the Chinese economic slowdown stem from worries about generalized trends unrelated to the CPP. To address these concerns, I conduct placebo tests by randomly assigning treated status to observations in the sample and plotting the interaction coefficient. This process is repeated 10,000 times, with placebo coefficients binned into 1,000 buckets. Figure B.15 presents the results, with the bottom and top 2.5% percentiles in yellow. The realized coefficient is marked by a vertical red line. Notably, none of the placebo values approaches the estimated coefficient (t-statistic: -4,962.883), further supporting the reasonableness of the parallel trends assumption.

1.6 Damages and Potential Policy Responses

In this section, I provide a rough estimate of damages, benefits to firms, and sketch out possible policy responses to be covered in more detail in future work.

Calculating Damages

To provide further intuition regarding the results' magnitude, I provide some rough calculations. Starting with the PM2.5 results, I follow the methodology outlined in Muller, Mendelsohn, and Nordhaus 2011, calculating marginal damages for five air pollutants (NH_3, NO_x, PM2.5, SO_2, and VOCs) at each coal power plant in the contiguous U.S. using the AP4 model. I then merge these calculations with annual emissions data from the EPA's

National Emissions Inventory to estimate yearly emissions in metric tons for each coal plant. By applying the estimated effect from column (5) of Table C.4, I calculate the additional emissions attributed to the Green Paradox. Lastly, I multiply these excess emissions by the marginal damage estimates for each pollutant to quantify the total social cost.

Starting with PM2.5, I estimate that the Green Paradox resulted in approximately $361 million in annual damages, totaling $609 million from the Clean Power Plan's announcement to the issuance of the stay.[29] These damages include the effects of PM2.5 on health, agricultural yields, visibility, increased damage to fixed assets, and reduced recreation. If we assume a fixed relationship between PM2.5 and the production of the other four pollutants (NH_3, NO_x, SO_2, and VOCs), total damages rise to over $9 billion annually.

Damages from greenhouse gases are harder to quantify, so I present a figure that accounts for different carbon pricing and discount rate assumptions. Figure B.16 displays the net present social cost of carbon for each day over a decade. The social cost of carbon represents the discounted stream of total damages from one ton of carbon emissions. However, these damages must be discounted again because the Green Paradox suggests that the emissions are not avoided, but rather shifted forward in time.

The yellow lines represent the EPA's current social cost of carbon ($190), while the dark green lines use the social cost at the time of the Clean Power Plan's issuance ($48) (U.S. Environmental Protection Agency 2015b). Solid lines are discounted at the risk-free rate from December 2015 (0.375%), and the dashed lines use the higher risk-free rate from December 2023 (5.375%). The vertical dotted line indicates the number of days from the EPA's initial announcement to regulate power plant emissions until the stay was issued. Estimated damages range from over $13 billion in the high-cost, low-discount rate scenario to $3.5 billion in the low-cost, high-discount rate scenario. These estimates imply total damages of between $6 and $21 billion over the life of the Clean Power Plan.

---

[29] This period spans 615 days.

Policy Responses

Policy options to address the Green Paradox are particularly challenging. The Clean Air Act grants the EPA authority to regulate new, modified, and reconstructed sources, while for existing sources under Section 111(d), the EPA sets standards that states must address through their own policy plans. This process takes many years, during which the Green Paradox continues causing social harm. Any effective policy solution will need to be implemented either before the new rule is proposed or work more swiftly than regulation under Section 111(d) of the Clean Air Act.

One area of low-hanging fruit is incorporating the Green Paradox into cost-benefit analyses, enabling the EPA to account for short-term emission increases by setting stricter future targets and reducing the length of transition periods. The Regulatory Impact Analysis for the Clean Power Plan and the 2023 power plant standards did not consider the Green Paradox's impact on short-term emissions and its associated social costs (U.S. Environmental Protection Agency 2015b; U.S. Environmental Protection Agency 2023). However, the recent Phase 3 Heavy Duty Trucks rule considered a related phenomenon known as "pre-buying," where sales of current heavy-duty truck models increase ahead of regulation (U.S. Environmental Protection Agency 2024a). This precedent suggests that the EPA could incorporate the Green Paradox into their analyses if sufficient evidence of the Green Paradox is available, at which point two adjustments to rules should be considered.

First, the estimated compliance costs for the Clean Power Plan were approximately $2 billion annually in the initial years, eventually rising to around $4 billion (U.S. Environmental Protection Agency 2015b). Notably, these costs represent roughly 10–20% of the estimated damages associated with the Green Paradox, suggesting that the expenses tied to shorter grandfathering periods may be cost-benefit justified. This is especially relevant for PM2.5-related damages, which are non-linear and even short-term exposures lead to adverse health effects.

Second, the EPA could adopt the approach advocated by Revesz and Kong 2011, setting standards for new and existing sources jointly. This strategy would be particularly effective in addressing CO_2 emissions. Because CO_2 accumulates in the atmosphere where it remains for extended periods, today's emissions add to the existing stock, intensifying their harmful impact. To counteract the Green Paradox, the EPA could estimate the amount of excess emissions produced during the transition period and adjust the final emissions targets downward for both types of sources. By factoring the Green Paradox into both the implementation period and the final emissions targets, the EPA could potentially maintain a stable total stock of CO_2 emissions despite the Green Paradox.

More boldly, the EPA could consider using its expansive authority under Section 303 of the Clean Air Act to enjoin coal power plant owners from engaging in the most egregious forms of the Green Paradox. Section 303 lays out emergency powers for the EPA to respond to and prevent air pollution if it poses "an imminent and substantial endangerment to public health or welfare, or the environment."[30] The EPA can issue a temporary administrative order or bring a civil action when it has reasonable cause for concern that damage might happen and a reasonable basis that a party is contributing to the danger.[31] Logistically, Section 303 is beneficial because the EPA can seek to enjoin all actors under a single order, meaning that the EPA could seek civil action against all facilities affected by the Green Paradox in a single proceeding (Hardy et al. 1974).

Indeed, tackling the Green Paradox is exactly the type of problem Congress sought to address when it granted the EPA this authority under the Clean Air Act and its predecessors. For example, the House Report for the Air Quality Act of 1967 specified that the "authority is necessary...due to the necessary passage of time which will occur prior to establishment of enforceable standards."[32] Further, when adopting the 1970 version of Section 303, the

---

[30] 42 U.S.C. § 7603.

[31] Memorandum on transmittal of Guidance on Section 303 of the Clean Air Act from Eric V. Schaeffer (Apr. 1, 1999), `https://www.epa.gov/sites/default/files/2021-05/documents/transmittalofguidanceonsection303ofcaa040199.pdf`.

[32] H.R. Rep. No. 90-728 (1967), reprinted in 1967 U.S.C.C.A.N. 1938, 1954-55.

Senate Report stated that "an emergency situation exists whenever there is a perceptible increase in the mortality rate,"[33] which is accounted for in the damage calculations above. Congress would go on to further expand the EPA's Section 303 powers in 1977 and 1990. In sum, the EPA wields broad discretionary powers under Section 303 to prevent air pollution before the implementation of regulatory standards, and Congress has been relatively clear that it wishes for the EPA to possess these powers.

More concretely, part of the challenge in regulating against the Green Paradox is that power demand and supply can have unpredictable swings. The stochastic nature of these changes makes it hard to develop tight ex ante rules specifying how generators can operate without also reducing their incentives to meet demand in case of an emergency. Using Section 303, the EPA could seek to enjoin coal power plants from engaging in "unreasonable" changes in utilization from their pre-regulatory change baseline and what would be expected given fluctuating supply and demand. Such an ex post rule is likely to mitigate the most egregious actors, but still provide adequate incentives to intervene to ensure grid reliability. Additionally, implementation would be relatively easy for the EPA, given that it can seek to impose the rule in a given proceeding and it entails little marginal legal risk, relative to codifying a new regulation, given Congress' explicit grants of power.

Implementing both of these policy recommendations together should mitigate the role of the Green Paradox. Section 303 can tamp down on the worst actions, while setting tighter initial limits can adjust policy for the smaller temporal shifts of the Paradox. Importantly, these tools require no additional authority or funding, and could be implemented within the timeline of state plans under Section 111(d) of the Clean Air Act, allowing the EPA to fix the problem independently when making a rule within its granted authority.

---

[33] S. Rep. No. 91-1196 (1970), at 35-36.

1.7 Concluding Remarks

When energy market participants anticipate future policy interventions, economic theory suggests they adjust their activities to maximize the net present value of future profits. Specifically, if regulation is expected to raise the cost of carbon emissions, carbon-intensive fuel sources may become cheaper as producers accelerate extraction to maximize present value. This front-loading of extraction reduces the value of future reserves and leads to increased emissions in the short term—potentially undermining the policy's objectives. This phenomenon is known as the Green Paradox. In this paper, I leverage a plausibly exogenous shock from a change to legal procedure, the Supreme Court's stay of the Clean Power Plan, to empirically investigate the Green Paradox, enabling me to make several contributions to the literature.

First, this paper identifies a novel mechanism driving the Green Paradox: the depreciation channel. In this channel, owners of long-lived assets increase their consumption of fixed assets once regulation limits an asset's lifespan. I demonstrate this effect using out-of-merit generation, which captures inefficient generation after accounting for the price effects of the Green Paradox. My findings indicate that inefficient generation decreased by over 7.5% following the stay of the Clean Power Plan. I also show that this is driven by power plant owners increasing consumption of their plant's stock of economic depreciation by increased ramping.

Second, building on the finding above that focusing only on the price channel understates the Green Paradox's impact, I provide the first direct measurements of the paradox's emissions effect. Using a new method to identify spatial spillovers, I find that PM2.5 pollution increased by approximately 6.5% in zip codes with coal-fired plants before the Clean Power Plan was stayed. Additionally, prior to the stay, coal-fired power plants increased emissions by 7.7%. The cost of the increased pollution is estimated at more than $22 billion when accounting for other pollutants emitted alongside PM2.5.

Third, this paper proposes a novel solution to the Green Paradox: including the effect of the Green Paradox in the EPA's Regulatory Impact Analysis, paired with the use of Section 303 of the Clean Air Act to block unreasonable increases in utilization. In doing so, policymakers should jointly shorten the transition period and lower final emissions targets to account for emissions brought forward in time while simultaneously restraining the worst actors.

Overall, this paper provides valuable guidance for regulators on the optimal rollout of regulatory policies, shedding light on the unintended consequences that such policies may generate. When calculating phase-in periods, regulators must balance firms' adjustment costs against the risk of counterproductive behavior. Specifically, in the context of climate policy, policymakers should be aware of the legal risks that proposed regulations face. If not, as demonstrated by the Clean Power Plan, they risk incentivizing market participants to accelerate greenhouse gas emissions before a rule is overturned, which paradoxically hastens climate change rather than delaying it.

REFERENCES

Air Quality Life Index, “How Does Particulate Air Pollution Impact Health?,” (2023).

Bakke, Tor-Erik., Hamed. Mahmudi, Aazam. Virani, and Song. Zhang, “Politicization of the Supreme Court and Firm Value: Evidence from Ruth Bader Ginsburg’s Death,” *SSRN*, (2023).

Borrero, Miguel., Gautam. Gowrisankaran, and Ashley. Langer, “Ramping costs and coal generator exit,” (2024).

Cicala, Steve., “When Does Regulation Distort Costs? Lessons from Fuel Procurement in US Electricity Generation,” *American Economic Review*, 105 (2015), 411–444.

Cicala, Steve., “Imperfect Markets Versus Imperfect Regulation in US Electricity Generation,” *American Economic Review*, 112 (2022), 409–441.

Cusworth, Daniel H., Riley M. Duren, Andrew K. Thorpe, Michael L. Eastwood, Robert O. Green, Philip E. Dennison, Christian. Frankenberg, Joseph W. Heckler, Gregory P. Asner, and Charles E. Miller, “Quantifying Global Power Plant Carbon Dioxide Emissions with Imaging Spectroscopy,” *AGU Advances*, 2 (2021).

Cusworth, Daniel H., Andrew K. Thorpe, Charles E. Miller, A. Ayasse, R. Jiorle, Riley M. Duren, Ray. Nassar, Jon-Paul. Mastrogiacomo, and Robert R. Nelson, “Two Years of Satellite-based Carbon Dioxide Emission Quantification at The World’s Largest Coal-fired Power Plants,” *EGUsphere*, (2023), 1–35.

Di Maria, Corrado., Ian. Lange, and Edwin. Van der Werf, “Should We Be Worried about The Green Paradox? Announcement Effects of the Acid Rain Program,” *European Economic Review*, 69 (2014), 143–162.

Ebenstein, Avraham., Maoyong. Fan, Michael. Greenstone, Guojun. He, and Maigeng. Zhou, “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 (2017), 10384–10389.

Fabrizio, Kira R., Nancy L. Rose, and Catherine D. Wolfram, “Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on US Electric Generation Efficiency,” *American Economic Review*, 97 (2007), 1250–1277.

Gerlagh, Reyer., “Too Much Oil,” *CESifo Economic Studies*, 57 (2011), 79–102.

Glicksman, Robert L., and Johanna. Adashek, “Agency Authority to Address Chemicals of Emerging Concern: EPA’s Strategic Use of Emergency Powers to Address PFAS Air Pollution,” *Harv. Env’t L. Rev.*, 48 (2024), 369.

Grafton, R. Quentin., Tom. Kompas, Ngo. Van Long, and Hang. To, “US Biofuels Subsidies and CO2 Emissions: An Empirical Test for a Weak and a Strong Green Paradox,” *Energy Policy*, 68 (2014), 550–555.

Hall, Robert E., “The Dynamic Effects of Fiscal Policy in an Economy with Foresight,” *The Review of Economic Studies*, 38 (1971), 229–244.

Hammond, Emily., and Richard J. Pierce Jr, “The Clean Power Paln: Testing the Limits of Asdministrative Law and the Electric Grid,” *Geo. Wash. J. Energy & Envtl. L.*, 7 (2016), 1.

Hardy, George E. Jr., P. Pate, C. B. Robison, and W. T. Willis, “First use of the Federal Clean Air Act’s emergency authority: A local analysis,” *American Journal of Public Health*, 64 (1974), 72–76.

Henneman, Lucas., Christine. Choirat, Irene. Dedoussi, Francesca. Dominici, Jessica. Roberts, and Corwin. Zigler, “Mortality risk from United States coal electricity generation,” *Science*, 382 (2023), 941–946.

Hernandez-Cortes, Danae., and Kyle C. Meng, “Do Environmental Markets Cause Environmental Injustice? Evidence from California’s Carbon Market,” *Journal of Public Economics*, 217 (2023), 104786.

Hoel, Michael., “Carbon Taxes and The Green Paradox,” *Climate Change and Common Sense: Essays in Honour of Tom Schelling*, (2012), 203–24.

Hotelling, Harold., “The Economics of Exhaustible Resources,” *Journal of Political Economy*, 39 (1931), 137–175.

Ikeuchi, Hiroaki., Michio. Murakami, and Satoshi. Watanabe, “Scavenging of PM2. 5 by Precipitation and the Effects of Precipitation Pattern Changes on Health Risks Related to PM2. 5 in Tokyo, Japan,” *Water Science and Technology*, 72 (2015), 1319–1326.

Irwin, Neil., “The Most Important Least-Noticed Economic Event of the Decade,” (2018).

Judd, Kenneth L., “Short-run Analysis of Fiscal Policy in a Simple Perfect Foresight Model,” *Journal of Political Economy*, 93 (1985), 298–319.

Kaplow, Louis., “An economic analysis of legal transitions,” *Harvard Law Review*, (1986), 509–617.

Klass, Alexandra B., and Rebecca. Wilton, “Local Power,” *Vand. L. Rev.*, 75 (2022), 93.

Kumar, Nikil., Philip. Besuner, Steven. Lefton, Dwight. Agan, and Douglas. Hilleman, “Power plant cycling costs,” (2012).

Lemoine, Derek., “Green Expectations: Current Effects of Anticipated Carbon Pricing,” *Review of Economics and Statistics*, 99 (2017), 499–513.

Levmore, Saul., “Changes, anticipations, and reparations,” *Colum. L. Rev.*, 99 (1999), 1657.

Liptak, Adam., “Justice Antonin Scalia, Who Led a Conservative Renaissance on the Supreme Court, Is Dead at 79,” (2016).

Liptak, Adam., and Coral. Davenport, “Supreme Court Blocks Obama E.P.A. Coal-Emissions Regulations,” (2016).

Macey, Joshua., “Zombie Energy Laws,” *Vand. L. Rev.*, 73 (2020), 1077.

Macey, Joshua., and Jackson. Salovaara, “Bankruptcy as Bailout: Coal Company Insolvency and the Erosion of Federal Law,” *Stan. L. Rev.*, 71 (2019), 879.

Muller, Nicholas Z., Robert. Mendelsohn, and William. Nordhaus, “Environmental accounting for pollution in the United States economy,” *American Economic Review*, 101 (2011), 1649–1675.

Nash, Jonathan Remy., and Richard L. Revesz, “Grandfathering and environmental regulation: the law and economics of new source review,” *Nw. UL Rev.*, 101 (2007), 1677.

Norman, Maya A., and Wolfram. Schlenker, “Empirical Tests of the Green Paradox for Climate Legislation,” (2024).

Okullo, Samuel J., Frédéric. Reynès, and Marjan W. Hofkes, “(Bio-) Fuel Mandating and the Green Paradox,” *Energy Economics*, 95 (2021), 105014.

Parton, Lee C., and Steven J. Dundas, “Fall in the Sea, Eventually? A Green Paradox in Climate Adaptation for Coastal Housing Markets,” *Journal of Environmental Economics and Management*, 104 (2020), 102381.

Plumer, Brad., “How Obama’s Clean Power Plan actually works — a step-by-step guide,” (2015).

Preonas, Louis., “Market power in coal shipping and implications for US climate policy,” *Review of Economic Studies*, 91 (2024), 2508–2537.

“Recent Regulation: The Clean Power Plan,” *Harvard Law Review*, 129 (2016).

Reguant, Mar., “Complementary bidding mechanisms and startup costs in electricity markets,” *The Review of Economic Studies*, 81 (2014), 1708–1742.

Revesz, Richard L., and Allison L Wesfahl. Kong, “Regulatory Change and Optimal Transitional Relief,” *Nw. UL Rev.*, 105 (2011), 1581.

Shallenberger, Krysti., “Oregon passes 50% renewables standard into law,” (2016).

Shallenberger, Krysti., and Gavin. Bade, “Ohio PUC approves AEP, FirstEnergy plant subsidy requests,” (2016).

Shavell, Steven., “On optimal legal change, past behavior, and grandfathering,” *The Journal of Legal Studies*, 37 (2008), 37–85.

Sinn, Hans-Werner., “Public Policies Against Global Warming: A Supply Side Approach,” *International Tax and Public Finance*, 15 (2008), 360–394.

Stavins, Robert N., “Vintage-differentiated environmental regulation,” *Stan. Envtl. LJ*, 25 (2006), 29.

Stein, AF., Roland R. Draxler, Glenn D. Rolph, Barbara JB. Stunder, Mark D. Cohen, and Fong. Ngan, “NOAA’s HYSPLIT Atmospheric Transport and Dispersion Modeling System,” *Bulletin of the American Meteorological Society*, 96 (2015), 2059–2077.

U.S. Department of Energy, “Balancing Authority Backgrounder,” (2023).

U.S. Environmental Protection Agency, “FACT SHEET: Overview of the Clean Power Plan,” (2015).

U.S. Environmental Protection Agency, “Regulatory Impact Analysis for the Clean Power Plan Final Rule,” (2015).

U.S. Environmental Protection Agency, “Regulatory Impact Analysis for the Proposed New Source Performance Standards for Greenhouse Gas Emissions from New, Modified, and Reconstructed Fossil Fuel-Fired Electric Generating Units; Emission Guidelines for Greenhouse Gas Emissions from Existing Fossil Fuel-Fired Electric Generating Units; and Repeal of the Affordable Clean Energy Rule,” (2023).

U.S. Environmental Protection Agency, “Light-Duty Vehicle Greenhouse Gas Emissions Standards: Regulatory Impact Analysis,” (2024).

U.S. Environmental Protection Agency, “New Source Performance Standards for Greenhouse Gas Emissions from New, Modified, and Reconstructed Fossil Fuel-Fired Electric Generating Units; Emission Guidelines for Greenhouse Gas Emissions from Existing Fossil Fuel-Fired Electric Generating Units; and Repeal of the Affordable Clean Energy Rule,” 89 (2024), 39798–40064.

Velazco, V. A., M. Buchwitz, H. Bovensmann, M. Reuter, O. Schneising, J. Heymann, T. Krings, K. Gerilowski, and J. P. Burrows, “Towards Space Based Verification of CO_2 Emissions from Strong Localized Sources: Fossil Fuel Power Plant Emissions as Seen by a CarbonSat Constellation,” *Atmospheric Measurement Techniques*, 4 (2011), 2809–2822.

Walton, Robert., "N.J. Senate Committee Passes Bill for 80% Renewable Energy Standard," (2016).

Wei, Y., X. Xing, and A. Shtein, "Daily and Annual PM2. 5, O3, and NO2 Concentrations at ZIP Codes for the Contiguous US, 2000-2016, v1.0," *NASA Socioeconomic Data and Applications Center*, (2022).

Welton, Shelley., "Public energy," *NYUL Rev.*, 92 (2017), 267.

Wolf, Richard., "Supreme Court Blocks President Obama's Climate Change Plan," (2016).

Yerramilli, Anjaneyulu., Venkata Bhaskar Rao. Dodla, Venkata Srinivas. Challa, LaToya. Myles, William R. Pendergrass, Christoph A. Vogel, Hari Prasad. Dasari, Francis. Tuluri, Julius M. Baham, and Robert L. Hughes, "An Integrated WRF/HYSPLIT Modeling Approach for the Assessment of PM 2.5 Source Regions over the Mississippi Gulf Coast Region," *Air Quality, Atmosphere & Health*, 5 (2012), 401–412.

APPENDIX A

SUPPLEMENTARY RESULTS AND INSTITUTIONAL DETAILS

A.1 Additional Litigation Details

Litigation against the CPP became increasingly unusual following the stay. Before oral argument, on October 4th, 2017, Reuters reported that the Trump administration sought to repeal the CPP and replace it with the Affordable Clean Energy (ACE) Rule.[1] The EPA was sued to block implementation of the ACE, and the D.C. Circuit ruled on January 19, 2021, that the ACE exceeded the EPA's authority.[2] Since the Trump administration had rescinded the CPP to implement the ACE, the D.C. Circuit's decision effectively reinstated the Clean Power Plan. The ACE was similar to the CPP but significantly less restrictive.

Following the D.C. Circuit's ruling, the EPA faced another lawsuit to prevent enforcement of the Clean Power Plan. This time, the Supreme Court heard the case. On June 30, 2022, the Supreme Court invalidated the CPP, ruling that the EPA did not have the authority to promulgate the rule. This marked a novel precedent as the Court had never before struck down a regulation under the "Major Questions" doctrine.[3] This final blow to the Clean Power Plan is what one would typically use in an event study. However, because the stay was so unusual market participants already expected the rule to be overturned; thus, no new information was provided when it was finally resolved.

---

[1] V. Volcovici, Trump EPA to Propose Repealing Obama's Climate Regulation: Document, *Reuters* (Oct. 4, 2017), `https://www.reuters.com/article/idUSKBN1C92CI/`.

[2] S. Mufson, Federal Court Scraps Trump Administration's Power Plant Rule, *Washington Post* (Jan. 19, 2021), `https://www.washingtonpost.com/climate-environment/2021/01/19/federal-circuit-court-scraps-trump-administration-power-plant-rules/`.

[3] Under the major questions doctrine, the Supreme Court has rejected agency claims of regulatory authority when (1) the underlying claim of authority concerns an issue of "vast 'economic and political significance,'" and (2) Congress has not clearly empowered the agency with authority over the issue. Util. Air Regul. Grp. v. EPA, 573 U.S. 302, 324 (2014).

A.2 Additional Tests: Purchasing Behavior and Coal Extraction

In the main text, I show that treated units reduced their operating time, CO_2 production, and PM2.5 emissions. As a robustness check, I use a separate data set to study whether power plant operators had a corresponding change in fuel purchasing behavior. If operators altered their behavior in response to the regulation in the direction of the Green Paradox, they would likely need to supplement their existing contractual fuel purchases. Firms can choose to do so in the spot market or via contracts (Cicala 2015). Firms may prefer spot purchases as they save on contracting costs and reduce lock-in, while futures markets can manage price risk. However, given that the Green Paradox decreases prices, firms may wish to enter into supply contracts to lock in the lower prices and transfer risk to their suppliers. Which force wins out is ambiguous; therefore, while I cannot sign a differential effect between spot and contract purchases, the overall effect should be a reduction in fuel purchases following the issuance of the stay as firms reduce their utilization of treated fuels.

I test this using data from Schedule 2 of EIA Form 923. Form 923 provides monthly data on generation and fuel consumption at the power plant level, and Schedule 2 collects plant-level fuel receipts and cost data. I use the energy content of each purchase to standardize the units in mmBTUs. It also details whether the fuel purchases are spot, under an existing contract, or from a new contract.

I explore changes in purchasing behavior by disentangling whether they come from increases in spot or contractual purchases. The results are in Table C.8. Following the stay, power plant operators reduced their fuel purchases across the board, with more of the adjustment conducted in the spot market, consistent with firms boosting their consumption as predicted by the Green Paradox.

Next, I look to the coal market for evidence of the Green Paradox in coal production. I use weekly coal production data from the EIA for the entire United States. To examine the evolution of the production path, I estimate an ARIMA(3, 1, 0) model. The results are in Figure B.19. After the stay in *West Virginia*, coal production continued to fall over the following weeks, consistent with the Green Paradox. However, this data is low frequency and at a high level of aggregation, so it should interpreted as further suggestive evidence consistent with the prior results rather than a causal estimate.

A.3 Additional Tests: Equity and Coal Prices Paid

To test the price channel, I follow prior work (e.g., Lemoine 2017) and look at abnormal returns for electricity generators and firms involved in coal mining, defined by SIC codes 4911 and 1220/1221, respectively. If the price channel is in effect, we should see positive abnormal returns for coal mining firms following the stay.

Isolating the effect on coal miner returns during this period is challenging. Throughout 2016, coal producers were in significant financial distress due to the fuel's secular decline and debt loads that had not yet adjusted to that new reality. The average yield to maturity for the sector was well above 20%, and many firms would file for bankruptcy around 2017 (Macey and Salovaara 2019). In light of this issue, I leverage the fact that the stay was announced after the market closed on February 9, 2016, allowing me to calculate the close-to-open abnormal returns. Close-to-open returns are the preferred measure as they do not incorporate any information that arrived during the trading day, which would be unrelated to the Clean Power Plan.

Further, given the secular trend for coal producers, I estimate a Fama-French Three Factor plus Momentum model using an estimate window of -200 to -15 days before the stay. I calculate abnormal returns at the value-weighted industry level since the stay should affect the sector as a whole. Figure B.20 plots the abnormal daily returns for the [-2,2] event window. Confidence intervals are Bonferroni corrected. Table C.9 reports the abnormal returns and t-statistics for the event day.

The returns paint an interesting picture. Measured on a close-to-open basis, coal producers have significant positive abnormal returns, consistent with the Green Paradox channel. Given the importance of fuel prices to their variable costs, electric utilities have significant negative abnormal returns on the event date. However, the close-to-close returns are significant and negative for coal miners and insignificantly different from zero for utilities. As mentioned above, the close-to-close tests are more likely to be impacted by intra-day news given the distressed state of the industry, making them a less well-identified measure. To examine this possibility, I check Factiva for additional coal-related news. I found no mention of coal that day other than to note that those stocks opened higher.[4]

Given the mixed results above, I also use information from prices paid on EIA Form 923. This data provides information on the quantity and prices paid by power plants at a monthly frequency. I restrict my analysis to spot purchases because contracts are more likely to have sticky prices, whereas spot prices should respond in near real-time. Figure B.21 plots the evolution of spot prices, calculated as the monthly weighted average price paid per mmBTU by all reporting power plants, in panel A. Panel B plots the distribution of the monthly log returns of spot coal prices from Form 923 from 2008 to 2023. The dotted line is the return for the first reporting date after the stay. Consistent with theory and the close-to-open returns tests, the average price paid per mmBTU of coal increased following the stay.

To formalize the analysis slightly, Table C.10 provides estimates from regressing the log returns on event dummy, year, and quarter fixed effects using data from 2015 and 2016. Standard errors are bootstrapped with 10,000 repetitions. Consistent with the tests above, the log returns are positive and marginally significant in two of the three specifications.

Thus, on balance, the evidence appears consistent with the Green Paradox being the channel through which this effect operates.

---

[4] https://www.reuters.com/article/usa-court-carbon/coal-company-shares-jump-after-supreme-court-blocks-emission-plan-idUSL2N15P16Q/.

A.4 Supreme Court Vacancy Event Studies

On February 13, 2016, just four days after the stay was issued in *West Virginia v. EPA*, Justice Antonin Scalia passed away leaving a vacancy on the Supreme Court. When a vacancy occurs, the President nominates a candidate which must be approved by a majority vote of the Senate. In this case, the Presidency was held by the Democrats and the Senate by the Republicans. Any nominee then would have to be amenable to both parties. This creates some ambiguity as to whether Scalia's passing would have helped or hurt the odds of the Clean Power Plan being upheld.

On the one hand, a Democratic-nominated justice likely would have been more favorable to upholding the rule. On the other, a nominee that would garner approval from the Republicans, a necessary condition to be appointed, likely would have been favorable to the removal of the Clean Power Plan. Thus, to answer this question I repeat the event studies from the issuance of the stay, with similar, albeit much more muted, results which can be seen in Table C.7 and Figure B.18. The close-to-open returns are again positive for coal producers, but the close-to-close are negative. This result could be read as the vacancy increasing the odds that the Clean Power Plan would not be enacted. However, the close-to-close returns are significant and negative as in Bakke et al. 2023. The negative returns could be due to a belief that the new nominee would uphold the Clean Power Plan, or could be the result of market participants selling coal producers upon positive news. Given that the close-to-open measure is partially insulated from this effect, it is arguably more accurate. The vacancy at the Supreme Court then appears to have made it more likely that the Clean Power Plan would be held invalid.

A.5 Government Ownership

One potential response to the Green Paradox is for the government to engage in public-private bargaining with coal-fired power plant owners. Through voluntary agreements, it may be possible to act more swiftly than by implementing modified source standards under the Clean Air Act. Private owners increase plant utilization because their profit-maximizing strategies do not account for the social costs of PM2.5 and GHG emissions. In contrast, government entities internalize some of these social costs through reduced tax revenue or increased public expenditures (Welton 2017; Klass and Wilton 2022). Alternatively, the government may simply lack the strong incentives that drive the Green Paradox in the private sector. Thus, a government owner might either partially internalize the costs associated with the Green Paradox or lack the incentive to respond aggressively, making it less vulnerable to its effects. To test this hypothesis, I re-run the extensive and intensive margin tests in column (5), incorporating a triple interaction among the treatment indicator, post-period indicator, and an indicator for government ownership (whether federal, state, or municipal).

Consistent with this intuition, the triple difference for the time operating yields a coefficient of 0.094 (t-statistic: 3.42), while the regular treated-post interaction produces -0.058 (t-statistic: -3.59). This suggests that government-owned plants are largely unaffected by the Green Paradox and are displaced by private generators.[5] In Figure B.17, I present event-time plots for the triple interaction of event time, treatment status, and government ownership. The plots show no evidence of extensive margin adjustments, and any intensive margin effect is significantly diminished.

On the intensive margin, there is no statistically significant change in the CO_2 emissions of government-owned plants (the coefficient estimate on the triple difference is -0.079, with a t-statistic of -1.22), implying that public ownership primarily influences the decision of whether to generate, rather than the intensity of operations. Overall, this suggests that

---

[5] The estimated coefficient on the treated-post interaction is 0.042, indicating that government ownership nearly offsets the Green Paradox effect seen in private generators.

government acquisition of stranded assets could potentially mitigate the Green Paradox. To enhance the impact of this approach, the government could link executive compensation for managers of government-owned plants to the frequency with which these plants operate within the merit order, based on relative efficiency that excludes input price fluctuations.

Government ownership of power plants would represent a significant shift away from the trend of increasing market forces in electricity markets. A less drastic alternative could be to expand and reform Community Choice Aggregation (CCA) (Welton 2017; Klass and Wilton 2022). CCAs enable local governments to procure electricity for residents while using the existing private utility for distribution and grid maintenance. These organizations can contractually specify energy sources, allowing customers to opt out of coal-fired generation. If CCAs control sufficient demand, they might mitigate the Green Paradox by limiting the ability of generators to meet demand by increasing coal generation.

However, CCAs are currently authorized in only nine states, and many of these states do not grant CCAs the authority needed to maximize their effectiveness.[6] Two specific features of CCA statutes are especially relevant to addressing the Green Paradox. First, since a critical mass of demand would need to exclude coal generation, statutes should enable local governments to collaborate, thereby reducing bargaining costs. Lower costs should lead to lower prices and, more importantly, shorten the time needed to amass sufficient demand to prevent generation shifting by power plant operators.

The second essential feature is the authority to enter into long-term contracts, which serves two purposes. First, as widely recognized in the literature, long-term contracts can incentivize the construction of cleaner energy sources, ensuring enough non-coal capacity to maintain grid reliability during demand spikes. Additionally, long-term contracts can act as a commitment mechanism, preventing CCAs from contracting for coal-based generation even if the Green Paradox makes coal temporarily cheaper. Thus, long-term contracts help

---

[6] John Farrell, Community Choice Energy: An Alternative to Utility Energy Monopolies Gives Communities More Control, Inst. for Local Self-Reliance (Feb. 2020), `https://ilsr.org/wp-content/uploads/2020/02/CommunityChoiceEnergyReportILSR.pdf`.

CCAs maintain their commitment to exclude coal generation despite potential Green Paradox incentives. Community Choice Aggregation is not without drawbacks. In particular, cost allocation for past investments and future investments to increase capacity have been sticking points for existing CCAs.[7]

---

[7] Trieu Mai et al., Electrification Futures Study: Scenarios of Electric Technology Adoption and Power Consumption for the United States, Nat'l Renewable Energy Lab'y, at 1 (2018), `https://www.nrel.gov/docs/fy19osti/72195.pdf`.

APPENDIX B

FIGURES

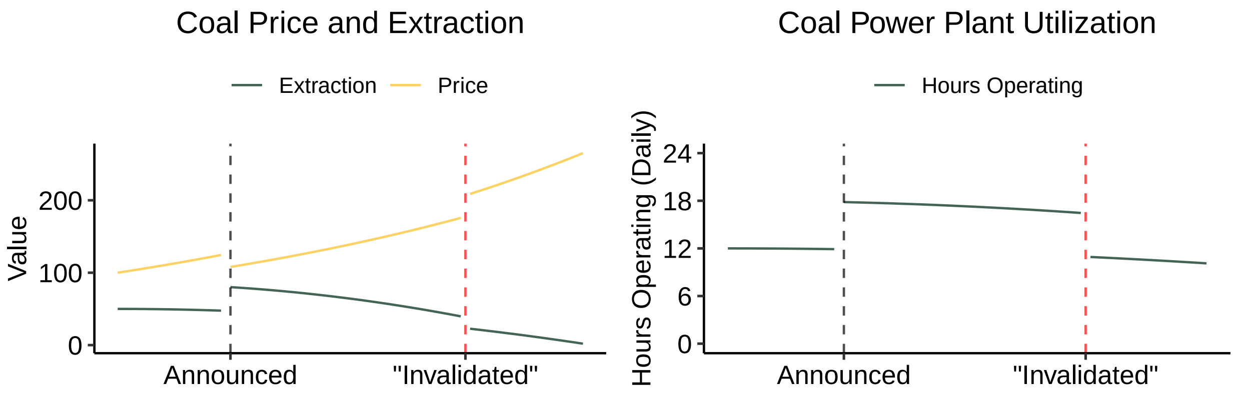

(a) Green Paradox - Coal Mining

(b) Green Paradox - Coal Generation

Figure B.1: The Green Paradox in Pictures

Panel A plots the hypothesized price and extraction paths for the Green Paradox. The black dotted line represents the increase in extraction and reduction in price corresponding to the announcement of a proposed regulation that brings forward the terminal period for a given fossil fuel (i.e., coal). Panel B plots this effect for coal power plants, which increase their operations after the policy is announced. The dotted red line corresponds to the empirical construct tested in this paper, the plausible exogenous “invalidation” of the regulation announced at the black dotted line. Prices should increase, and extraction and coal power plant utilization should fall at the time of “invalidation”.

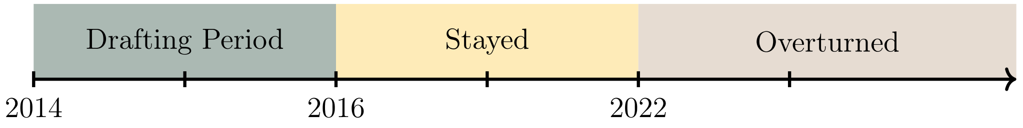

Figure B.2: Clean Power Plan Litigation Timeline

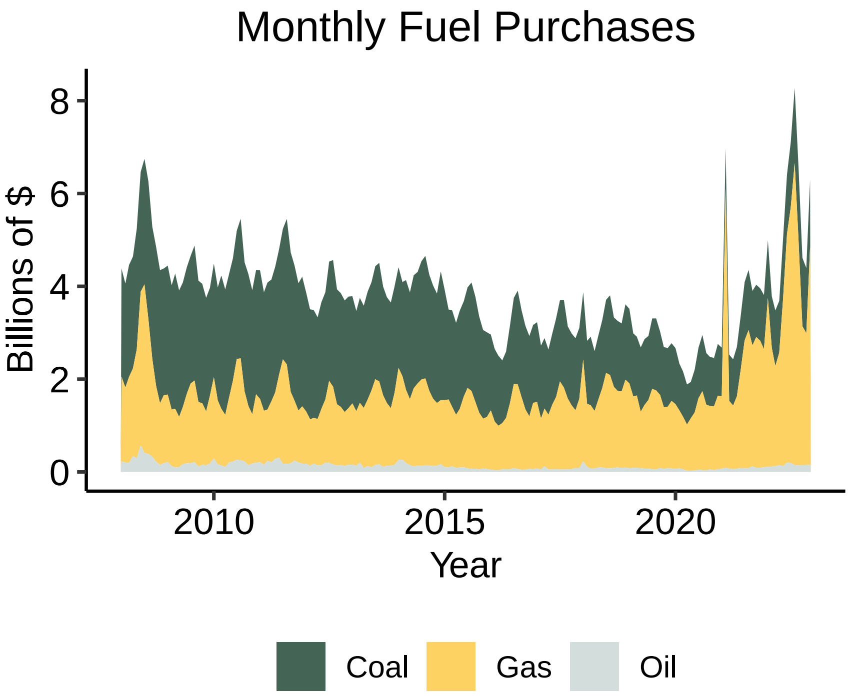

Figure B.3: Monthly Fuel Purchases

This figure plots the total value of monthly fuel purchases for coal, gas, and oil as reported on Schedule 2 of EIA Form 923.

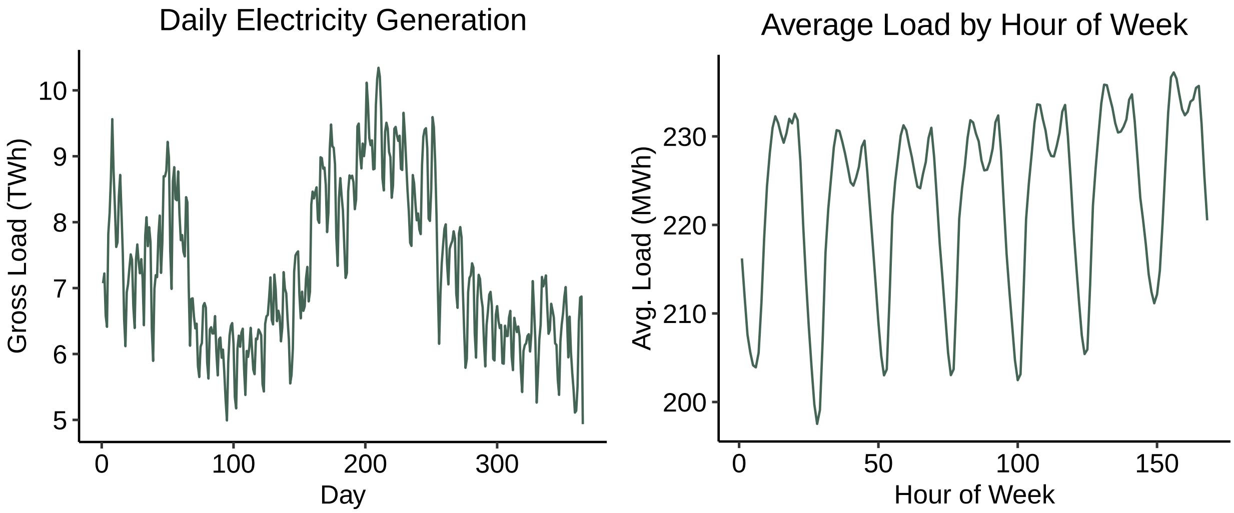

(a) Daily Electricity Generation

(b) Average Hourly Load

Figure B.4: Electricity Generation and Average Hourly Load

Panel A plots the total gross amount of electricity generated each day in 2015, highlighting seasonality in generation. Panel B plots the average gross generation for each hour of the week over the 2015-2016 period, demonstrating the intra-week fluctuations in generation.

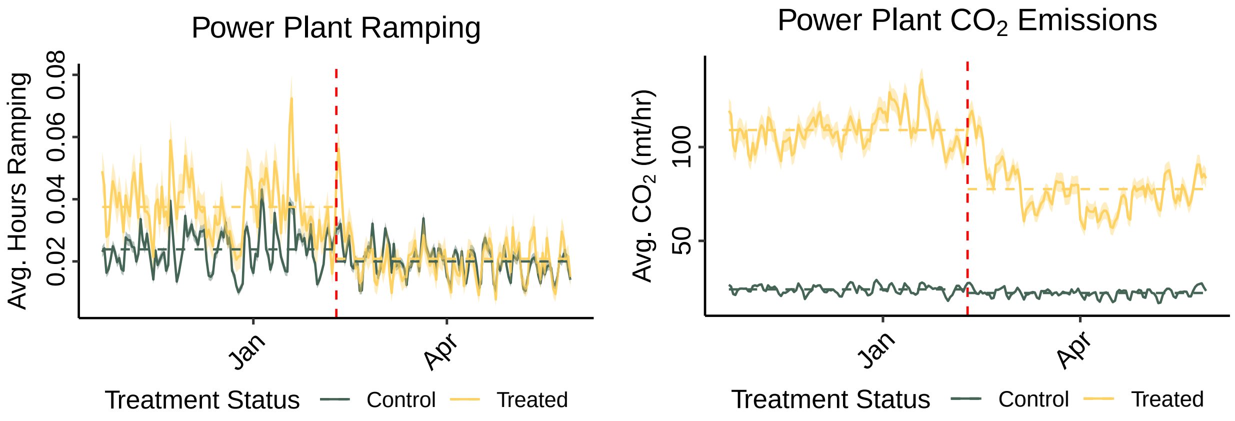

(a) Mean Hourly Ramping (b) Mean CO_2 Emissions

Figure B.5: Raw Ramping and CO_2 Emissions Data

Panel A of this figure plots the mean ramping by coal power plants, the treated group, relative to natural gas power plants, the control group, in the twelve weeks around the stay of the Clean Power Plan. Panel B plots the mean CO_2 emissions for treated and control plants over the same twelve-week window. Standard errors are plotted in each figure.

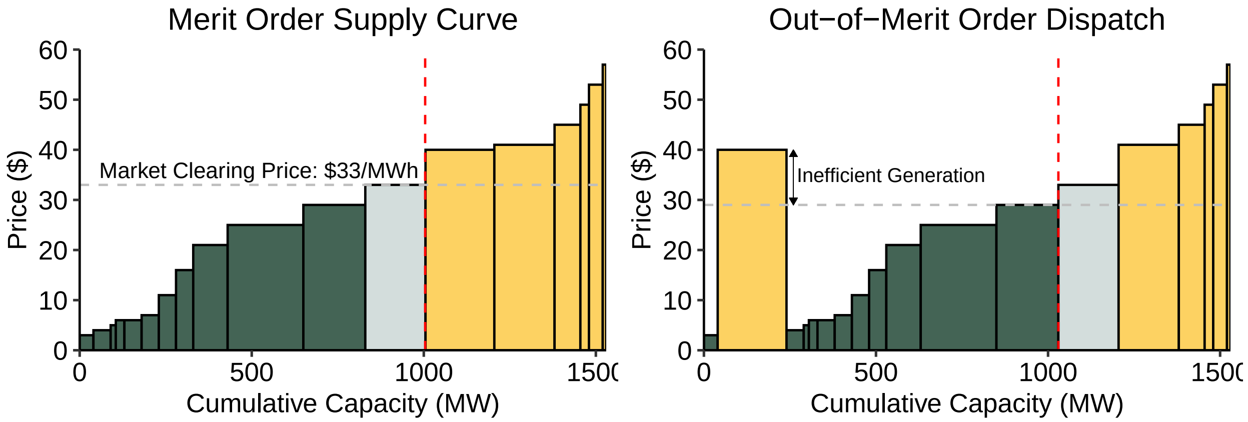

(a) Merit Order Generation (b) Out-of-Merit Order Generation

Figure B.6: Electricity Supply and Merit Order

Panel A plots the ideal merit order generation. The red line represents electricity demand in megawatts, the dotted grey line is the market clearing price, the green bars depict inframarginal generators, the grey bar represents the marginal unit, and the yellow bars show high-cost generators that do not run. Panel B provides an example of out-of-merit generation, indicated by the yellow bar on the left. This unit is used to meet demand despite having a higher marginal cost compared to the first non-utilized plant to the right of the red demand line.

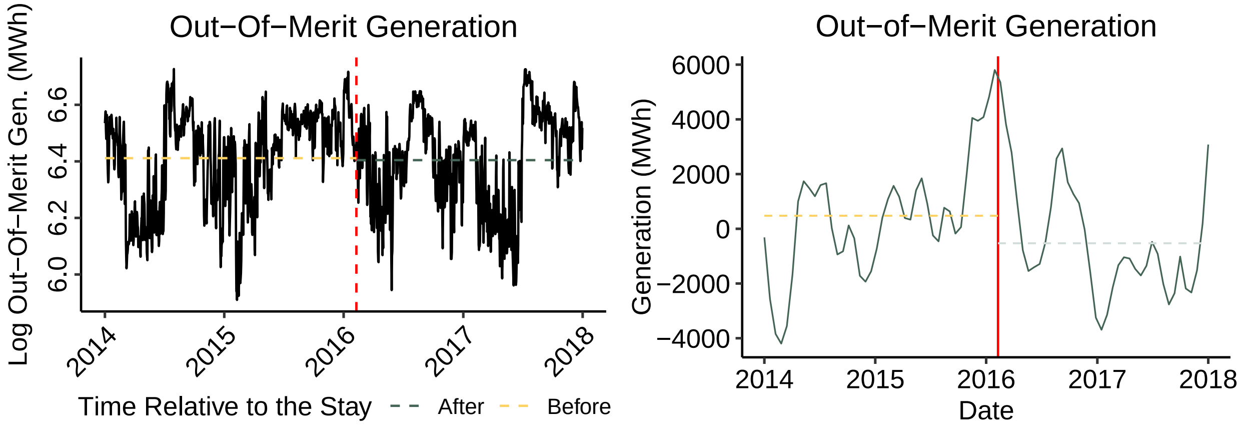

(a) Out-of-Merit Order Generation

(b) Merit Order Generation

Figure B.7: Electricity Supply and Merit Order

Panel A plots the raw time series of the log of out-of-merit generation from 2014 through 2017. Panel B plots the realized out-of-merit generation residualized for weather controls and day-of-week-hour and day-of-year-balancing authority fixed effects. The vertical red line marks the day the CPP was stayed. The dotted yellow line indicates the pre-stay mean out-of-merit generation and the dotted grey line marks the post-stay mean out-of-merit generation.

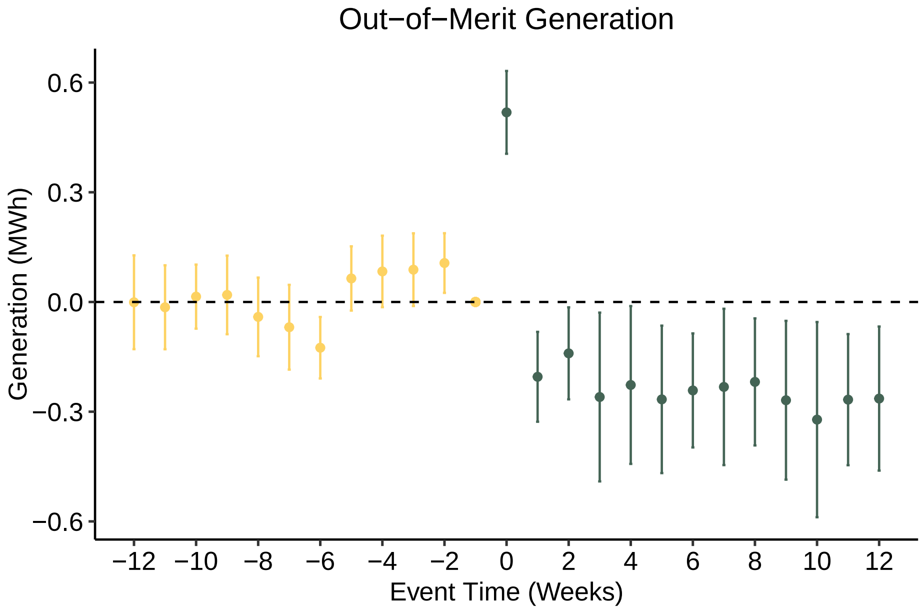

Figure B.8: Weekly Out-of-Merit Generation Event Plot

This figure takes the specification from column (4) in Table C.2 and aggregates it to the weekly level. The dependent variable is the log of out-of-merit generation over the week. Confidence intervals are at the 95% level.

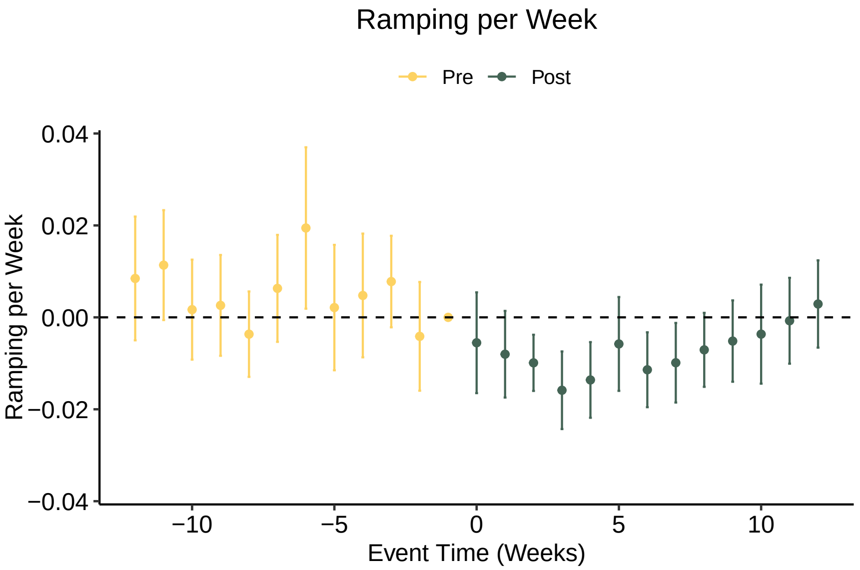

Figure B.9: Weekly Ramping Event Plot

This figure takes the specification from column (5) in Table C.3 and aggregates it to the weekly level. The dependent variable is an indicator variable for whether a plant ramps by more than 10% from its minimum or maximum generation capacity. Confidence intervals are at the 95% level.

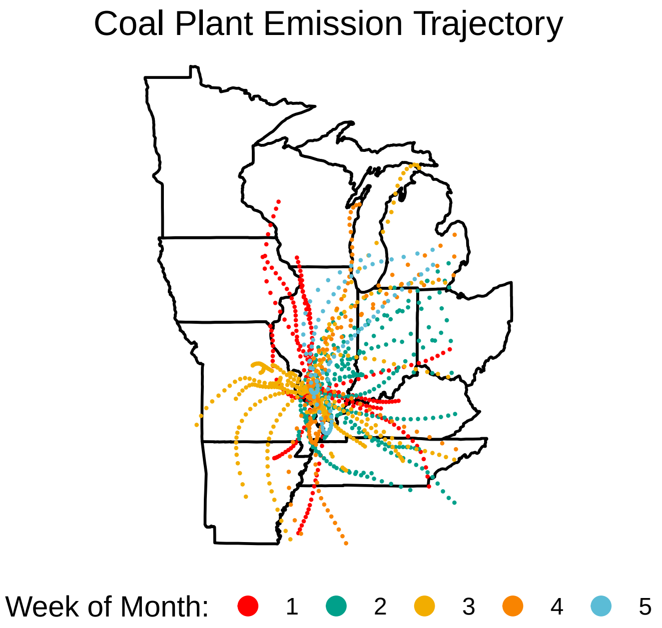

Figure B.10: Coal Plant Emission Trajectory

This figure provides an example of the HYSPLIT trajectory analysis using the trajectories from the Prairie State Energy Campus in January 2016. It plots the four daily trajectories with their corresponding hourly positions represented by each dot. The different colors correspond to the trajectories for a given week in January 2016.

Fifth Draft/Images/PM25.pdf

Figure B.11: PM2.5 Concentration by ZCTA

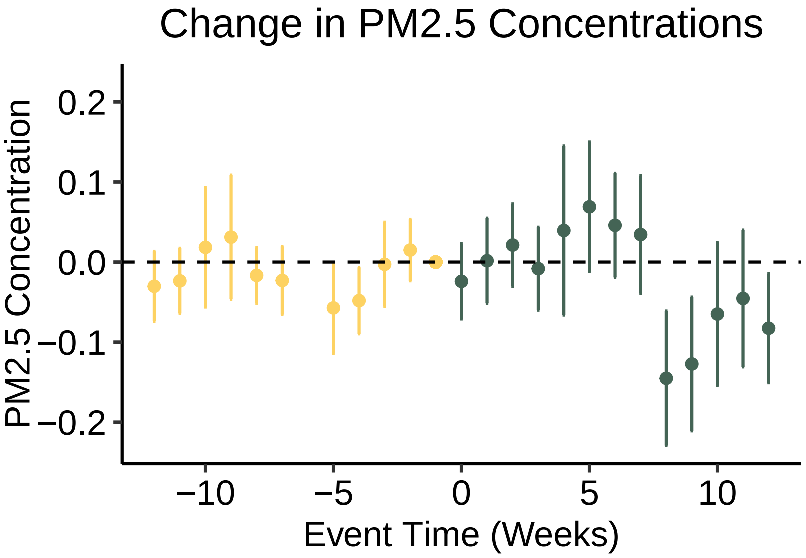

Figure B.12: Change in PM2.5 Concentrations in Zip Codes with Coal Power Plants
This figure takes the specification of column (5) in Table C.4 and aggregates it to the weekly level. The outcome variable is the log of the PM2.5 concentration. Each treatment variable is interacted with weekly dummies in event time. Confidence intervals are at the 95% level.

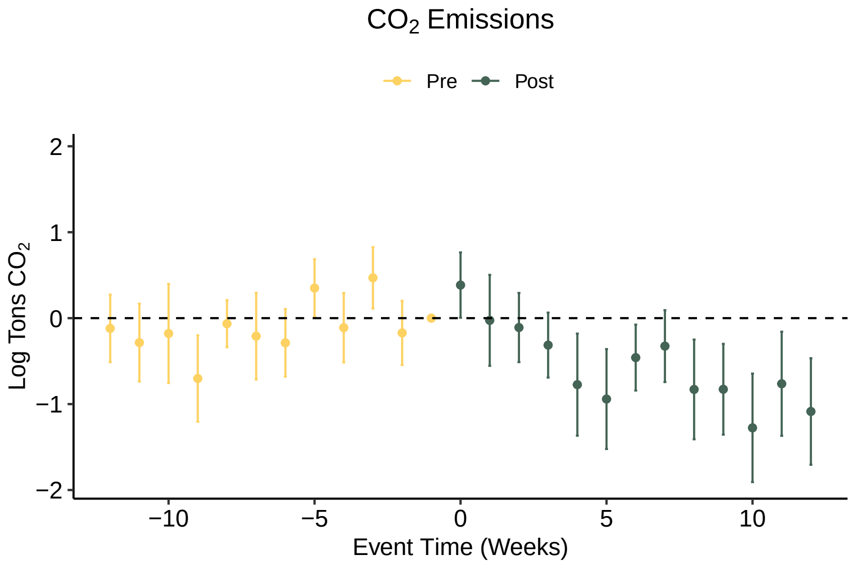

Figure B.13: Change in CO_2 Emissions from Coal Power Plants

This figure takes the specification from column (5) in Table C.5 and aggregates it to the weekly level. The dependent variable is the log tons of carbon dioxide emissions over timeweek. Each treatment variable is interacted with weekly dummies in event time. Confidence intervals are at the 95% level.

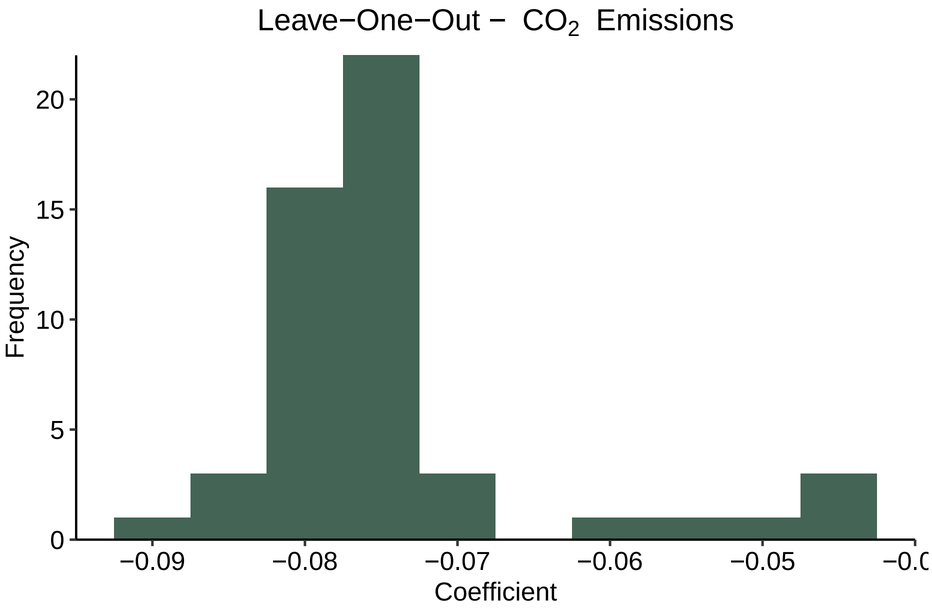

Figure B.14: Leave-One-Out Tests

Each figure plots the results of a leave-one-out test where I drop one state and then reestimate the treatment effect for column (5) in Tables C.5.

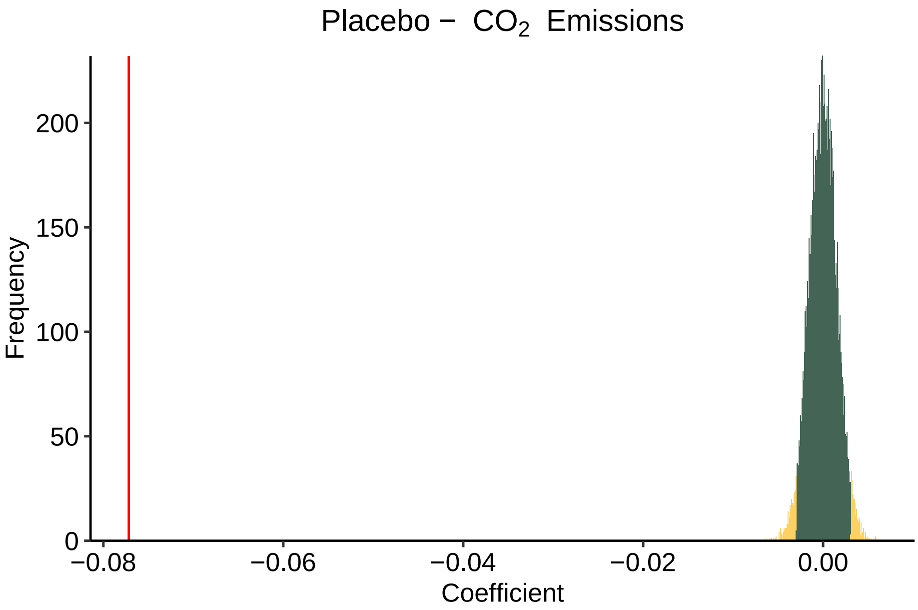

Figure B.15: CO_2 Emissions Placebo Test

Each figure plots the results of a placebo test where treatment is randomly assigned, and the model from column (5) in Tables C.5. The vertical red line corresponds to the estimated treatment effect from those models. The upper and lower 2.5% of placebo coefficients are in yellow. The t-statistic of the observed coefficient is -4,962.883.

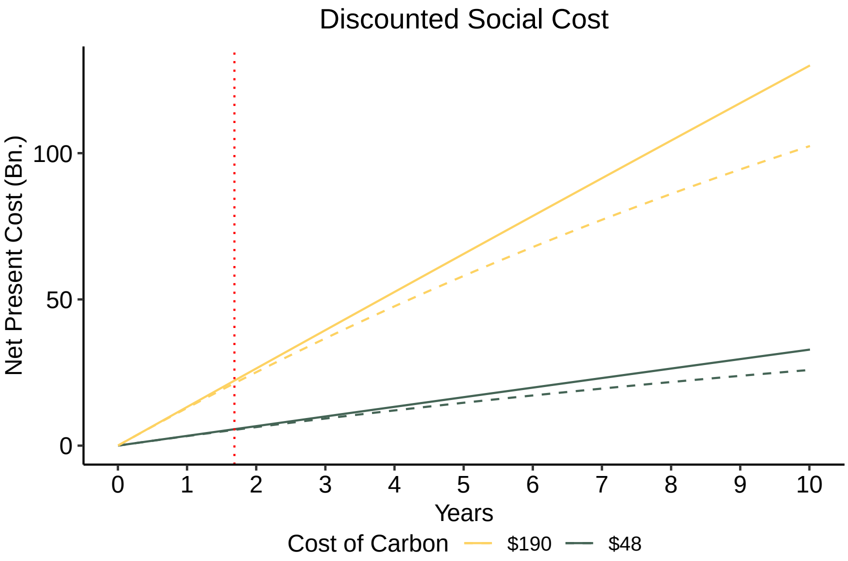

Figure B.16: Net Present Social Cost

This figure plots the net present social cost for each day over a decade using the effects estimated from the hourly power plant data. The yellow lines represent the cost using the EPA's current social cost of carbon of $190, and the dark green uses the EPA's social cost of carbon when the Clean Power Plan was issued ($48). Further, the solid lines are discounted at the risk-free rate from December 2015 (0.375%), while the dashed lines are discounted at the risk-free rate from December 2023 (5.375%). The dotted line represents the number of days from when the EPA announced it would seek to regulate power plant emissions to the stay (615).

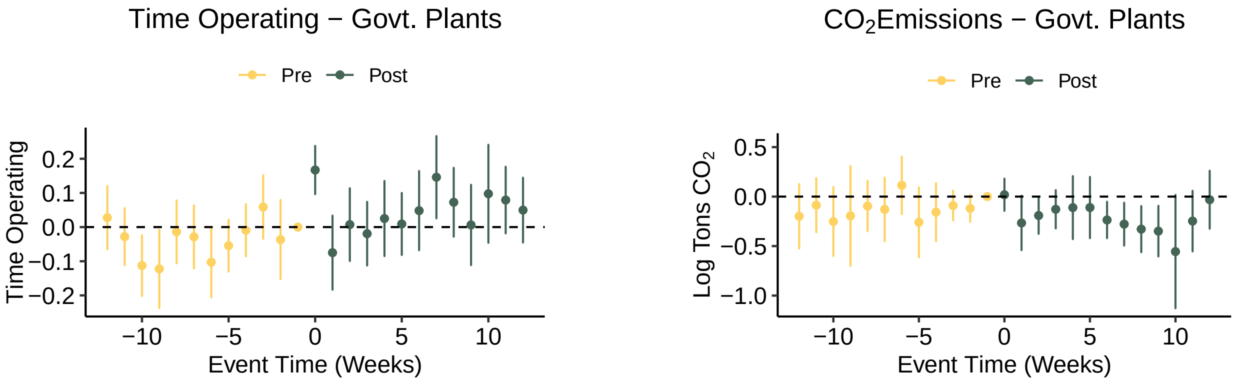

(a) Extensive Margin - Govt. Plants

(b) Intensive Margin - Govt. Plants

Figure B.17: Government Ownership Tests

This figure plots the triple interaction of the event time, treated variable, and a dummy for whether a plant is government-owned using the specification from column (5) in Table C.5 and aggregates it to the weekly level. The dependent variable in panel A is the time operating each week and in panel B is the log tons of carbon dioxide emissions. Confidence intervals are at the 95% level.

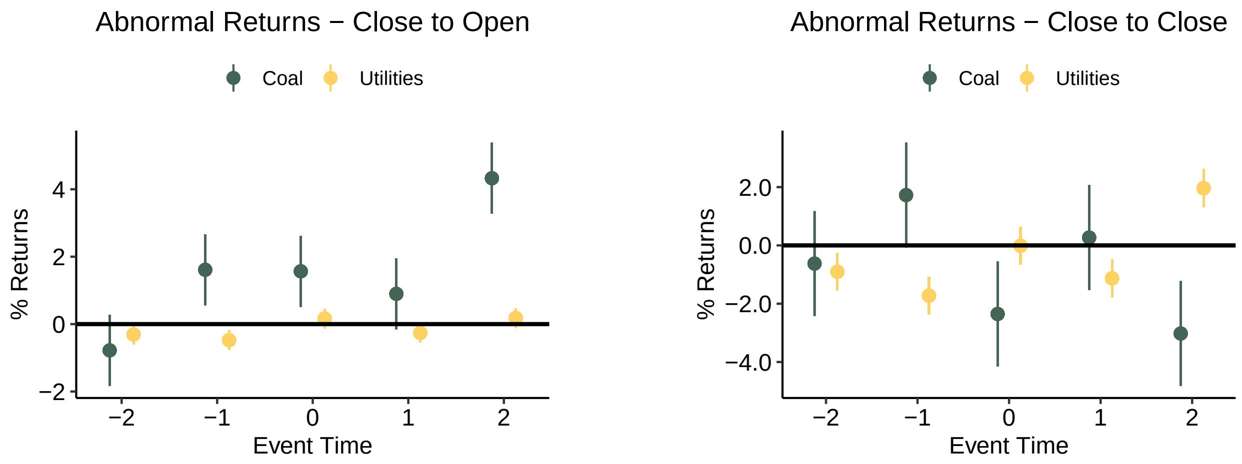

(a) Close-to-Open Returns

(b) Close-to-Close Returns

Figure B.18: Daily Abnormal Returns - Scalia's Passing

Coefficients are the abnormal returns of the value-weighted industry portfolio calculated using a Fama-French Three Factor plus Momentum model. Error bars are clustered by date, and Bonferroni corrected at the 95% level.

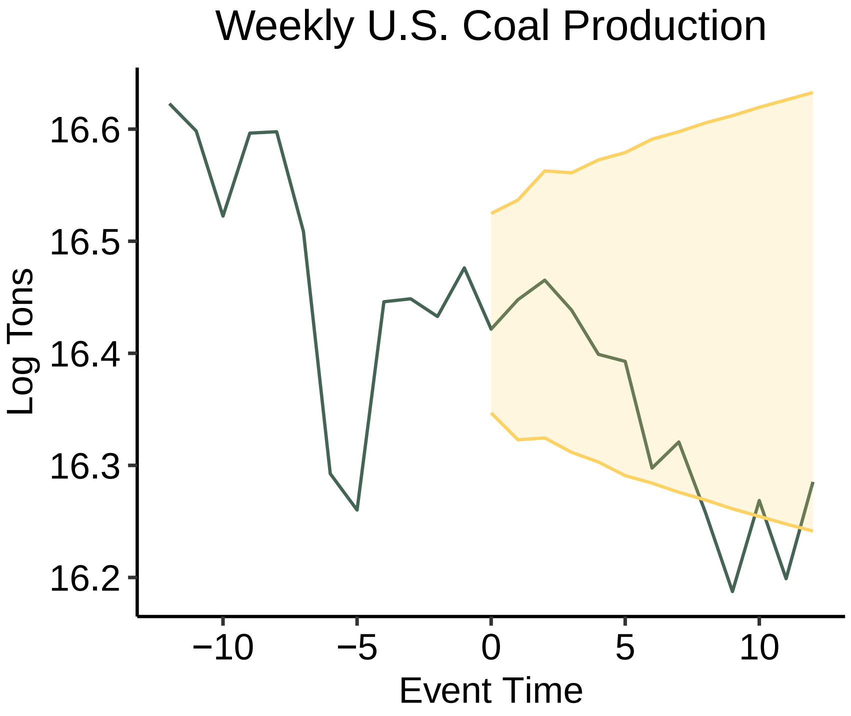

Figure B.19: Log Tons of Coal Production

This figure plots the log tons of coal produced each week as the dark line. The 95% confidence intervals from the ARIMA(3,1,0) model are projected in yellow.

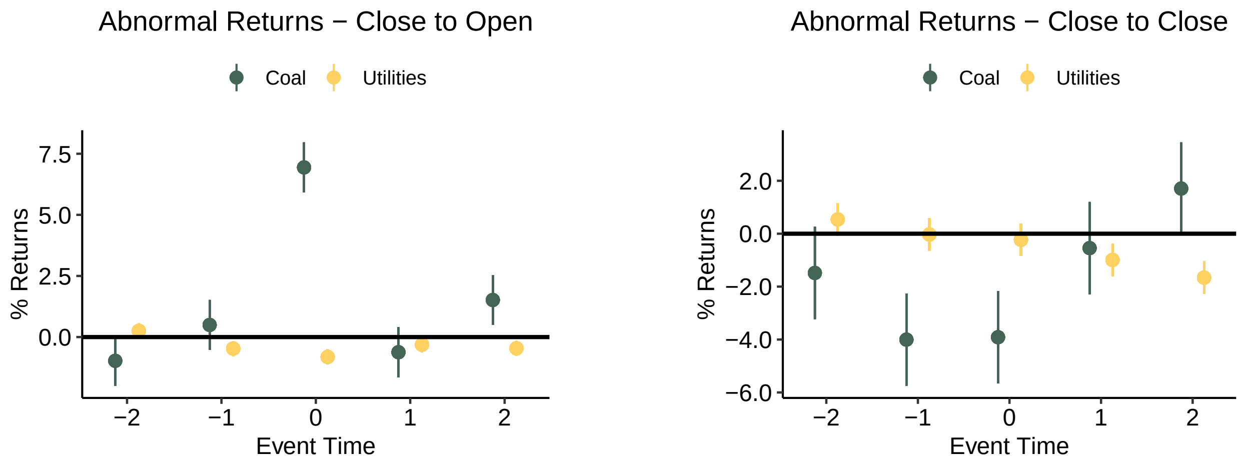

(a) Close-to-Open Returns

(b) Close-to-Close Returns

Figure B.20: Daily Abnormal Returns

Coefficients are the abnormal returns of the value-weighted industry portfolio calculated using a Fama-French Three Factor plus Momentum model. Error bars are clustered by date, and Bonferroni corrected at the 95% level.

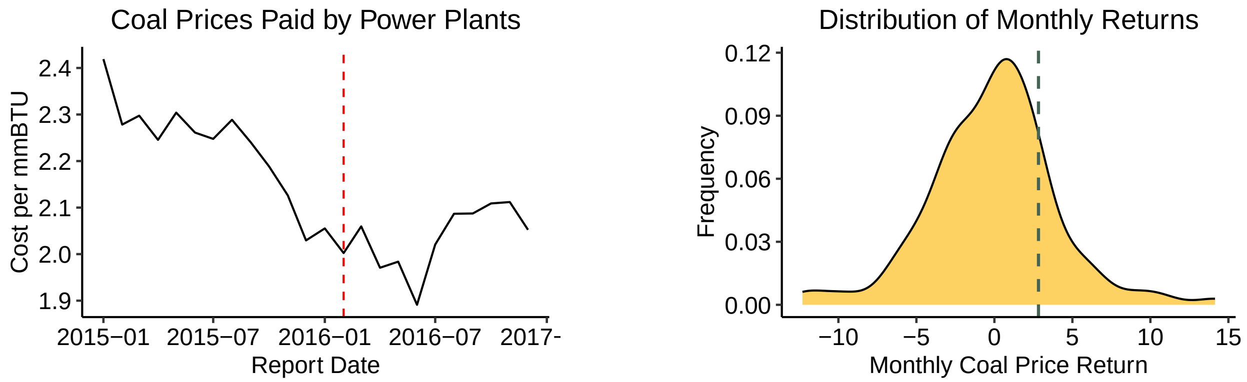

(a) Coal Prices Paid

(b) Distribution of Log Returns to Coal

Figure B.21: Prices Paid and Coal Price Returns

Panel A presents the raw time series of prices paid by power plants for coal as reported on Schedule 2 of EIA Form 923. Panel B plots the distribution of the log monthly returns from coal prices paid between 2014 and 2016. The dotted black line indicates the monthly log return in February 2016, when the Clean Power Plan was stayed.

APPENDIX C

TABLES

Table C.1: Summary Statistics

Panel A: Out-of-Merit Generation Sample

	Pre						Post					
	N	Mean	SD	P25	P50	P75	N	Mean	SD	P25	P50	P75
Out of Merit Gen. (MwH)	787863	6.6	1.5	5.6	6.5	7.5	716143	6.6	1.4	5.7	6.5	7.5
Avg. Precip.	787863	2.5	5.7	0.0	0.3	2.4	716143	2.5	6.4	0.0	0.2	2.4
Avg. Temp.	787863	14.9	9.7	8.4	16.2	22.9	716143	16.4	9.0	10.3	17.8	23.6

Panel B: Ramping Sample

	Pre						Post					
	N	Mean	SD	P25	P50	P75	N	Mean	SD	P25	P50	P75
Ramping (Indic.)	3727224	0.0	0.2	0.0	0.0	0.0	4029900	0.0	0.1	0.0	0.0	0.0
Avg. Precip.	3727224	3.6	8.5	0.0	0.2	2.8	4029900	2.9	6.4	0.0	0.2	3.0
Avg. Temp.	3727224	8.9	8.1	3.7	9.1	14.6	4029900	13.6	7.9	8.9	14.8	19.4

Panel C: PM 2.5 Sample

	Control						Treated					
	N	Mean	SD	P25	P50	P75	N	Mean	SD	P25	P50	P75
Log PM2.5	357431	1.8	0.5	1.5	1.9	2.2	6720	1.8	0.5	1.5	1.8	2.2
Avg. Precip.	357431	3.0	8.7	0.0	0.0	1.4	6720	2.9	7.4	0.0	0.0	1.9
Avg. Temp.	357431	9.6	8.8	3.6	9.8	16.1	6720	7.4	7.8	2.0	7.8	13.2
Traffic Vol.	357431	9.7	1.6	8.6	9.8	10.8	6720	9.2	1.3	8.5	9.2	10.2

Panel D: CEMS Emissions Sample

	Control						Treated					
	N	Mean	SD	P25	P50	P75	N	Mean	SD	P25	P50	P75
Log CO_2	6614046	1.1	1.9	0	0.0	2.8	1091714	1.6	2.5	0.0	0.0	4.3
Operating	6614046	0.3	0.4	0	0.0	1.0	1091714	0.3	0.5	0.0	0.0	1.0
Avg. Precip.	6614046	3.3	7.5	0	0.2	2.9	1091714	3.3	7.0	0.0	0.3	3.2
Avg. Temp.	6614046	11.4	8.3	6	12.0	17.6	1091714	10.1	8.5	4.4	10.4	16.4

Notes. This table provides summary statistics, broken out by treated (pre period) and control (post period) observations, for each sample used in the main specifications. Panel A contains the summary statistics used in the out-of-merit generation analysis. Panel B contains summary statistics for the ramping sample; Panel C for the PM 2.5 sample; Panel D for CEMS emissions sample.

Table C.2: Change in Out-of-Merit Generation

	Log Out-of-Merit Generation (MW)			
	(1)	(2)	(3)	(4)
(Intercept)	6.611***	6.816***		
	(155.202)	(93.752)		
Post	−0.011	−0.091	−0.091	−0.075***
	(−0.173)	(−0.823)	(−0.823)	(−4.340)
Num.Obs.	1 504 006	1 504 006	1 504 006	1 504 006
R2 Adj.	0.000	0.010	0.010	0.937
Weather Controls		X	X	X
Day-of-Week × Hour			X	X
Day-of-Year × Balancing-Authority				X

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes. The dependent variable is the log of out-of-merit generation in megawatts. *Post* is an indicator variable equal to one for all observations after the issuance of the stay in *West Virginia v. EPA*. Column (1) provides the results absent any controls. Column (2) adds weather controls interacted with the *Post* indicator. Column (3) adds day-of-week-hour fixed effects to control for intra-week cyclicality in electricity demand. Column (4) adds day-of-year-balancing authority fixed effects to control for seasonal variation. Standard errors are clustered by balancing authority × month × year.

Table C.3: Change in Coal Power Plant Ramping Behavior

	Ramping (Indicator)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(Intercept)	0.024***	0.024***						
	(19.061)	(15.542)						
Treated	0.014**	0.023***			0.006	0.011+		
	(2.684)	(3.639)			(0.984)	(1.917)		
Post	−0.004*	−0.004*	0.000	0.000	0.000	−0.001	0.000	−0.003*
	(−2.324)	(−2.301)	(0.167)	(0.167)	(−0.162)	(−0.494)	(0.167)	(−2.311)
Treated × Post	−0.013*	−0.007	−0.011***	−0.011***	−0.011***	−0.011**	−0.005**	−0.002
	(−2.168)	(−1.341)	(−4.529)	(−4.529)	(−4.465)	(−3.268)	(−2.864)	(−0.617)
Post × Top Decile							0.005	
							(1.118)	
Treated × Post × Top Decile							−0.030**	
							(−3.234)	
Post × Most Treated								0.007***
								(4.515)
Treated × Post × Most Treated								−0.017***
								(−3.373)
Num.Obs.	7 533 381	7 533 381	7 533 381	7 533 381	7 533 381	7 533 381	7 476 169	7 533 381
R2 Adj.	0.001	0.002	0.105	0.105	0.102	0.081	0.106	0.106
Weather Controls		X	X	X	X	X	X	X
State				X	X	X	X	X
Unit × Day-of-Week × Hour			X	X			X	X
Plant × Day-of-Week × Hour					X			
Utility × Day-of-Week × Hour						X		

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes. The dependent variable is an indicator equal to one if a given power-generating unit is ramping, as measured by coming off of its minimum/maximum generation capacity by at least 10%, in a given hour. *Treated* is an indicator variable that equals one for all fossil fuel-powered generating units that do not use natural gas. *Post* is an indicator variable set to one for all observations after the issuance of the stay in *West Virginia v. EPA*. *Top Decile* is an indicator variable set to one if the unit was in the top decile of most polluting plants in the 2012-2014 period. *Most Treated* is an indicator variable equal to one for generating units located in a state that required above-median reductions in GHG emissions under the Clean Power Plan. Column (1) presents the results without controls, column (2) adds controls, and column (3) adds Unit-Day-of-Week-Hour fixed effects to control for unobserved unit heterogeneity and its interaction with intra-week cyclicality in electricity demand. Column (4) adds State fixed effects to control for state-level policies and macroeconomic fluctuations. Column (5) uses Plant-Day-of-Week-Hour fixed effects to isolate changes between generating units at the same power plant facility. Relatedly, column (6) uses Utility-Day-of-Week-Hour fixed effects to isolate within-utility changes on the extensive margin. Columns (7) and (8) explore heterogeneous treatment effects for the dirtiest power units and the most treated states. Standard errors are clustered by balancing authority × month × year. T-statistics are reported in parentheses.

Table C.4: Change in Zip Code PM2.5 Concentrations

	Log PM2.5 Concentration					
	(1)	(2)	(3)	(4)	(5)	(6)
(Intercept)	1.891***	1.883***	1.843***			
	(75.542)	(72.599)	(44.461)			
Treated	0.001	0.009	0.032			
	(0.042)	(0.389)	(0.959)			
Post	0.004	0.003	−0.041*	0.003	−0.141***	−0.180*
	(0.202)	(0.152)	(−2.054)	(0.143)	(−3.587)	(−2.435)
Treated × Post	0.009	0.010	0.011	0.017	−0.065*	−0.118+
	(0.833)	(0.907)	(0.786)	(1.564)	(−2.177)	(−1.707)
Spillover		0.039*	0.062***	0.011*	0.001	−0.010
		(2.647)	(3.979)	(2.025)	(0.268)	(−1.450)
Traffic Volume						0.052***
						(4.898)
Num.Obs.	8 697 714	8 697 714	7 151 090	7 151 090	7 151 090	363 766
R2 Adj.	0.000	0.001	0.063	0.389	0.572	0.609
Weather Controls			X	X	X	X
Zip Code				X		
Zip Code × Week					X	X

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes. The dependent variable is the log of the PM2.5 concentration for a given zip code on a particular day. *Treated* is an indicator variable set to one for zip codes that contain a coal-fired power plant. *Post* is an indicator variable equal to one for all observations after the issuance of the stay in *West Virginia v. EPA*. *Spillover* is an indicator variable equal to one if the emissions trajectory from a power plant crossed into a zip code within the last 24 hours. *Traffic Volume* is the log of the count of vehicles passing through monitoring stations in a given zip code on that day. Column (1) provides the results absent any controls. Column (2) adds controls for the spillover group, and column (3) adds the controls for weather. Column (4) adds zip code fixed effects to control for unobserved zip code heterogeneity, and column (5) adds zip code-week fixed effects to control for zip code-level trends. Column (6) adds traffic controls to column (5). Standard errors are clustered by state. T-statistics are reported in parentheses.

Table C.5: Change in Coal Power Plant CO_2 Emissions

	Log Tons of CO_2							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(Intercept)	4.228***	4.097***						
	(109.415)	(86.246)						
Treated	1.237***	1.338***			0.550*	1.772***		
	(16.917)	(14.628)			(2.489)	(8.496)		
Post	−0.048	−0.103*	−0.002	−0.002	0.000	−0.003	−0.002	−0.015
	(−0.924)	(−2.117)	(−0.207)	(−0.207)	(0.024)	(−0.144)	(−0.221)	(−1.091)
Treated × Post	−0.050	−0.025	−0.057*	−0.057*	−0.077**	−0.057	0.001	−0.008
	(−0.524)	(−0.263)	(−2.549)	(−2.549)	(−2.671)	(−0.632)	(0.024)	(−0.184)
Post × Top Decile							0.036	
							(0.447)	
Treated × Post × Top Decile							−0.152+	
							(−1.660)	
Post × Most Treated								0.041*
								(2.551)
Treated × Post × Most Treated								−0.096*
								(−2.135)
Num.Obs.	2 051 557	2 051 557	2 051 557	2 051 557	2 051 557	2 051 557	2 051 557	2 051 557
R2 Adj.	0.178	0.185	0.771	0.771	0.724	0.571	0.771	0.771
Weather Controls		X	X	X	X	X	X	X
State				X	X	X	X	X
Unit × Day-of-Week × Hour			X	X			X	X
Plant × Day-of-Week × Hour					X			
Utility × Day-of-Week × Hour						X		

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes. The dependent variable is the log tons of carbon dioxide emitted by a power unit for each hour. *Treated* is an indicator variable that equals one for all fossil fuel-powered generating units that do not use natural gas. *Post* is an indicator variable set to one for all observations after the issuance of the stay in *West Virginia v. EPA*. *Top Decile* is an indicator variable set to one if the unit was in the top decile of most polluting plants in the 2012-2014 period. *Most Treated* is an indicator variable equal to one for generating units located in a state that required above-median reductions in GHG emissions under the Clean Power Plan. Column (1) presents the results without controls, column (2) adds controls, and column (3) adds Unit-Day-of-Week-Hour fixed effects to control for unobserved unit heterogeneity and its interaction with intra-week cyclicality in electricity demand. Column (4) adds State fixed effects to control for state-level policies and macroeconomic fluctuations. Column (5) uses Plant-Day-of-Week-Hour fixed effects to isolate changes between generating units at the same power plant facility. Relatedly, column (6) uses Utility-Day-of-Week-Hour fixed effects to isolate within-utility changes on the intensive margin. Columns (7) and (8) explore heterogeneous treatment effects for the dirtiest power units and the most treated states. Standard errors are clustered by balancing authority × month × year. T-statistics are reported in parentheses.

Table C.6: Change in Coal Power Plant CO2 Emissions (Robustness)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(Intercept)	4.231***	4.082***						
	(104.885)	(81.095)						
Treated	1.225***	1.344***			0.546*	1.786***		
	(16.322)	(14.423)			(2.470)	(8.277)		
Post	−0.049	−0.111*	−0.003	−0.003	0.000	−0.008	−0.003	−0.016
	(−0.895)	(−2.186)	(−0.217)	(−0.217)	(−0.041)	(−0.396)	(−0.232)	(−1.128)
Treated × Post	−0.040	−0.007	−0.057*	−0.057*	−0.077**	−0.055	0.000	−0.006
	(−0.409)	(−0.078)	(−2.534)	(−2.533)	(−2.630)	(−0.604)	(0.015)	(−0.153)
Post × Top Decile							0.037	
							(0.460)	
Treated × Post × Top Decile							−0.153+	
							(−1.673)	
Post × Most Treated								0.043*
								(2.589)
Treated × Post × Most Treated								−0.099*
								(−2.180)
Num.Obs.	1 941 443	1 941 443	1 941 443	1 941 443	1 941 443	1 941 443	1 941 443	1 941 443
R2 Adj.	0.182	0.190	0.771	0.771	0.723	0.579	0.771	0.771
Weather Controls		X	X	X	X	X	X	X
State				X	X	X	X	X
Unit × Day-of-Week × Hour			X	X			X	X
Plant × Day-of-Week × Hour					X			
Utility × Day-of-Week × Hour						X		

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes. The dependent variable is the log tons of carbon dioxide emitted by a power unit for each hour after dropping New Jersey, Ohio, and Oregon from the sample. *Treated* is an indicator variable that equals one for all fossil fuel-powered generating units that do not use natural gas. *Post* is an indicator variable set to one for all observations after the issuance of the stay in *West Virginia v. EPA*. *Top Decile* is an indicator variable set to one if the unit was in the top decile of most polluting plants in the 2012-2014 period. *Most Treated* is an indicator variable equal to one for generating units located in a state that required above-median reductions in GHG emissions under the Clean Power Plan. Column (1) presents the results without controls, column (2) adds controls, and column (3) adds Unit-Day-of-Week-Hour fixed effects to control for unobserved unit heterogeneity and its interaction with intra-week cyclicality in electricity demand. Column (4) adds State fixed effects to control for state-level policies and macroeconomic fluctuations. Column (5) uses Plant-Day-of-Week-Hour fixed effects to isolate changes between generating units at the same power plant facility. Relatedly, column (6) uses Utility-Day-of-Week-Hour fixed effects to isolate within-utility changes on the intensive margin. Columns (7) and (8) explore heterogeneous treatment effects for the dirtiest power units and the most treated states. Standard errors are clustered by balancing authority × month × year. T-statistics are reported in parentheses.

Table C.7: Supreme Court Vacancy - Event Day Abnormal Returns

	Abnormal Event Day Returns (%)			
	Close-to-Open	Close-to-Close		
	FF-Momentum	Market Model	Fama-French	FF-Momentum
Coal	1.567	0.921	0.689	-2.352
	(3.834)	(1.910)	(1.154)	(-3.373)
Utilities	0.165	-0.369	-0.056	-0.012
	(1.484)	(-2.454)	(-0.282)	(-0.049)

Notes. Coefficients are the abnormal returns of the value-weighted industry portfolio on the day of Justice Scalia's passing. T-statistics are reported in parentheses using standard errors clustered by date.

Table C.8: Purchases Channel

	Log mmBTUs Purchased		Indicator Variable	
	Spot Purchases	Contract Purchases	(Spot Purchase)	(New Contract)
Post	0.444***	0.268*	0.014+	−0.020***
	(3.936)	(2.539)	(1.735)	(−4.364)
Avg. Load	−0.008***	−0.003*	0.000	0.000
	(−3.847)	(−2.330)	(0.732)	(−1.028)
Treated × Post	−0.835***	−0.255+	−0.048***	−0.015**
	(−4.548)	(−1.756)	(−3.465)	(−3.209)
Num.Obs.	15 079	15 079	15 079	15 079
R^2 Adj.	0.817	0.918	0.819	0.498
Year	X	X	X	X
Plant	X	X	X	X

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes. This table presents changes in the purchasing behavior of power plants. The dependent variable in the first column is the log tons of spot purchases in a month. Column 2's dependent variable is the log tons of monthly contractual purchases. Column 3 has a dependent variable that is an indicator equal to one if the power plant conducts a spot purchase in a given month. Column 4 dependent variable is an indicator set to one if the power plant entered a new contract in a given month. *Post* is an indicator variable equal to one for all months after the stay in *West Virginia v. EPA*. *Treated* is an indicator variable equal to one if the power plant is coal-fired. All specifications include year and plant fixed effects to control for plant characteristics and macroeconomic shocks. Standard errors are clustered by plant. T-statistics are reported in parentheses.

Table C.9: West Virginia v EPA - Event Day Abnormal Returns

	Abnormal Event Day Returns (%)			
	Close-to-Open	Close-to-Close		
	FF-Momentum	Market Model	Fama-French	FF-Momentum
Coal	6.942	-4.937	-3.816	-3.912
	(17.493)	(-11.328)	(-6.153)	(-5.779)
Utilities	-0.814	-0.229	-0.227	-0.228
	(-6.795)	(-1.710)	(-1.119)	(-0.958)

Notes. Coefficients are the abnormal returns of the value-weighted industry portfolio on the day the stay was issued in *West Virginia v. EPA*. T-statistics are reported in parentheses using standard errors clustered by date.

Table C.10: Log Coal Returns - Price Paid

	Log Returns		
	(1)	(2)	(3)
(Intercept)	−0.850	−1.488*	−1.878**
	(−1.402)	(−2.444)	(−3.123)
Stay	3.676+	2.981	3.407+
	(2.008)	(1.630)	(1.863)
Num.Obs.	24	24	24
R2 Adj.	0.020	0.027	0.142
Year		X	X
Quarter			X

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes. This table presents changes in log returns of coal prices paid by power plants following the event date. The dependent variable is the log return of the monthly weighted average coal price per mmBTU. *Stay* is an indicator variable equal to one for March 2016, the month after the stay in *West Virginia v. EPA*. Standard errors are bootstrapped with resampling conducted 10,000 times. T-statistics are reported in parentheses.