

Policy Uncertainty in the Market for Coal Electricity: The Case of Air Toxics Standards*

Gautam Gowrisankaran[†] Ashley Langer[‡] Wendan Zhang[§]

January 28, 2023

Abstract

Uncertainty over government policy affects important and irreversible decisions, including firm technology adoption, entry, and exit. The real options literature has found that higher outcome variance may delay irreversible decisions. This paper considers the effect of changes in the timing of uncertainty resolution surrounding the Mercury and Air Toxics Standard (MATS). We estimate a dynamic oligopoly model of technology adoption and exit for coal-fired electricity generators that recovers generators' beliefs regarding future MATS enforcement. We develop a computationally tractable equilibrium concept called Approximate Belief Oligopoly Equilibrium in which all agents make decisions incorporating their impact on a reduced set of aggregate market states. Generators subject to MATS experienced substantial policy uncertainty: the perceived enforcement probability was as low as 43%. Removing policy uncertainty increases expected discounted generator profits by \$0.930 billion, but increases pollution by \$0.809 to \$2.206 billion. Resolving uncertainty early attenuates extreme outcomes in expectation, in this case reducing expected exit.

JEL Codes: L13, Q48, Q52

Keywords: environmental policy, air pollution, mercury, MATS, dynamic equilibrium, moment-based Markov equilibrium

*We thank Mike Abito, James Archsmith, Nick Buchholz, Liran Einav, Jonathan Elliott, David Evans, Matt Grennan, Ryan Kellogg, Derek Lemoine, Louis Preonas, Michael Riordan, Nick Ryan, Gabriel Weintraub, and seminar and conference participants at multiple institutions for helpful comments. We are grateful to Kiran Dwivedi, Utkarsh Kumar, and Lorelei Loraine for excellent research assistance.

[†]Columbia University, CEPR, and NBER

[‡]University of Arizona and NBER

[§]Renmin University of China

1 Introduction

Uncertainty over government policy affects important and irreversible decisions, such as firm technology adoption, entry, and exit. In many settings, the process of forming and implementing policies creates uncertainty. For instance, in the U.S., the government frequently enacts new policies by passing legislation that empowers agencies to develop specific regulations. This approach may lead to more responsive policies by allowing existing legislation to respond to new circumstances and technologies. However, developing regulations takes time and they may be subject to court challenges and executive leadership changes that generate policy uncertainty.

The well-developed real options literature has investigated the impact of uncertainty as characterized by outcome variance, finding that increasing outcome variance may lead to delayed irreversible decisions and higher costs (Teisberg, 1993; Dixit and Pindyck, 1994). However, in policy settings, uncertainty is often characterized by whether the policy will be enforced, rather than by outcome variance alone. In this case, the effect of earlier uncertainty resolution on irreversible decisions such as investment is less clear. For instance, consider a firm that must make a drilling investment at a given point in time and faces a 50% chance of an environmental policy that would move oil prices from \$50 to \$70 per barrel if adopted (Kellogg, 2014; Herrnstadt et al., 2020). If the uncertainty surrounding the policy is resolved before the drilling decision point rather than after, the probability of drilling could either decrease or increase: if drilling costs \$55 per barrel, early uncertainty resolution would decrease the expected probability of drilling from 100 to 50%,¹ while if it is \$65, it would increase the expected probability from 0 to 50%. While the direction of the change in investment varies, resolving uncertainty early here always attenuates extreme outcomes in expectation.

This paper considers a major pollution regulation in the electricity sector where enforcement was uncertain: the Mercury and Air Toxics Standard (MATS). MATS required that electricity generators adopt substantial pollution abatement equipment in order to remain in the market. We estimate agents' beliefs regarding the likelihood of MATS enforcement by modeling generators' technology adoption and exit decisions within a dynamic oligopoly

¹With uncertainty resolved after the drilling decision point, drilling will occur because costs are less than the expected oil price of \$60/barrel.

framework. We then use our estimates to simulate how the timing of uncertainty resolution affects counterfactual outcomes in the industry, including generator exit, equilibrium costs, and pollution. Because firms in many sectors—including healthcare, telecommunications, and finance—are oligopolists making irreversible decisions in the face of similarly uncertain policies, our methods and approaches may be more broadly applicable.

Coal-fired electricity generating units (which we refer to as generators) are the primary emitters of air toxics from electricity production. These pollutants, which include mercury, benzene, and arsenic, cause cancer, birth defects, and other serious illnesses. Despite their dangers, federal regulation of air toxics has come relatively recently and been highly uncertain: the Environmental Protection Agency (EPA) did not release the final MATS rule until 2012 with a scheduled enforcement date in 2016. In the MATS regulatory impact analysis, the EPA calculated that compliance via technology adoption would cost generators \$9.6 billion but provide substantial pollution reduction benefits (Environmental Protection Agency, 2011). Partially because of these costs, MATS has been subject to extensive judicial and administrative review—including by the U.S. Supreme Court—but ultimately survived. Given the large pollution externalities of coal generation, policy uncertainty surrounding MATS potentially has important financial and environmental ramifications.

To answer our research questions, we estimate a dynamic oligopoly model of coal generator actions and beliefs over the period 2006-17 and then perform policy counterfactuals on our estimated model. We focus on merchant generators—or independent power producers (IPP)—because they face market incentives rather than rate-of-return regulation. In our model, each year, generators are characterized by a capacity, heat rate (i.e., fuel inefficiency), location, and whether they have adopted abatement technology. Those generators potentially subject to MATS form an expectation about the probability of 2016 enforcement. They then simultaneously decide whether to adopt abatement technology (if they have not already adopted), exit, or continue operating without adopting. Following this choice, generators earn profits by supplying electricity to hourly markets within the year. Profits are equal to revenues from electricity sales minus the costs of fuel, ramping (increasing their generation level), and operations & maintenance (O&M).

Equilibrium effects are potentially important in our context. For instance, one generator's

exit will increase rivals' profits and decrease their likelihoods of exit. Further, a generator may adopt abatement technology partially to signal a commitment to rivals to remaining in the market (Riordan, 1992; Schmidt-Dengler, 2006). Yet the potentially large number of generators in each market results in a curse of dimensionality that makes estimation and the computation of equilibria difficult.

We therefore develop an equilibrium concept called Approximate Belief Oligopoly Equilibrium (ABOE). In an ABOE, all agents are oligopolists that compete in a Markov equilibrium and recognize that their actions affect market states. However, rather than keeping track of each competitor's status, agents keep track of aggregate state information. Each generator's beliefs about how these states evolve—conditional on its own actions—are consistent with competitors' equilibrium actions as approximated by conditional moments, in our case AR(1) processes. Thus, in our setting, generators use approximately correct beliefs to make technology adoption and exit decisions. ABOE builds on approximate equilibrium techniques such as moment-based Markov equilibrium (MME, Ifrach and Weintraub, 2017) by treating all agents as oligopolists.²

In our application, the aggregate states are the market coal capacity, abatement technology adoption share, and natural gas to coal fuel price ratio. The first two aggregate states capture information about expected future profits that occurs in the oligopoly context. They also model generators' potential preemption strategies—of adopting abatement technology and not exiting—for which it is crucial that they understand that their actions affect the market state evolution, as modeled by the ABOE. Finally, the fuel price ratio is a major determinant of the expected costs and revenues from electricity sales. We model generators' beliefs about these three states as AR(1) processes conditional on the generator's own actions that approximate the aggregate outcomes of rivals' unobservable cost shocks. Beyond aggregating information into a market state, a limitation of our specification is that we do not model ownership linkages across generators.

In general, it might be difficult to separately identify our key parameters—the perceived probabilities of future MATS enforcement—from the exit scrap value, since increases in either

²ABOE also relates to the game theoretic literature that discusses equilibrium convergence when agents have approximately correct beliefs (Esponda and Pouzo, 2016; Bohren and Hauser, 2021).

would encourage generators to exit. However, while the federal government was formulating air toxics policies, some U.S. states mandated air toxics reductions for generators within their borders. These standards were either legislative or developed with input from local power producers, and hence were largely not subject to the same level of uncertainty. Our estimator therefore compares exit rates between generators subject to these standards and those subject to MATS—after controlling for other differences—to identify the perceived enforcement probabilities.

Relationship to the Literature: This paper builds on three main literatures. First, we extend a recent literature that measures economic and policy uncertainty and evaluates the impact of this uncertainty on economic outcomes. Baker et al. (2016), Handley and Li (2020), and Langer and Lemoine (2020) develop measures of uncertainty using newspaper text, SEC filing text, and options prices, respectively. Among other papers, Handley and Limão (2017), Dorsey (2019), and Johnston and Parker (2022) examine the impact of uncertainty on trade flows, energy industry investment, and home prices, respectively. We add to this literature by recovering generators’ beliefs over enforcement probabilities and using the estimates to perform counterfactual simulations on the timing of uncertainty resolution.

Second, we contribute to a literature that estimates structural models of electricity markets, e.g. Fowlie (2010), Abito et al. (2022), Linn and McCormack (2019), Scott (2021), and Elliott (2022). Our paper develops a dynamic oligopoly model of the electricity industry that incorporates policy uncertainty. It also relates closely to two papers by an overlapping set of co-authors. First, we use the approach to estimating ramping and O&M costs from Borrero et al. (2023) to calculate annual generator profits. Second, like Gowrisankaran et al. (2023) we model how policies affect dynamic decisions in the electricity industry. However, Gowrisankaran et al. consider the role of U.S. state rate-of-return regulation in energy transitions, while this paper focuses on policy uncertainty faced by generators subject to market incentives.

Finally, we advance the literature on dynamic oligopoly models with approximately correct beliefs. In this literature, MME (Ifrach and Weintraub, 2017) builds on Oblivious Equilibrium (Weintraub et al., 2008), by allowing for aggregate shocks and equilibrium computation when an industry is not in a steady state. ABOE builds on MME by allowing all actors

to understand that they affect the aggregate market states, whereas in MME, the researcher classifies each actor as either fringe or dominant. Recent empirical MME applications include Gerarden (2022), Vreugdenhil (2020), Corbae and D’Erasmus (2021), and Jeon (2022).

Summary of findings: We estimate that generators’ perception of the probability of MATS enforcement started at 100% in 2012, dropped to as low as 43.0% in 2014, and then rose to 99.9% in 2015, for an average probability of enforcement of 78.27%. Further, we estimate that exit and technology adoption are both costly, with an exit costing the generator \$196 million, technology adoption to comply with U.S. state standards costing \$151 million, and technology adoption to comply with MATS costing \$550 million. Our model predicts that generators subject to MATS would spend \$7.3 billion on technology adoption and \$19.24 billion on total exit costs, discounted to 2012.

Our counterfactual analyses investigate the impact of resolving policy uncertainty earlier and reducing generator exit costs. We calculate the effect of reducing policy uncertainty by investigating the differences in equilibrium outcomes under a setting where policy implementation is decided in 2016 versus one where there is a commitment in 2012 to whether the standard will be enforced in 2016, both with the mean generator perceived probability of enforcement. Resolving uncertainty immediately would increase expected generator profits by \$930 million in present discounted value. Nevertheless, eliminating uncertainty also *increases* expected pollution by about 65 million pounds of SO₂ over 30 years, valued between \$809 million and \$2.206 billion dollars, in part by decreasing generator exit.

As in the oil drilling example above, ex ante uncertainty resolution in our model attenuates extreme outcomes in expectation because it takes the mean of generator decisions over multiple policy environments. We find that early policy uncertainty resolution leads to less exit and more pollution, implying that many coal generators were close to the exit margin.

Removing exit costs—for instance by having the government pay for site remediation—reduces the number of generators in 2016 by 15.1% and increases generator profits by 49.6%. Thus subsidizing exit costs increases exit, but requires substantial government transfers to coal generators.

The remainder of this paper is organized as follows. Section 2 discusses the institutional framework, data, and construction of key variables. Section 3 specifies our structural model

of technology adoption and exit. Section 4 explains our approach to estimation and identification. Section 5 presents our results and counterfactuals. Finally, Section 6 concludes.

2 Institutional Framework and Data

2.1 Background on Regulation of Air Toxics

The EPA regulates 187 air toxics, which are also called hazardous air pollutants,³ under the 1990 Clean Air Act Amendments (CAAA). The EPA’s first attempt to regulate generators’ mercury emissions was the Clean Air Mercury Rule (CAMR), which was finalized in 2005. The courts vacated CAMR in 2008 under *New Jersey v. EPA*,⁴ which found that the EPA should have regulated mercury under a maximum achievable control technology (MACT) standard, instead of CAMR which was a voluntary cap-and-trade regulation (Hudson, 2010). Although the rule was vacated, our data show that some generators did install mercury abatement technologies during the CAMR period.

At approximately the same time that *New Jersey v. EPA* vacated CAMR, the courts found in *Sierra Club v. EPA* that the EPA would have to regulate mercury and other air toxics together, rather than starting with mercury alone.⁵ In response to these decisions, the EPA finalized MATS in 2012, after releasing earlier versions of the proposed rule in 2011. The final MATS rule required generators to comply with MATS by 2015, but extensions to 2016 were built into the rule and were widely granted.

The investments necessary to achieve compliance with MATS are irreversible and costly, implying that generators may not want to adopt these technologies unless they are fairly certain that compliance will be required. MATS compliance technologies convert pollutants into water-soluble forms, bind them to larger particles, and precipitate the new compounds with a particulate matter catcher.⁶ This basic process can be achieved with different technologies, making it difficult to determine compliance from technology adoption data alone.

³<https://www.epa.gov/haps/what-are-hazardous-air-pollutants>.

⁴517 F.3d 574 (D.C. Cir. 2008).

⁵551 F. 3d 1019 (D.C. Cir. 2008), also known as the “Brick MACT” decision.

⁶Compliance may also potentially be achievable by fuel switching to cleaner coal.

EPA regulations have been vulnerable to two key sources of uncertainty. First, as with CAMR, EPA rule-making has been subject to substantial legal challenge, up to and including Supreme Court review. Further, changes in executive leadership have drastically altered the EPA’s focus. These leadership changes interact with legal challenges, e.g., a new administration may change legal approaches.

In the context of MATS, uncertainty arose from both of these sources. The final rule was challenged by several U.S. states’ attorneys general. The result of these challenges was that, in 2015, the Supreme Court remanded MATS to the EPA for additional justification that MATS was “appropriate and necessary.” However, the order left MATS in place, which effectively meant that generators needed to comply by the 2016 deadline. In 2017, the incoming Trump administration did not file the justification but left MATS in place nonetheless.⁷

U.S. states started to develop their own mercury abatement policies during the CAMR period. An early report by the Congressional Research Service lists U.S. states with their own policies, along with preliminary announcement and enforcement dates (Congressional Research Service, 2007). CAMR encouraged the development of these policies, which varied substantially across U.S. states. In some cases—e.g., Florida—these policies were cap-and-trade systems broadly similar to CAMR while in other cases—e.g., New Hampshire—they were effectively voluntary. We are aware of only one U.S. state where the state policy faced a substantial court challenge: Pennsylvania legislators opposed the regulation put forward by the state agency and ultimately the state court overturned it.

We define our sample of U.S. states with state air toxics standards as those with announced rules or enacted legislation with enforcement dates prior to 2016, excluding Pennsylvania. Starting with the Congressional Research Service report, we use a combination of newspaper articles, state environmental agency press releases, and state statutes to verify and update our list of U.S. states with standards and their announcement and enforcement years. On-line Appendix A1 provides details of U.S. state standards, and Table A1 in On-line Appendix A4 lists announcement and enforcement years for these standards.

In contrast to federal regulations, these U.S. state standards were generally subject to very little uncertainty for at least two reasons. First, in some states (e.g. CT and MD), these

⁷In 2021, the Biden Administration did file the justification of MATS.

standards were passed into law by state legislatures (Halloran, 2003; Pelton, 2006). Second, even in states such as IL and MA where the standards were created as rules issued by the state environmental agency, they were generally developed in tandem with the owners of large coal generators (Hawthorne and Tribune staff reporter, 2006; United Press International, 2004), which led to substantially fewer and weaker judicial challenges. For this reason, we can use the decisions of generators subject to U.S. state enforcement to identify the costs of generator exit and compliance in the absence of policy uncertainty.

One complication is that U.S. state air toxics regulations were weaker than MATS in some cases, for at least three reasons. First, the specified standards for mercury levels were sometimes higher than under MATS. Second, the state standards covered mercury rather than all air toxics. Finally, enforcement in some cases consisted of the regulator approving generators' abatement technology adoption plans rather than monitoring ex post outcomes, as occurs under MATS. These factors motivate our modeling assumption that compliance costs vary across the two sets of U.S. states.

Supporting our assumptions on the timing, salience, and importance of MATS to generators not subject to U.S. state standards, many of these generators responded immediately to the 2012 MATS announcement by reporting to the Energy Information Administration (EIA) that they planned to retire. Specifically, at the beginning of 2012, 8.6% of generators subject to MATS newly reported that they would retire between 2012 and 2015. The analogous 2012 increase for generators subject to U.S. state standards was only 1.4%.⁸

Finally, emissions from coal generators were also subject to other pollution regulations during our analysis period. The most important of these is the Cross-State Air Pollution Rule (CSAPR), which regulated SO₂ and NO_x emissions from generators in certain eastern states. Unlike for MATS, generators could comply with CSAPR by purchasing emissions permits. The market price of these permits has generally been quite low, so we do not model the costs of CSAPR compliance.

⁸Calculations from EIA form 860 based on coal generators in Eastern Interconnection.

2.2 Data Sources

While our primary analysis data set is at the generator-year level, we calculate annual generator profits using data at the generator-hour level. Fossil fuel power plants are generally made up of a collection of generators that may have different costs, capacities, and abatement technologies. Because of the differences across generators within a plant, we focus on decisions at the generator, and not plant, level.⁹ Our data sources include information on generators’ production, costs, emissions, market conditions, demand, prices, and abatement technology adoption.

Our primary data source is the EPA’s Continuous Emissions Monitoring System (CEMS) database. These data are at the generator-hour level. Our sample covers 2006 to 2017 for U.S. states in the Eastern Interconnection, which includes the vast majority of IPP coal generators in the U.S.¹⁰ Each observation in the CEMS data provides the heat input of the fuel used (in MMBtu), electricity production (in MWh), and CO₂, NO_x, and SO₂ emissions (in pounds/MMBtu) for each generator that the EPA monitors with a CEMS. The CEMS data further report a facility identifier and the location of each generator. As discussed in On-line Appendix A2, we use SO₂ as a proxy to measure the adoption of air toxics abatement technology and to measure annual generator emissions.

We merge the hourly CEMS data with several other data sources. First, the EPA provides an annual-level data set that includes generator characteristics. We define a generator as using coal if the primary fuel variable includes the word “COAL.” Our definition therefore includes generators that primarily use coal, but also use other fuels.

Second, we merge in annual data from the EIA Form 923 on whether the facility is an independent power producer (IPP) or not. We consider only IPP generators because other generators, which are owned by load-serving entities, may face non-market incentives (Gowrisankaran et al., 2023). We define a coal generator to be an IPP if the facility at which it is located is an IPP at any point in our sample.

Third, we create a data set of annual, U.S. state-level natural gas and coal prices from

⁹We define generators using the EPA’s definition of units, which is based on emissions release points. This corresponds roughly, but not exactly, to the Energy Information Administration’s definition of “boilers,” which is based on fuel combustion.

¹⁰We drop coal generators in Oklahoma, since they all appear to enter after MATS was announced.

EIA Form 423. This form reports fuel prices by generator and year. We aggregate fuel prices to the U.S. state-year level by taking the mean weighted by annual generation at each generator. We use price data at the U.S. state-year level because it measures the opportunity cost of fuel faced by generators.

Fourth, we merge in hourly wholesale electricity prices by U.S. state. We obtain prices for nodes in each Regional Transmission Organization (RTO) or Independent System Operator (ISO) in the Eastern Interconnection.¹¹ For some U.S. states, the data report prices for multiple nodes. For these states we take the mean over the nodes for each hour. For other states, e.g. Georgia, there is no reported electricity price. In these cases, we assign the price from the node that is geographically closest to the state.

Fifth, we deflate these prices to January, 2006 dollars. To do this, we use the Bureau of Labor Statistics' chain-weighted consumer price index for urban consumers.

Sixth, we recover hourly U.S. state-level electricity load from the Public Utility Data Liberation (PUDL) database, which derives its data from the Federal Energy Regulatory Commission (FERC) Form 714. PUDL reports multiple measures of load. In particular, we use the reported load scaled to match the total annual load at the state level in EIA Form 861.

Finally, we use county-level weather data from PRISM. We aggregate these data to the U.S. state level by calculating the population-weighted mean of the daily minimum and maximum temperatures, using annual population data from the U.S. Census. Following Schlenker and Taylor (2021), we recover daily heating degrees, which measure the amount by which population-weighted state average daily temperature exceeds 65 degrees (if at all), and analogously cooling degrees.

We use these data to construct a number of key variables at the generator level, specifically the year of abatement technology adoption and exit, minimum and maximum generating capacity, and heat rate. On-line Appendix A2 presents details.

Table 1: Generator Descriptive Statistics by Regulatory Regime

	U.S. State Standard	MATS
Capacity (MW)	279.69 (213.78)	245.74 (282.60)
Heat Rate (MMBtu/MWh)	10.51 (2.25)	11.90 (4.64)
Coal Fuel Price (\$/MMBtu)	1.70 (0.19)	2.13 (0.40)
Marginal Fuel Costs (\$/MWh)	17.85 (4.44)	26.23 (13.94)
Generators	93	226
Generator-years	841	2040

Note: Authors' calculations based on annual analysis sample of IPP coal generators. Standard deviations are in parentheses.

2.3 Descriptive Statistics

Table 1 presents generator-level descriptive statistics on our analysis data separately for generators subject to U.S. state air toxics standard enforcement and MATS. Our analysis data contain 319 IPP coal generators, of which 93 are subject to U.S. state enforcement and 226 are subject to MATS. They include a total of 2,881 generator-year observations. These statistics suggest that it is reasonable to assume that generators are technologically similar across the two sets of U.S. states. Specifically, generator mean capacity levels and heat rates are similar across the two sets of U.S. states, though there is substantial variation in these characteristics within each set of U.S. states. Generators in U.S. states with air toxics standards do face lower average coal fuel prices than generators subject to MATS (\$1.70 vs \$2.13 per MMBtu), which then feeds into lower marginal fuel costs (\$17.85/MWh vs \$26.23/MWh).

While generators' underlying characteristics are fairly similar across enforcement regimes, their technology adoption and exit decisions are quite different. Table 2 reports the number of generators that adopt air toxics abatement technology or exit the market by years to enforcement. These data show that generators subject to state standards were much more likely to adopt pollution abatement technology than exit, while the reverse is true for generators subject to MATS. Although this could reflect MATS being stricter than U.S. state standards, it could also be explained by differences in underlying market conditions. Specifically, natural

¹¹Specifically, we retrieve electricity prices from the New England ISO, New York ISO, PJM, Midcontinent ISO, and the Southern Power Pool.


Table 2: Counts of Generators Adopting Abatement Technology or Exiting

Years to Enforcement	State Standard			MATS		
	Adoptions	Exits	Share Complied	Adoptions	Exits	Share Complied
4	12	0	0.34	2	23	0.30
3	5	1	0.51	1	19	0.54
2	9	0	0.77	2	7	0.64
1	4	4	1.00	9	21	1.00
Total	30	11		14	84	

Note: Authors' calculations based on annual analysis sample of IPP coal generators.

gas prices were substantially lower during the MATS enforcement period than earlier.

Figure 1: Share of Generators in Compliance Relative to Base Year, Among Adopters



Note: Share of generators that have adopted air toxics abatement technology among generators subject to U.S. state air toxics standards, relative to the total number of generators that had not adopted abatement technology five years before air toxics standard enforcement. Calculations based on annual analysis sample.

These differences in market conditions at the time of air toxic standard implementation make a simple comparison in outcomes between the two sets of U.S. states difficult. Nonetheless, Figure 1 provides some intuition behind the identification of the perceived probabilities of MATS enforcement over time by comparing generators that adopted abatement technology within four years of enforcement across the two sets. The dashed line shows the share of generators subject to U.S. state standards that had adopted, and the solid line shows this share for generators subject to MATS, both by years to enforcement. Four years prior to enforcement, adoption rates are fairly similar, but at three and two years prior to enforcement, generators subject to MATS had adopted abatement technologies at substantially lower rates than generators subject to state standards, suggesting that, in those years, generators may have believed the probability of 2016 MATS enforcement to be substantially below one.

Table 3: Descriptive Statistics of Hourly Generation

Share of Hours at Maximum Generation	0.44 (0.50)
Share of Hours at Minimum Generation	0.25 (0.43)
Share of Hours at Zero Generation	0.31 (0.46)
Generation (MWh)	174.14 (248.30)
Electricity Price (\$/MWh)	39.14 (25.20)
U.S. State Electricity Demand (GWh)	13.32 (5.35)
Daily Heating Degrees	13.90 (14.88)
Daily Cooling Degrees	2.83 (4.82)
N	26,014,143

Note: Authors' calculations based on hourly generation data of IPP coal generators. Standard deviations are in parentheses.

Finally, Table 3 presents descriptive statistics on the hourly generation data we use to calculate annual generator profits. Over our sample period, IPP coal generators operated at maximum generation for 44% of hours, minimum generation for 25% of hours, and were off for 31% of hours, for a mean hourly generation level of 174.14 MWh. Generators in our sample faced wholesale electricity market prices with a mean of \$39.14/MWh. Mean hourly U.S. state electricity demand was 13.32 GWh, which was driven by a mean of 13.90 heating degree days and 2.83 cooling degree days. All of these variables have substantial variation both over time and across generators.

3 Model

We develop an infinite-horizon dynamic equilibrium model of abatement technology adoption and exit for coal independent power producers (IPPs), which we refer to as “generators” for brevity. Each year, t , there is a set of generators that are currently operating. Each generator, $j = 1, \dots, J_t$, has a time-invariant heat rate, $heat_j$, and capacity, K_j , and an indicator for whether it has active air toxics abatement technology, $Tech_{jt}$. We assume that each U.S. state forms one electricity market. For brevity, our notation considers one market and hence does not include a market index.

We model generators as competing annually in a dynamic oligopoly through their technology adoption and exit decisions. They also compete with natural gas, renewable, utility-owned coal, and other sources. We do not directly model other sources’ entry, exit, technology adoption, and production decisions, but treat them as exogenous, though state-contingent and time-varying. Section 3.1 discusses the state space and equilibrium.

Each year t proceeds as follows. First, the policy environment updates, with policymakers announcing new standards and generators obtaining information about previously announced standards. In some year t_0 , the regulator announces that an air toxics standard will be enforced τ_0 years in the future. Before this year, generators do not expect to be subject to any air toxics regulation. In years when enforcement is $0 < \tau \leq \tau_0$ years away, generators use new information to update their common belief of the probability that enforcement will occur.¹² We denote the full set of perceived probabilities P_{τ_0}, \dots, P_1 . Upon forming beliefs P_τ , generators believe that they will continue to perceive the probability of enforcement to be P_τ until the announced air toxics standard enforcement date. For U.S. states which implemented their own air toxics standards, we assume that $P_{\tau_0} = \dots = P_1 = 1$.

Second, generators make adoption and exit decisions.¹³ Generators that have not yet adopted abatement technology must decide whether to adopt the technology and pay an adoption cost $A - \varepsilon_{jat}$, continue operating without adopting the technology and receive a

¹²While we allow for uncertainty about whether the standard will be enforced, we assume that the level of the standard and the date of potential enforcement is certain.

¹³We do not model generators’ decisions to enter. Coal entry during our sample period is very limited and entry that occurred resulted from prior decisions.

payment ε_{jct} , or exit and earn a scrap value $X + \varepsilon_{jxt}$. The cost shocks to generator j , $\vec{\varepsilon}_{jt} \equiv (\varepsilon_{jat}, \varepsilon_{jct}, \varepsilon_{jxt})$, are type 1 extreme value *i.i.d.* across options, years, and generators, and are generator j 's private information at the decision point. These shocks arise as generators may have idiosyncratic maintenance needs or contracts that cause variation in the costs of adopting technology or exiting. Generators that have already adopted abatement technology or are in years before air toxics standard announcement only need to choose between continuing to operate or exiting. Section 3.2 details annual dynamic optimization.

Third, conditional on the technologies and capacities of generators and other sources, generators earn annual profits from competing in hourly electricity markets. Their annual revenues from selling electricity are the sum of the hourly wholesale electricity market prices times the quantity supplied. Generators bear three types of costs: fuel; ramping; and operation & maintenance (O&M).

Finally, any exit or abatement adoption decisions made in this year are realized. At this point, if this is the final year before potential enforcement, the regulator enforces the air toxics standard with probability P_1 . If it is enforced, generators that have not adopted are forced to exit immediately.

A limitation of our analysis is that we model each generator as an independent optimizer. We believe that coordination in exit and adoption decisions across generators is limited because the market for fossil-fuel electricity is relatively unconcentrated. In particular, we find a mean HHI of 2,321 (with a standard deviation of 1,832) across the states and years that we study.

3.1 Equilibrium and State Space

As discussed above, equilibrium effects are potentially important in our setting, but keeping track of all rivals' choices results in a curse of dimensionality that substantially complicates estimation and equilibrium computation. Accordingly, we develop and apply a concept called Approximate Belief Oligopoly Equilibrium (ABOE). In an ABOE, each player forms perceptions of market evolution conditional on its actions, where these perceptions are consistent with equilibrium play when fit to a simple functional form. Each player chooses dynamically

optimal actions based on these perceptions. Therefore, an ABOE consists of sets of strategies for every player, aggregate market states, and moments on state transitions conditional on the player’s actions such that: 1) the state-contingent strategies reflect optimizing behavior given the aggregate states and moments, and 2) simulated state transitions that condition on the actions of the player but otherwise use the candidate ABOE strategies yield these moments.

Our approach is similar to moment-based Markov equilibrium (MME). In an MME, players also base their expectations of future states on moments. However, in an MME, there are two types of players—dominant and fringe—and only dominant players understand that they can affect the market state. In an ABOE, every player recognizes that it can influence the aggregate state. We account for this awareness by allowing the expected distribution of state transitions to condition on a player’s own actions.

In our case, the aggregate market states include three characteristics: (1) the (combined IPP and non-IPP) coal capacity relative to the 95th percentile of hourly load in the market, (2) the share of IPP coal capacity that has adopted abatement technology, and (3) the natural gas to coal fuel price ratio. Additionally, generator j ’s state includes its non-time-varying characteristics: heat rate, $heat_j$, capacity, K_j , coal price, f^C , and fixed market characteristics; its time-varying characteristics: annual profits, Π_{jt} , $Tech_{jt}$, the belief year, τ , years to potential air toxics standard enforcement, $\tilde{\tau}$,¹⁴ and its cost shocks, $\vec{\varepsilon}_{jt}$.

We include coal capacity and adoption share as market states because they will affect pricing in hourly markets and dynamic oligopoly preemption incentives. We include the fuel price ratio as a market characteristic because relative natural gas fuel price affects generators’ current and expected future profits and therefore their adoption and exit decisions. For instance, if natural gas prices are low, then coal generators expect to operate less, receive lower prices when they operate, pay higher ramping costs, and ultimately earn lower profits.¹⁵

Consistent with an ABOE, generators understand that their actions influence the first two

¹⁴If an air toxics standard has not yet been announced or the enforcement year has already passed, then $\tau = \tilde{\tau} = 0$.

¹⁵We focus on natural gas as the alternative fuel source since other fossil fuels such as distillate fuel oil are more expensive than both coal and natural gas generation during this period. Further, we do not include renewable generation during this period because generation was fairly low in the markets we study.

market state variables. Specifically, a generator which chooses not to exit understands that the coal capacity next period will be higher than if it chooses to exit. Similarly, a generator recognizes that the share of capacity having adopted abatement technology next period will be higher if it has adopted abatement technology than if it has not. Modeling expectations of market evolution in this way allows generators to use their adoption and exit decisions as costly preemptive signals. While we allow for the coal capacity and adoption share variables to be determined in equilibrium, we assume that the fuel price ratio evolves exogenously, meaning that it does not respond to coal technology adoption or exit.

We assume that generators believe that the conditional moments of the three aggregate market state transitions reflect simple autoregressive processes. Specifically, we specify separate AR(1) regressions with normally distributed residuals for the evolution of each of the state variables. These AR(1) processes approximate the combination of exogenous market-level unobservables and the structural unobservables, $\vec{\varepsilon}$, for all generators. Because each generator believes that the market technology adoption share will depend on whether it adopts, we specify different AR(1) regressions for adoption share conditional on that generator’s adoption versus non-adoption choice.¹⁶ Because market fundamentals may vary across U.S. state and belief year τ , we further disaggregate our AR(1) regressions to this level. Following the definition of an ABOE, for every U.S. state and belief year we compute the equilibrium as a fixed point of dynamic optimization decisions and the AR(1) regression coefficients.

3.2 Generator Dynamic Optimization

Generator j makes adoption and exit decisions based on its state, which includes the three aggregated market characteristics noted above—denoted Ω_t , its own time-varying characteristics, $(Tech_{jt}, \tau, \tilde{\tau}, \vec{\varepsilon}_{jt})$, and its non-time-varying characteristics. When $\tilde{\tau} = 0$, generators face a relatively simple choice between exiting or continuing. When $\tilde{\tau} > 0$, generators that have not yet adopted abatement technology face an additional decision of whether to adopt the technology. In this case, the value of continuing depends fundamentally on whether $\tilde{\tau} = 1$

¹⁶We do not specify different AR(1) regressions for the coal capacity ratio since generators that exit no longer value the market state.

or $\tilde{\tau} > 1$ because in years with $\tilde{\tau} > 1$, there will be no enforcement decision at the end of the year and so the generator will get to continue even if it has not adopted.

Consider first the dynamic decision in a year with $\tilde{\tau} = 1$ (one year from enforcement) in the case where the generator has not previously adopted abatement technology.¹⁷ The generator faces the three choices of continuing without adopting, adopting, or exiting. If it continues without adopting and the air toxics standard is enforced, it will be forced to exit and receive X . We can write the Bellman for this case as:

$$\begin{aligned}
& V_j(\Omega, Tech = 0, \tau, \tilde{\tau} = 1, \vec{\varepsilon}_j) = & (1) \\
& \max \{ \Pi_j(\Omega) + P_\tau X + (1 - P_\tau) \beta E[V_j(\Omega', 0, \tau, 0, \vec{\varepsilon}_j') | \Omega, \text{No Standard}] + \sigma \varepsilon_{jc}, \\
& \quad \Pi_j(\Omega) - A + \beta \{ P_\tau E[V_j(\Omega', 1, \tau, 0, \vec{\varepsilon}_j') | \Omega, \text{Standard}] \\
& \quad + (1 - P_\tau) E[V_j(\Omega', 1, \tau, 0, \vec{\varepsilon}_j') | \Omega, \text{No Standard}] \} + \sigma \varepsilon_{ja}, \\
& \quad \Pi_j(\Omega) + X + \sigma \varepsilon_{jx} \},
\end{aligned}$$

where a prime indicates next period's value of the variable. Equation (1) shows that the perceived probability of air toxics standard enforcement enters into the $\tilde{\tau} = 1$ value function.¹⁸ In (1), the first choice is to continue operating without adopting abatement technology. With this choice, with probability P_τ , the air toxics standard will be enforced and the generator will be forced to exit, while with probability $1 - P_\tau$, the air toxics standard will not be enforced and the generator will never be forced to comply. The second choice is to invest in abatement technology this year, which we indicate by updating $Tech$ to 1 in the future state. In this case, the generator is not forced to exit regardless of whether the air toxics standard is enforced. For both of these choices, the continuation values are dependent upon whether the standard is enforced because enforcement will potentially change the set of rivals remaining in the market. The third choice is exit. Generators that have already adopted face a choice between continuing and exiting that mirrors equation (1), except with no adoption choice

¹⁷Because generators believe that their enforcement beliefs will remain constant in the future, we must consider instances where belief year, $\tau > 1$, even though $\tilde{\tau} = 1$.

¹⁸We assume that Π is a function of Ω but not $Tech$ for two reasons. First, our O&M cost estimates are similar for generators that have adopted technology and those that have not. Second, while heat rates do increase slightly following technology adoption, these effects are small (though they could be added to the calculation of profits in a straightforward way).

and no probability of being forced from the market if the standard is enforced.

Turning to the case of $\tilde{\tau} > 1$, we can write the Bellman equation for a generator that has not previously adopted as:

$$V_j(\Omega, 0, \tau, \tilde{\tau}, \vec{\varepsilon}_j) = \max \left\{ \Pi_j(\Omega) + \beta E[V_j(\Omega', 0, \tau, \tilde{\tau} - 1, \vec{\varepsilon}_j) | \Omega] + \sigma \varepsilon_{jc}, \right. \\ \left. \Pi_j(\Omega) - A + \beta E[V_j(\Omega', 1, \tau, \tilde{\tau} - 1, \vec{\varepsilon}_j) | \Omega] + \sigma \varepsilon_{ja}, \Pi_j(\Omega) + X + \sigma \varepsilon_{jx} \right\}. \quad (2)$$

In (2), generators believe that air toxics standard enforcement will be revealed in $\tilde{\tau}$ years and that they will continue to perceive an enforcement probability of P_τ for the next $\tilde{\tau} - 1$ years. Unlike with $\tilde{\tau} = 1$, there is no chance of enforcement occurring in this year.

3.3 Annual Profits

The Bellman equations depend fundamentally on state-contingent annual profits, $\Pi_j(\Omega)$. Generators earn these profits by competing in hourly electricity markets over the year. Following Linn and McCormack (2019), we model the generator as having a choice in each hour between: (1) generation at capacity, K_j ; (2) minimum generation $L_j K_j$ for $L_j \in (0, 1)$; and (3) not generating. Each year, t , includes hours $h = 1, \dots, H$, with generator j choosing a generation quantity $q_{jh} \in \{0, L_j K_j, K_j\}$ in each of these hours.

As modeled in Borrero et al. (2023), we assume that hourly profits are revenues net of fuel, ramping, and operations and maintenance costs, plus *i.i.d.* idiosyncratic, choice-specific unobservables. Annual profits are thus:

$$\Pi_{jt} = \sum_h \pi_{jh} \equiv \sum_h q_{jh} \times [p_h - \text{heat}_j \times f^C - \text{om}] - \mathbb{1}\{\tilde{q}_{jh} < q_{jh}\} r_{j,\tilde{q},q} + \sigma^g \varepsilon_{jhq}^g, \quad (3)$$

where p_h is the per-MWh electricity price, f^C is the coal fuel price (so that $\text{heat}_j \times f^C$ are fuel costs per MWh), om is the per-MWh O&M cost, \tilde{q}_{jh} is the generation quantity in the previous hour, $r_{j,\tilde{q},q}$ is the cost to generator j of ramping from \tilde{q} to q , and ε_{jhq}^g is the idiosyncratic unobservable ramping cost shock that we assume is distributed type 1 extreme value with a standard deviation, σ^g .

We include ramping costs in profits because they have been shown to be important in the

context of electricity generation (Cullen, 2014; Reguant, 2014; Linn and McCormack, 2019), and are particularly critical during our analysis period because of changes in the electricity industry. Specifically, the advent of hydraulic fracturing (“fracking”) led to sharp declines in the price of natural gas fuel starting around 2009. This, in turn, led to coal generators frequently having fuel and O&M costs above wholesale electricity prices, which increased generator cycling and hence the importance of ramping costs. As an example, in 2008, coal generators in our data averaged 31.4 hours at their maximum generation level each time they ramped to maximum generation, but this dropped to 21.0 hours in 2017.

4 Estimation and Identification

4.1 Adoption and Exit Parameters

We estimate our model using a full solution, nested-fixed-point approach, where the unit of observation is a generator in a year.¹⁹ The structural parameters of the model are the probabilities of MATS enforcement, P_4, \dots, P_1 , exit scrap value, X , abatement technology adoption cost, A , and standard deviation of the unobservable, σ . Because of the differences between U.S. state air toxics regulations and MATS discussed in Section 2.1, we allow A to vary based on whether the generator is subject to U.S. state or MATS enforcement. We assume that generators that comply with their U.S. state standards do not face additional costs from MATS compliance.²⁰

For generators in the years between air toxics standard announcement and enforcement which have not yet adopted, the dependent variable is the choice of continuing to operate, exiting, and adopting. For generators in other years, the dependent variable takes on two values, continuing to operate and exiting.²¹

¹⁹We choose a nested-fixed-point approach rather than a CCP approach, because generators subject to MATS base their decisions on subjective enforcement probabilities that they revise over time, implying that the data at time $t + \tau$ will not inform us about generators’ expectations at time t .

²⁰Two pieces of evidence support this assumption. First, as discussed in Section 2.1, generators subject to U.S. state standards largely did not change their announced retirement decisions in response to MATS’ announcement. Second, using our analysis data, we find that SO₂ emissions rates for generators complying with U.S. state standards were not significantly lower in 2016 relative to their enforcement years.

²¹We do observe adoption in years before standards’ announcements. Although we do not directly model the choice of adoption in these years, our data reflect each generator’s accurate adoption status at the start

We search over values of these structural parameters. For each candidate parameter vector, U.S. state, and belief year, we solve for the ABOE, which we use to calculate a likelihood. The likelihood for any generator and year is the probability of the equilibrium strategy given the observed action of continue, adopt, or exit, evaluated at the ABOE.

As discussed in Section 3.1, market evolution conditional on a generator’s actions is governed by three continuous states that evolve according to AR(1) processes. We initialize these processes to the values in the observed data. Since the fuel price ratio evolves exogenously to our model, this AR(1) regression remains constant throughout our solution process. We update the coal capacity and adoption share AR(1) processes by solving for the fixed points of generator Bellman equations, simulating data given optimizing choices, and rerunning the regressions that underlie these processes. We repeat this process until we reach a fixed point in both the regression coefficients and the Bellman equations. We calculate standard errors via a parametric bootstrap, which we discuss in more detail in On-line Appendix A3.

For generators subject to U.S. state enforcement, A and X are identified by the state-contingent rates at which generators choose each action given their expected profits. The scale parameter, σ , is in turn identified by the inclusion of annual profits. Identification of the MATS enforcement probabilities stems from the intuition underlying Figure 1: to the extent that generators subject to MATS delay exit and adoption relative to generators subject to U.S. state enforcement all else equal, our model will recover lower estimates of their enforcement probabilities. While we allow A to vary based on whether the generator is subject to MATS, identification is based on the exclusion restriction that exit costs are the same for all generators. Finally, we do not include an annual fixed cost of operation, since these are difficult to identify separately from exit costs (Collard-Wexler, 2013). Thus, our estimates of X will capture the present discounted value of any fixed costs of operation.

4.2 Annual Profit and Pollution Surfaces

Our structural estimation relies on state-contingent annual profits, $\Pi_j(\Omega)$, which in turn depend on revenues and costs. We estimate ramping and O&M costs—including the expected

of any year.

value of the idiosyncratic component of ramping costs— using the approach in Borrero et al. (2023), and then calculate Π_{jt} for every generator and year in the data. We allow ramping costs to vary by generator capacity bins. We then calculate each generator’s annual profits as:

$$\begin{aligned} \Pi_{jt} = \sum_{h=1}^H [q_{jh}p_h - q_{jh} \times heat_j \times f^C - \hat{r}_{j,\tilde{q}_h,q_h} \mathbb{1}\{\tilde{q}_{jh} < q_{jh}\} - \hat{o}m \times q_{jh} \\ - \hat{\sigma}^g Pr(q_{jh}|\tilde{q}_{jh}) \log(Pr(q_{jh}|\tilde{q}_{jh}))], \end{aligned} \quad (4)$$

where we indicate our estimated parameters with hats. The second line in (4) captures the expected value of the ramping cost unobservable, ε_{jhq}^g , conditional on choice q_{jh} . We calculate the probability of each action at each hour, $Pr(q_{jh}|\tilde{q}_{jh})$, from our cost estimation.

We do not observe $\Pi_j(\Omega)$ for every generator j and potential market state Ω . We therefore predict generators’ expected annual profits across their potential market states. We perform this prediction by calculating Π_{jt} for each generator-year in our sample and then regressing it on gas fuel price, relative coal capacity, their interaction, and generator j ’s fixed characteristics, K_j and $heat_j$. By assuming this functional form for the impact of these characteristics on expected profits, we can interpolate—and potentially extrapolate—for our estimation and counterfactuals to states that we do not observe in the data.²² This allows us to avoid directly modeling hourly generation decisions and evaluating how they might change in counterfactuals. In addition, since our predicted profits are a function of our state, which includes aggregate moments, we assume that each firm in a given state believes that it will get the predicted profits, in expectation. Together, these assumptions are required for identification.

We also want to understand how SO₂ pollution would vary under counterfactual policy environments. We focus on SO₂, since the value of MATS pollution reductions as calculated by the EPA (Environmental Protection Agency, 2011) is dominated by SO₂ and, consistent with our measure of MATS compliance, reductions in other pollutants will likely be approximately proportional to SO₂ reductions. For these simulations, we approximate a pollution

²²We have also investigated alternative specifications of this regression, such as a logged functional form, finding similar results.

surface with a similar regression to our profit regression, but where we use the log of annual pollution as the dependent variable.

5 Results and Counterfactuals

5.1 Cost, Profit, and Pollution Estimates

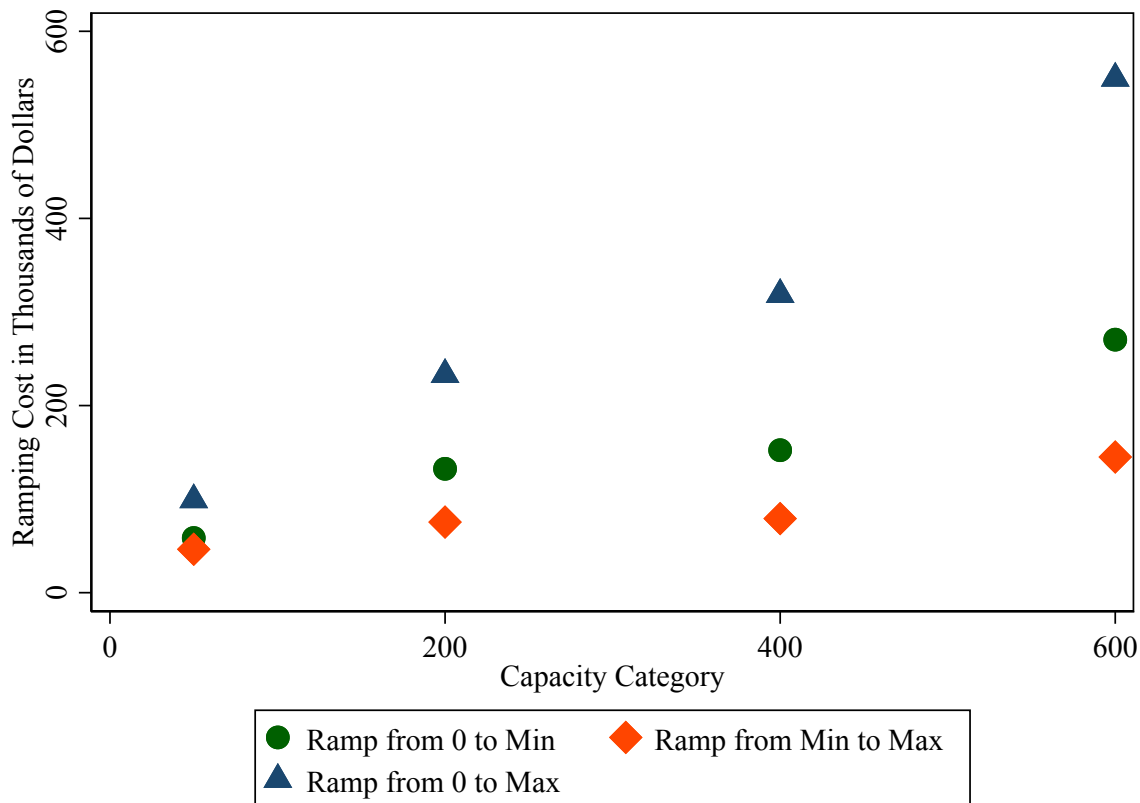
We estimate ramping and O&M costs following Borrero et al. (2023), separately for generators in 5 capacity bins (0–100MW, 100–300MW, 300–500MW, 500–700MW, and >700MW). Figure 2 reports the ratios of the ramping coefficients, $\frac{\hat{r}}{\hat{\sigma}_g}$, to the operating revenue coefficient, $\frac{\hat{1}}{\hat{\sigma}_g}$, for all ramping parameters and all but the highest capacity bin.²³ We find that ramping from minimum to maximum generation (represented by orange diamonds) is the least costly, ramping from off to minimum generation (green circles) is more costly, and ramping directly from off to maximum generation in one hour (blue triangles) is the most costly. The fact that ramping from zero to minimum is more costly than from minimum to maximum likely reflects startup costs, which are captured by our ramping cost estimates. Figure 2 further makes clear that ramping costs are strongly increasing in generator capacity. In fact, the largest generators have extremely high ramping costs, with ramping from off to minimum generation costing approximately \$1.05 million.

Our ramping cost estimates are consistent with the finding that starting up from off is particularly costly. Engineering estimates suggest that start-up costs for large coal plants may reach \$500,000 (Kumar et al., 2012). Engineering estimates conceptually differ from ours in that they are based on models of excess fuel, non-fuel inputs, wear-and-tear costs rather than revealed preferences, and hence may provide a lower bound through omitted categories. Estimates based on revealed preferences from Cullen (2014) similarly find extremely large start-up costs. Reguant (2014) finds start-up costs of €15-20k for a 150MW coal generator and approximately €30k for a 350MW coal generator in Spain.

Turning to operations and maintenance costs, we find that generators on average pay O&M costs of \$15.18 per MWh of generation, corresponding to just over \$3,000 for a 200MW

²³The highest capacity bin has ramping estimates that follow the same pattern as the smaller bins, but has a substantially larger scale. We do not show the results for this bin for scaling reasons.

Figure 2: Variation in Ramping Cost Estimates with Capacity



Note: Ramping costs are ratios of estimated ramping coefficient to operating revenue coefficient from separate regressions by capacity bin. Symbols are placed at the midpoints of the capacity bins.

generator operating at full capacity. This estimate is close to the EIA’s National Energy Modeling System estimates of approximately \$14/MWh, which is used in Linn and McCormack (2019).

Once we have recovered estimates of ramping and O&M costs, we calculate the generator profits for generator-years in our analysis data. The 5th percentile of calculated profits is -\$15 million and the 95th is \$56 million, with 65% of annual profits above zero. The sizable share of calculated profits that are below zero suggest that there may be substantial option value to remaining in operation or sizable exit costs. On-line Appendix Figure A1 presents a histogram of these calculated profits.

Turning to our estimated profit and pollution surfaces, Table 4 presents the results of our regressions of calculated profits and pollution on model states. The profit regression R^2

Table 4: Profit and Pollution Surface Regression Results

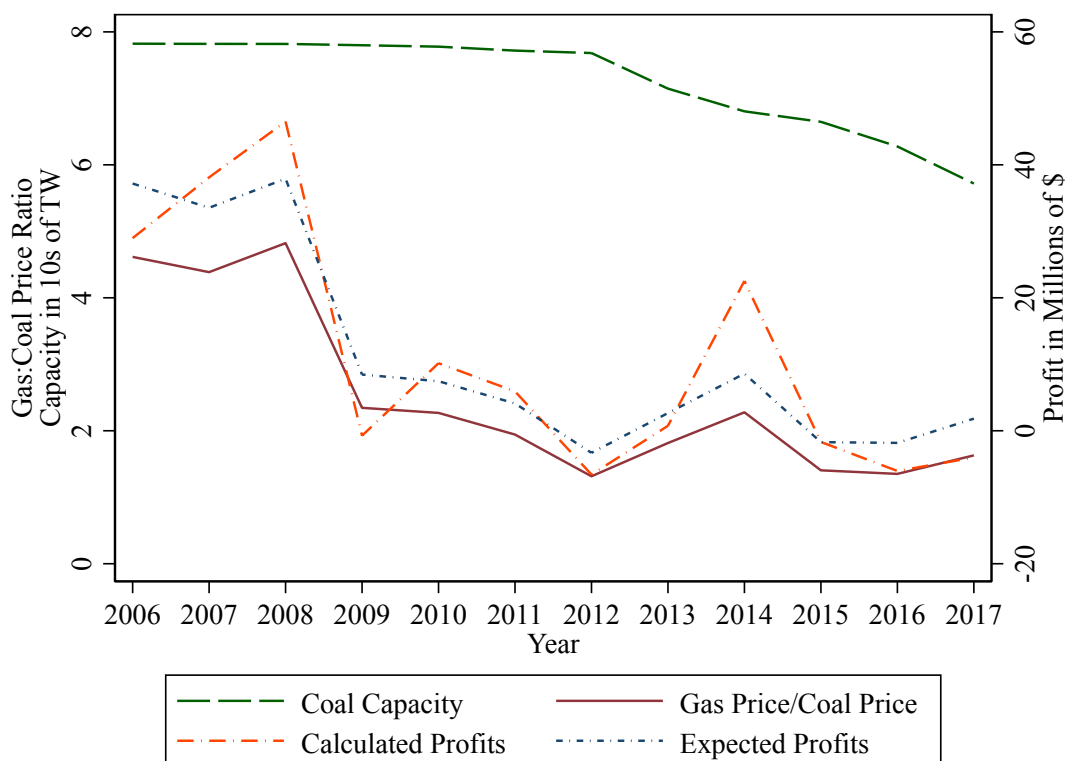
	Annual Profit (millions of \$)	Log Annual Pollution (lbs of SO ₂)
Compliant		-1.285*** (0.054)
U.S. State Coal Share	-5.505*** (1.701)	0.335*** (0.064)
Gas to Coal Fuel Price Ratio	9.371*** (0.819)	0.798*** (0.042)
Interaction of Coal Share and Price Ratio	2.681*** (0.773)	-0.182*** (0.064)
U.S. State Coal Fuel Price (\$)	16.231*** (1.230)	-0.728*** (0.182)
Heat Rate (MMBtu/MW)	-0.648*** (0.133)	1.278*** (0.244)
Capacity (MW)	0.052*** (0.004)	0.985*** (0.046)
Constant	-55.418*** (4.109)	7.847*** (0.809)
Observations	3035	2819
R^2	0.4778	0.9938

Note: Regression of calculated profits and log pollution from observed data on dynamic model states. The pollution regression also uses logs of the independent variables other than compliance and excludes generator-years with zero pollution. Standard errors are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

is relatively high at 0.4778, and most of the parameters have the expected sign. Profits are negatively correlated with U.S. state coal capacity (which includes both IPP and non-IPP coal capacity). They are positively correlated with the fuel price ratio, and this correlation is stronger when the U.S. state coal capacity is higher. Generators with higher heat rates—that burn coal less efficiently—have lower profits. However, one coefficient with a counterintuitive sign is the U.S. state coal fuel price. U.S. states with high coal fuel prices have higher coal generator profits, which likely captures other attributes of these U.S. states. Because this variable is fixed over time for any generator and does not vary in counterfactuals, this counterintuitive sign is not concerning by itself.

Figure 3 shows that expected mean profits—as predicted with our regression of calculated profits on model states—follow calculated actual profits well over time. The solid red line shows the gas to coal price ratio, which indicates that when gas prices dropped, coal generator profits also dropped substantially. The blue long-dashed line shows that coal capacity started falling approximately two years after the fall in gas prices, with gradual exit contin-

Figure 3: Variation in Coal Capacity, Mean Profits, and Fuel Prices Over Time



Note: Authors' calculations based on the analysis sample of IPP coal generators. Calculated profits use our estimated ramping and O&M costs and observed generation choices. Expected mean profits are predicted for generator-years in our sample using estimated profit surface regression parameters.

uing throughout the rest of the sample. Our structural model will attribute this pattern to idiosyncratic shocks to the value of exit.

The pollution surface regression, presented in column 2 of Table 4, has three differences relative to our profit surface regression. First, we include whether the generator is compliant with air toxics standards in the observation year, since the goal of compliance is to reduce pollution. Second, we log both pollution and the dependent variables other than compliance. Finally, we run this regression only for generator-years with non-zero pollution. This regression has a high R^2 of 0.9938, and all of the coefficients have the expected sign. Specifically, compliance and coal fuel prices are associated with lower pollution. Higher coal capacity, generator heat rate, and generator capacity are all associated with higher pollution. Importantly, we observe a positive baseline coefficient on gas prices relative to coal prices. This

is consistent with our priors since we would expect coal generators to run more when gas prices are higher. The fact that this effect is decreasing as the coal capacity increases is also consistent with the idea that this increased generation would be spread over more coal generators in this case.

5.2 Results of Adoption and Exit Model

Table 5: Structural Parameter Results

	Base Specification	Same Adoption Cost State vs. MATS
Predicted Enforcement Probabilities:		
Probability 2012	1.000*** (0.061)	0.999*** (0.080)
Probability 2013	0.699*** (0.120)	0.525*** (0.159)
Probability 2014	0.433*** (0.109)	0.306** (0.139)
Probability 2015	0.999*** (0.107)	0.997*** (0.103)
Generator Costs:		
Adoption Cost (million \$)	150.9** (75.1)	413.9*** (41.8)
Extra MATS Adoption Cost (million \$)	398.7*** (72.1)	—
Exit Scrap Value (million \$)	-196.4*** (37.4)	-196.8*** (42.6)
1/ σ (million \$)	63.6*** (5.7)	63.4*** (6.5)
Simulated Log Likelihood	-628.34	-637.88

Note: Structural parameter estimates from nested-fixed point estimation. Standard errors calculated via a parametric bootstrap are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 5 presents our structural results for two potential specifications. The first column presents our base specification, where generators subject to MATS are allowed to have different adoption costs than generators subject to state standards. We find that in 2012 generators perceived that MATS was essentially certain to be enforced. This probability drops to 70% in 2013 and to 43% in 2014 when the U.S Supreme Court agreed to hear arguments in the MATS case (*Michigan v. EPA*). By 2015, however, generators realized that MATS was again very likely to be enforced in 2016, with the probability again rising to nearly one.

As shown previously in Table 2, generators subject to MATS disproportionately exited or adopted in both 2012 and 2015 compared to 2013 and 2014. This is reflected in our perceived probability estimates. The high 2012 estimate may reflect the fact that we assume that the MATS enforcement probability was 0 prior to 2012, but generators likely had expectations of a positive enforcement probability. The large number of exits and adoptions in 2012 may be the result of generators anticipating an air toxics policy before 2012, but not knowing the details of the policy until 2012. Indeed, we see the number of generators subject to MATS that exit dip from 15 in 2010 to 6 in 2011, before rising to 23 in 2012. We therefore interpret our 2012 probability estimate with caution. The remaining estimates however, suggest that generators' actions revealed substantial uncertainty surrounding 2016 MATS enforcement.

Beyond generators' perceived probabilities of enforcement, we find that adopting air toxics abatement technology costs generators subject to state standards \$151 million. Adoption costs generators subject to MATS an additional \$399 million, consistent with the evidence that compliance with MATS may be more expensive due to more stringent standards and enforcement. These are substantial costs given that the mean generator profit during our sample period is approximately \$13 million. We further estimate that generators need to pay \$196 million in exit costs (negative scrap value) to shut down. One study has found decommissioning costs for coal generators of up to \$466,000 per MW, which is similar but somewhat smaller than our estimates (Raimi, 2017). Our estimates are therefore consistent with substantial site remediation costs and may also include avoided fixed costs (Collard-Wexler, 2013).

The second column of Table 5 requires abatement costs to be identical across generators subject to state standards and MATS. We find similar results, with adoption costs between U.S. state and MATS adoption costs from the first column. Though the probability estimates in 2013 and 2014 are smaller than in our baseline, the temporal patterns are quite similar, starting at approximately one in 2012, falling in 2013 and again in 2014, and then recovering to nearly one.²⁴

²⁴We also ran a model where adoption and exit costs are proportional to the generator's capacity. This model fits the data less well, but recovered broadly similar estimates of the probability of enforcement over time.

5.3 Counterfactual Results

Table 6: Counterfactual Results

	Data	Estimated Model	Same Mean Prob.: 0.7827	Uncertainty Resolved in 2012 w/ Mean Prob.	No Exit Cost
Adoption Costs (Bill. \$)		7.30	6.99	6.53	5.10
Exit Costs (Bill. \$)		19.24	19.15	18.74	0.00
Total Profits (Bill. \$)		46.74	47.73	48.66	69.90
Pollution (Mill. lbs. SO ₂)		867.52	881.30	946.50	740.91
Number of Generators:					
2012	191	191.0	191.0	191.0	191.0
2013	168	175.9	176.5	176.8	169.4
2014	149	163.0	163.1	163.6	151.6
2015	142	152.5	150.7	151.4	137.1
2016	121	129.5	131.0	135.0	111.0
Count Adopting	14	14.5	13.7	12.85	10.1

Note: Column 1 reports observed exit and adoption decisions in the data. Column 2 reports predicted outcomes at model estimates. Column 3 replaces the estimated probabilities of 2016 MATS enforcement with the mean estimated probability across years. Column 4 calculates the expect outcomes with uncertainty completely resolved in 2012, with the mean estimated probability across years. Column 5 sets exit costs to 0. First four rows of results report the total discounted profits or costs from 2012 through 2041.

We use the structural parameter estimates from our base specification to simulate a series of counterfactuals. In each of the counterfactuals, we re-solve for the fixed point between generators' expectations of the evolution of the market state and their adoption and exit decisions. Table 6 presents counterfactual discounted costs, profits, pollution, and exit and adoption outcomes for generators subject to MATS enforcement over the 30 years from 2012-2041.

The first two columns of Table 6 present the observed exit and adoption outcomes for generators subject to MATS enforcement and the predicted outcomes using our model estimates. The model generally reproduces the data well. In particular it predicts that 14.2

generators would adopt abatement technology while 14 adopted in practice. We slightly underpredict the exit rate, which may be due to the estimation sample being different from our counterfactual sample. Turning to the other outcomes in column 2, our model predicts that generators will pay \$7.3 billion in abatement technology costs, pay \$19.2 billion in exit costs, and earn \$46.7 billion in profits, discounted and summed over the 30 year period. Our estimates of generator abatement costs are similar to the EPA’s ex ante estimates of compliance costs, which were \$9.6 billion (Environmental Protection Agency, 2011). Generators will also produce 868 million discounted pounds of SO₂ pollution over this period.

The third column of Table 6 takes the average estimated probability (0.7827) from our model and applies it evenly in all years, including in the enforcement year, so that MATS is only actually enforced in 78.27% of simulation draws. This case is roughly similar to the baseline in column 2—as intended—but it provides a simplified ex post uncertainty resolution case that we use to compare to ex ante uncertainty resolution.

Specifically, column 4 of Table 6 similarly assumes a 78.27% probability of MATS enforcement, but this is decided randomly at the moment that MATS is announced in 2012. This means that the *level* of the standard is identical to the counterfactual in column 3 in expectation, but that there is no uncertainty after announcement over whether MATS will be enforced. Thus the comparison between columns 3 and 4 provides a clear description of the costs and benefits of resolving policy uncertainty earlier.

With uncertainty resolved at announcement, generator profits are \$930 million higher than when the uncertainty is resolved in 2016. Of this, \$460 million comes from lower adoption costs, and \$410 million comes from lower exit costs. The remainder of the savings accrue from generators timing their adoption and exit decisions to better take advantage of time-varying costs such as maintenance downtime that affect adoption and exit costs via the $\vec{\epsilon}_{jt}$ term.

While resolving uncertainty in 2012 increases profits by 1.9%, it also increases the number of generators remaining in 2016 by 3.1%, and increases pollution by 7.4%. We use the EPA’s regulatory impact analysis total pollution benefits range and SO₂ pollution reduction from MATS to calculate the range of benefits from reducing SO₂ and related pollutants (Environmental Protection Agency, 2011). This calculation implies that a one pound reduction in

SO₂ is worth between \$12.41 and \$33.83. This then implies that resolving uncertainty ex ante would result in damages of \$809 million to \$2.206 billion via the 65 million pound increase in pollution. This pollution increase is therefore critical to understanding the welfare impacts of removing policy uncertainty.

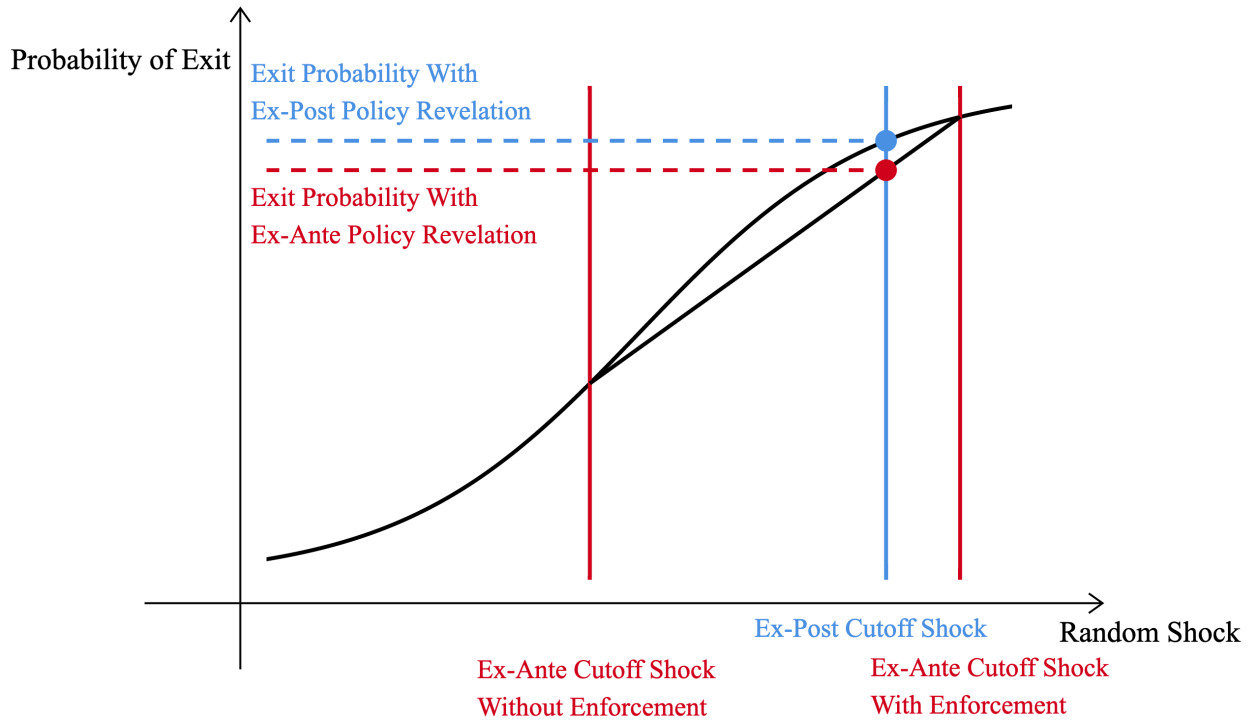
Finally, the last column of Table 6 keeps policy uncertainty the same as in our estimated model, but assumes that exit costs are fully subsidized. While this policy transfers exit costs to the government, it also reduces adoption costs, since some generators choose to exit rather than adopt pollution abatement technologies. This counterfactual results in 15.1% fewer generators in the market in 2016 and a drop in pollution of 14.6% relative to the baseline in the second column. In tandem, generators' discounted profits rise dramatically, by 49.6%. This column illustrates how achieving substantial reductions in coal capacity may be very costly to the government.

5.4 Understanding the Impacts of Early Uncertainty Resolution

In order to better understand these counterfactual results, Figure 4 focuses on a conceptual cumulative distribution function of a generator's random shocks between 2012 and 2016 that affect its exit decision. We focus on exit rather than adoption because exit was the primary means of compliance for generators subject to MATS (as shown in Table 2).

The blue vertical line in Figure 4 represents the minimum shock that would lead the generator to exit the market given the current policy with uncertainty resolution after the generator makes its exit decision (ex post resolution). The blue horizontal dashed line indicates the corresponding probability of exit. If, however, uncertainty was resolved at the time of announcement (ex ante), then the generator would receive the shock after uncertainty resolution. This would lead to different cutoff shocks depending on the enforcement decision, which we represent with the two vertical red lines. Figure 4 places the cutoff shock with ex post uncertainty resolution as 78.27% of the way between the two ex ante cutoff shocks to correspond to our estimated likelihood of enforcement. We mark the expected exit probability with ex ante uncertainty resolution with the horizontal red dashed line. This probability is the weighted average of the exit probabilities from the two red vertical lines.

Figure 4: Conceptual Exit Cumulative Distribution Function



Note: Conceptual cumulative distribution function for a single generator’s exit choice in a given year.

Our counterfactuals show that the exit probability with ex ante uncertainty resolution is lower than with ex post uncertainty resolution. As we can see in Figure 4, this occurs when the cutoff shock with ex post uncertainty resolution is on the concave portion of the CDF: with a single-peaked distribution of exit shocks, the concave portion coincides with a high probability of generator exit. This is consistent with the very low natural gas prices during this period, which likely meant that some coal generators were very close to the exit margin. Alternatively, if generators had seen substantially greater value to remaining in the market, then resolving uncertainty ex ante would *increase* the probability of exit. This is the continuous version of the oil drilling example in Section 1: if drilling costs \$55 per barrel, the expected future oil price is on the concave portion of its drilling CDF, implying that ex ante uncertainty resolution would decrease the expected drilling probability, while the opposite is true when drilling costs \$65 per barrel.

In general, earlier uncertainty resolution will increase generator profits by allowing them to better respond to random shocks. With a single-peaked distribution of exit shocks (i.e., S-shaped exit shock CDF), it will also attenuate extreme outcomes in expectation by decreasing high exit probabilities (which lie on the concave portion of the adoption CDF) and raising low exit probabilities (which lie on the convex portion). This again generalizes our oil drilling example: in both cases earlier uncertainty resolution attenuates irreversible decisions in expectation. This result contrasts with the real options literature’s result that reducing the variance of future profits may increase irreversible investment.

6 Conclusions

This paper investigates the policy uncertainty surrounding a major U.S. environmental regulation, the Mercury and Air Toxics Standard. Because abatement technology adoption and exit decisions are costly and irreversible, we estimate a dynamic oligopoly model of generator abatement technology adoption and exit behavior. Importantly, we estimate generators’ perceptions of the probability that MATS would be enforced in 2016 along with technology adoption and exit costs. Our model is identified by the difference in abatement technology adoption and exit decisions between coal electricity generators facing MATS and similar generators facing U.S. state air toxics standards. We develop a new approximate equilibrium concept that we call Approximate Belief Oligopoly Equilibrium (ABOE) that allows us to capture equilibrium effects and avoid the curse of dimensionality that would stem from keeping track of the actions of the many generators within a market.

We find that there was substantial uncertainty over whether MATS would be enforced, with generators’ perceived enforcement probability falling to 43% in 2014. The average expected enforcement probability was approximately 80% over the 2012-2015 period. In order to understand the impact of the timing of uncertainty resolution, we compare the observed pattern of uncertainty to a counterfactual environment where there is the same approximately 80% chance at the moment of MATS announcement that MATS would be enforced in 2016, but this uncertainty is resolved instantly for all generators in 2012 with full commitment. We find that resolving uncertainty earlier decreases the cost of complying with

MATS by \$930 million, but increases pollution damages by \$809 million to \$2.206 billion as it causes more generators to remain in the market in expectation.

Overall, these results highlight a difference with the real options literature, which shows that decreasing uncertainty leads to more irreversible actions (corresponding to exit and adoption for MATS). In contrast, resolving uncertainty earlier attenuates extreme outcomes in expectation. When generators are close to the exit margin, this will decrease exit and increase pollution. However, if generators were unlikely to exit, these factors would be reversed and resolving policy uncertainty earlier might have decreased pollution.

Our results highlight the fact that regulatory discretion, while potentially allowing infrequent legislation to adapt to changing circumstances and technologies, can also increase policy uncertainty. Recently, the 2022 Supreme Court case *West Virginia et al. v. Environmental Protection Agency* added to this uncertainty by calling into question agencies' roles in formulating regulations. Our results imply that movements toward less regulatory discretion will attenuate extreme outcomes in expectation, but the exact impact in any given setting requires understanding the incentives that actors face.

References

- Abito, Jose Miguel, Christopher R Knittel, Konstantinos Metaxoglou, and André Trindade**, “The Role of Output Reallocation and Investment in Coordinating Environmental Markets,” *International Journal of Industrial Organization*, 2022, *83*, 102843.
- Baker, Scott R., Nicholas Bloom, and Steven J. Davis**, “Measuring Economic Policy Uncertainty,” *The Quarterly Journal of Economics*, 2016, *131* (4), 1593–1636.
- Bohren, J Aislinn and Daniel N Hauser**, “Learning with Heterogeneous Misspecified Models: Characterization and Robustness,” *Econometrica*, 2021, *89* (6), 3025–3077.
- Borrero, Miguel, Gautam Gowrisankaran, and Ashley Langer**, “Ramping Costs and Coal Generator Exit,” *Working Paper*, 2023.
- Collard-Wexler, Allan**, “Demand Fluctuations in the Ready-Mix Concrete Industry,” *Econometrica*, 2013, *81* (3), 1003–1037.
- Congressional Research Service**, “Mercury Emissions from Electric Power Plants: States Are Setting Stricter Limits,” 2007. Available at https://www.everycrsreport.com/files/20070222_RL33535_9490786af904352a3618fa139135a3471dab06fb.pdf.
- Corbae, Dean and Pablo D’Erasmus**, “Capital Buffers in a Quantitative Model of Banking Industry Dynamics,” *Econometrica*, 2021, *89* (6), 2975–3023.
- Cullen, Joseph A.**, “Dynamic Response to Environmental Regulation in the Electricity Industry,” 2014.
- Dixit, Avinash K. and Robert S. Pindyck**, *Investment Under Uncertainty*, Princeton University Press, 1994.
- Dorsey, Jackson**, “Waiting for the Courts: Effects of Policy Uncertainty on Pollution and Investment,” *Environmental and Resource Economics*, 2019, *74* (4), 1453–1496.
- Elliott, Jonathan T**, “Investment, Emissions, and Reliability in Electricity Markets,” 2022.

- Environmental Protection Agency**, “Regulatory Impact Analysis for the Final Mercury and Air Toxics Standards,” 2011. Available at <https://www3.epa.gov/ttnecas1/regdata/RIAs/matsriafinal.pdf>.
- Esponda, Ignacio and Demian Pouzo**, “Berk–Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models,” *Econometrica*, 2016, *84* (3), 1093–1130.
- Fowle, Meredith**, “Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement,” *American Economic Review*, 2010, *100* (3), 837–69.
- Gerarden, Todd**, “Demanding Innovation: The Impact of Consumer Subsidies on Solar Panel Production Costs,” *Management Science*, 2022. forthcoming.
- Gowrisankaran, Gautam, Ashley Langer, and Mar Reguant**, “Energy Transitions in Regulated Markets,” *Working Paper*, 2023.
- Halloran, Liz**, “State Takes Action on Mercury,” *Hartford Courant*, 2003.
- Handley, Kyle and J. Frank Li**, “Measuring the Effects of Firm Uncertainty on Economic Activity: New Evidence from One Million Documents,” *National Bureau of Economic Research Working Paper*, 2020.
- **and Nuno Limão**, “Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States,” *American Economic Review*, 2017, *107* (9), 2731–83.
- Hawthorne, Michael and Tribune staff reporter**, “State Moves Against Mercury,” *Chicago Tribune*, 2006.
- Herrnstadt, Evan M, Ryan Kellogg, and Eric Lewis**, “The Economics of Time-Limited Development Options: The Case of Oil and Gas Leases,” Technical Report, National Bureau of Economic Research 2020.
- Hudson, J. Brian**, “Out of the Rabbit Hole: The DC Circuit Brings the EPA Back from Wonderland in *New Jersey v. EPA*,” *Vill. Envtl. LJ*, 2010, *21*, 45.

- Ifrach, Bar and Gabriel Y. Weintraub**, “A Framework for Dynamic Oligopoly in Concentrated Industries,” *The Review of Economic Studies*, 2017, 84 (3), 1106–1150.
- Jeon, Jihye**, “Learning and Investment Under Demand Uncertainty in Container Shipping,” *The RAND Journal of Economics*, 2022, p. forthcoming.
- Johnston, Sarah and Dominic Parker**, “Causes and Consequences of Policy Uncertainty: Evidence from McGirt vs. Oklahoma,” *Oklahoma (May 31, 2022)*, 2022.
- Kellogg, Ryan**, “The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling,” *American Economic Review*, 2014, 104 (6), 1698–1734.
- Kumar, N., P. Besuner, S. Lefton, D. Agan, and D. Hilleman**, “Power Plant Cycling Costs,” Technical Report, National Renewable Energy Laboratory 2012.
- Langer, Ashley and Derek Lemoine**, “What Were the Odds? Estimating the Market’s Probability of Uncertain Events,” *National Bureau of Economic Research Working Paper*, 2020.
- Linn, Joshua and Kristen McCormack**, “The Roles of Energy Markets and Environmental Regulation in Reducing Coal-Fired Plant Profits and Electricity Sector Emissions,” *The RAND Journal of Economics*, 2019, 50 (4), 733–767.
- Pelton, Tom**, “State Loosens Pollution Rule,” *Baltimore Sun*, 2006.
- Raimi, Daniel**, “Decommissioning US Power Plants: Decisions, Costs, and Key Issues,” Technical Report, Resources for the Future 2017.
- Reguant, Mar**, “Complementary Bidding Mechanisms and Startup Costs in Electricity Markets,” *The Review of Economic Studies*, 2014, 81 (4), 1708–1742.
- Riordan, Michael H.**, “Regulation and Preemptive Technology Adoption,” *The RAND Journal of Economics*, 1992, pp. 334–349.
- Schlenker, Wolfram and Charles A. Taylor**, “Market Expectations of a Warming Climate,” *Journal of Financial Economics*, 2021, 142 (2), 627–640.

- Schmidt-Dengler, Philipp**, “The Timing of New Technology Adoption: The Case of MRI,” *Manuscript, London School of Economics*, 2006.
- Scott, Jonathan B.**, “Positive Spillovers from Infrastructure Investment: How Pipeline Expansions Encourage Fuel Switching,” *The Review of Economics and Statistics*, 2021, pp. 1–43.
- Teisberg, Elizabeth Olmsted**, “Capital Investment Strategies Under Uncertain Regulation,” *The RAND Journal of Economics*, 1993, pp. 591–604.
- United Press International**, “Mass.: Toughest Mercury Emissions Limits,” *UPI*, 2004.
- Vreugdenhil, Nicholas**, “Booms, Busts, and Mismatch in Capital Markets: Evidence from the Offshore Oil and Gas Industry,” 2020.
- Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy**, “Markov Perfect Industry Dynamics with Many Firms,” *Econometrica*, 2008, 76 (6), 1375–1411.

On-Line Appendix

A1 Evidence on U.S. State Air Toxics Standards

In the years before the EPA released the final MATS rule, a number of U.S. states drafted and passed their own rules and legislation in order to tackle mercury emissions. This section first documents how we determine U.S. state announcement and enforcement dates. Second, we show supporting evidence for our assumption that U.S. state standards were certain once announced. Finally, we explain why we believe that compliance with U.S. state-level standards may have been less onerous than compliance with MATS.²⁵

A1.1 Implementation of U.S. State Standards

Connecticut's Public Act 03-72 was signed into law by Governor John Rowland in June 2003, after gathering unanimous approval in both the state Senate and House (Jones, 2002). The law specified an enforcement date of July 1, 2008 (Connecticut General Assembly, 2003).

Delaware's Department of Natural Resources and Environmental Control designed the "State Plan for the Control of Mercury Emission from EGUs." This set of regulations came into effect on December 11, 2006 (Hughes, 2006). The initial enforcement date was May 1, 2009 (Office of the Registrar of Regulations, 2006).

Illinois' Mercury Rule was proposed on January 5, 2006 by Governor Rod Blagojevich (Hawthorne and Tribune staff reporter, 2006). The regulation was formally adopted by the Illinois Pollution Control Board on December 21, 2006 (Gunn, 2006). The enforcement date was July 1, 2009 (Illinois Pollution Control Board, 2006).

In Maryland, a bill addressing mercury emissions was initially proposed in December 2005 (McIntire, 2005). After significant back-and-forth between Maryland's Democratic legislature and its Republican administration, the Healthy Air Act was passed into law on April 6, 2006 (Zibel, 2006). The enforcement date was January 1, 2010 (Maryland Department of the

²⁵We thank Jacob Felton, Amy Mann, Rory Davis, Alison Ray, Joanne Morin, Bruce Monson, and Anne Jackson from state agencies in Connecticut, Delaware, Illinois, Maryland, Massachusetts, and Minnesota for assistance in understanding U.S. state regulations.

Environment, n.d.).

The Commonwealth of Massachusetts adopted mercury standards on May 26, 2004 (United Press International, 2004). The enforcement date was January 1, 2008 (The Massachusetts Department of Environmental Protection, 2006).

Minnesota’s Mercury Emissions Reduction Act was signed into law on May 11, 2006. Firms had to comply with the new mercury emissions limitations by December 31, 2010 (Minnesota State Legislature, 2006).

In Wisconsin, the Department of Natural Resources proposed new mercury emission regulations pertaining to coal-fired utilities in September 2004 (Wisconsin Department of Natural Resources, n.d.). It later revisited these and updated them in 2008. The regulations were first enforced on April 16, 2016 (Legislative Reference Bureau Wisconsin, 2006).

As noted in Table A1, we specify the enforcement year as the first full year following the enforcement date. The one exception to this is Wisconsin, where we used the MATS enforcement year since MATS would have superseded the U.S. state policy.

In some cases, we assume that generators were subject to MATS instead of U.S. state enforcement, even though we found evidence of U.S. state policies. Florida developed a cap-and-trade policy based on CAMR rather than a technology standard (Congressional Research Service, 2007). Michigan’s regulation was a backstop to MATS and would only come into force if MATS was invalidated (Michigan Department of Environment, Great Lakes, and Energy, n.d.). New Hampshire’s policy was essentially voluntary (Congressional Research Service, 2007). North Carolina’s enforcement date was in 2018, substantially after MATS (Congressional Research Service, 2007). Finally, as discussed in Section 2.1, Pennsylvania’s standard was overturned by the state’s Supreme Court (Mustian and Demase, 2009).

A1.2 Enforcement Probabilities for U.S. State Standards

We assume that generators believed that U.S. state standards, once announced, would be enforced with certainty. In three U.S. states (Connecticut, Maryland, and Minnesota), the standard was passed as legislation. This reduces the potential for legal challenges as states have authority to pass legislation that regulates pollutants for sources located within their

borders. In other U.S. states, the standards were implemented via rule-making. In both cases, standards were generally developed in conjunction with industry groups. For instance, in Connecticut, the Senate chair of the environment committee commented that “[i]ndustry took a look at that battle and realized it was in their interest to collaborate and compromise” (Halloran, 2003). In Illinois, regulators included alternative emission standards at the behest of local industry, which limited emissions from power plants as a whole rather than from individual generating units (Illinois Pollution Control Board, 2006).

Overall, the U.S. state standard development process substantially decreased the potential for legal challenges. In particular, for the seven U.S. states that we define to have air toxics standards, we were able to uncover only one legal challenge: in Wisconsin, an industry group sued the Wisconsin Department of Natural Resources in May, 2008 (Pioneer Press, 2008). This challenge did not appear to be well-grounded since it was quickly dismissed in June, 2008 (Bauer, 2008)).

Despite the fact that U.S. state standards were essentially certain once passed, there was uncertainty surrounding their initial announcements. Maryland’s Healthy Air Act provides perhaps the most dramatic example of this uncertainty. Republican Governor Robert Erlich and the Democrat-controlled legislature drafted competing bills to address air pollution (Pelton, 2006). The eventual bill that the legislature passed was closer to the Democrats’ version. Erlich’s staff allegedly locked his office door at 4:30 on a Friday afternoon at the end of the legislative session to avoid him being presented with a set of bills that included the Healthy Air Act. Legislative aides slid the bill under the governor’s door. Maryland’s Attorney General determined that such action constituted presentment (Marimow and Mosk, 2006), and Erlich ultimately signed the bill (Zibel, 2006). Further illustrating that there was uncertainty surrounding the initial announcements, some U.S. states such as Georgia proposed mercury standards that were never put in place (Cash, 2006). We do not model this source of uncertainty. This may be an additional reason why we find that generators subject to federal standards perceived the probability of MATS enforcement to be 1 in 2012.

A1.3 Compliance Costs for U.S. State Standards

While U.S. state standards regulated mercury similarly to MATS, there are at least three differences between these standards that are important for our analysis. All three of these differences implied that it would be less costly for generators to comply with U.S. state standards than MATS.

First, most U.S. state standards covered substantially fewer air toxics than MATS. Specifically, MATS regulated both mercury and air toxics, as required by *Sierra Club v. EPA*. In contrast, none of the U.S. state laws or regulations referenced reductions in air toxics other than mercury.

Second, U.S. state standards often had higher allowable emissions limits than MATS and specified alternative compliance approaches. Specifically, MATS specified a mercury emissions limit of between 0.0002 and 0.0004 lbs/GWh, depending on the type of coal (EPA, 2012). Connecticut specified a limit of 0.6 lbs/TBtu or a 90% mercury removal rate (Connecticut General Assembly, 2003). At the mean heat rate for generators subject to U.S. state regulation in Table 1, 0.6 lbs/TBtu corresponds to 0.00063 lbs/GWh. Delaware specified a limit of 1 lb/TBtu or an 80% reduction by 2010 (Office of the Registrar of Regulations, 2006). Illinois specified a limit of 0.0080 lbs/GWh or 90% reduction (Illinois Pollution Control Board, 2006). Maryland specified 80% reduction by 2010 (Maryland General Assembly, 2005). Massachusetts specified a limit of 0.0075 lbs/GWh or 85% reduction by 2012 (The Massachusetts Department of Environmental Protection, 2006). Minnesota specified a 70% reduction (Minnesota State Legislature, 2006). Wisconsin specified the minimum of a limit of 0.0080 lbs/GWh and a 90% reduction (Legislative Reference Bureau Wisconsin, 2006).

Finally, enforcement of U.S. state standards was less rigorous in some cases. For instance, in Connecticut, the regulator could deem generators that installed appropriate mercury reduction technologies to be compliant even if they did not comply with the emissions limits (Connecticut General Assembly, 2003). In Wisconsin, a generator could argue for an extension if compliance would lead to a disruption of electricity supply (Legislative Reference Bureau Wisconsin, 2006).

A2 Data Construction

We use our data to construct a number of variables, focusing first on the adoption of technologies that lead to compliance with air toxic standards. Although compliance is a central variable in our analysis, adoption of technologies that lead to compliance is not directly reported in our data.

Many of the U.S. states that implemented air toxics standards early on specified that they would determine compliance by using a CEMS to measure mercury emissions. However, this technology was ultimately not reliable enough to use. States ended up measuring compliance with a combination of technology reporting, periodic stack tests, and other emissions reporting. MATS—which was implemented after state air toxics standards—did not attempt to measure air toxics compliance through a mercury CEMS.

Because the most cost-effective technologies that abate both mercury and other air toxics also reduce SO₂, one important way that the EPA determines MATS compliance is via SO₂ emissions rates. In particular, the MATS final rule (77 FR 9304) specifies that generators can comply with MATS by having SO₂ emissions rates below 0.2 lbs/MMBtu.

For most generators, we observed a large decline in SO₂ emissions rates in a particular year before air toxics enforcement. For generators subject to MATS, these declines frequently reduced annual average emissions rates below 0.2 lbs/MMBtu, although in a number of cases the post-decline emissions rates were between 0.2 and 0.4, with some variation across years. In contrast, pre-decline rates were typically well above 1. For this reason we define a generator as having adopted abatement technology in the first year when (i) its 3-year forward moving average SO₂ emissions rate falls below 0.4, or (ii) its annual emissions rate falls by 40% or more.²⁶

We use a similar method to determine compliance with U.S. state air toxics standards. The only difference is that we found that the post-decline emissions rates were typically below 0.7 but often above 0.4 lbs/MMBtu. Thus, we used a cutoff of 0.7 in our definition of compliance with state air toxics standards. This is consistent with the evidence in Section A1.3

²⁶In a small number of cases, we observe generators operating past the MATS enforcement date that did not meet this definition. We define these generators as having adopted abatement technology before our sample begins, to effectively remove their adoption choices from our data.

that state air toxics standards may be less stringent than MATS.

Exit decisions are also central to our model because generators may respond to air toxics standards by exiting the market. For our generator-year data set, we define a generator to have exited after the last year in which we observe it generating with coal in the CEMS data.²⁷ For our generator-hour data set, we define a generator to have exited after the last hour in which we observe it generating, unless there are fewer than 200 hours with zero generation at the end of its appearance in the CEMS data, in which case we simply use the end of its CEMS appearance as the exit hour.

Beyond adoption and exit, we also need to define generator fixed characteristics, specifically, minimum and maximum generation levels conditional on generating, capacity, and heat rate. We define a generator’s maximum generation level as the 95th percentile of its observed hourly generation conditional on operating in the CEMS data. We also use this as the generator’s capacity.²⁸ We define a generator’s minimum generation level as its modal generation level between the 5th and 60th percentile of capacity. We then bin hours with positive generation into minimum and maximum generation levels based on whichever level is nearer.

Finally, we calculate the heat rate of each generator at each hour using its heat input divided by its electricity production. Our analysis uses a time-invariant measure of the heat rate for each generator. Because generators operate most efficiently when generating near full capacity, we define each generator’s time-invariant heat rate as the mean hourly heat rate across hours in the maximum generation bin.

A3 Equilibrium Computation and Estimation

We estimate our model and conduct counterfactuals by solving for equilibria across candidate parameter vectors, U.S. states, and belief years. We recover the equilibrium as the fixed point of an algorithm that iterates between solving individual optimization decisions and estimating

²⁷Thus conversions from coal to natural gas—as analyzed by Scott (2021) in response to MATS—will appear as exits in our data.

²⁸We choose the 95th percentile because while generators can generate above listed capacity, this extra generation is extremely costly in the long-run.

market evolution regressions that are generated by these decisions.

This appendix begins by detailing our model’s state space and explaining generators’ dynamic optimization. We then specify market evolution regressions and outline our fixed point solution algorithm. Finally, we describe our bootstrap process for calculating standard errors and outline our counterfactual calculations.

A3.1 State Space

Recall that the Approximate Belief Oligopoly Equilibrium (ABOE) of our model specifies three time-varying aggregate market states which together form Ω —coal capacity relative to the 95th percentile of hourly load, the share of IPP coal capacity that has adopted abatement technology, and the ratio of natural gas to coal fuel price. The state for a generator also includes its time-varying characteristics—years to enforcement, belief year, and technology adoption; and non-time-varying characteristics—capacity, heat rate, and fixed market characteristics. We discretize the three continuous state variables into 1000 bins, with 10 grid points for each state variable. We choose these grid points to be evenly spaced between 0 and maximum levels that depend on the variable and U.S. state. For the fuel price ratio, this maximum is 120% of the maximum value observed in the data. For coal capacity, this maximum is the higher of 0.1 and 120% of the maximum value observed in the data. For the adoption share, this maximum is 1.

Units subject to MATS enforcement have discrete states $\{\tau, \tilde{\tau}, Tech\}$ based on the belief year $\tau \in \{0, 1, 2, 3, 4\}$, years to enforcement $\tilde{\tau} \in \{0, 1, 2, 3, 4\}$, and the unit’s technological adoption status $Tech \in \{0, 1\}$. There are 21 total states. There is one state for $\tau = 0$ as $Tech$ is not relevant in this case. There are two states for $\tau = 1$ where $\tilde{\tau} = 1$ and $Tech \in \{0, 1\}$, four states for $\tau = 2$ where $\tilde{\tau} \in \{1, 2\}$ and $Tech \in \{0, 1\}$, six states for $\tau = 3$ where $\tilde{\tau} \in \{1, 2, 3\}$ and $Tech \in \{0, 1\}$, and eight states for $\tau = 4$ where $\tilde{\tau} \in \{1, 2, 3, 4\}$ and $Tech \in \{0, 1\}$. At any moment in time, a generator has a given belief year, and so only perceives that a subset of these discrete states are relevant.

Units subject to U.S. state enforcement are certain about the standard being implemented once it is announced. Thus, belief year τ is not relevant for these U.S. states. For these units,

there are $1 + 2 \times (\text{Enforcement Year} - \text{Announced Year})$ discrete states.

A3.2 Details of Generators' Dynamic Optimization

We solve for generators' optimal dynamic adoption and exit decisions given their beliefs about current and future values of Ω using Bellman equations, as detailed in Section 3.2. For each generator, discrete state, bin of the continuous state, conditional AR(1) state transition moments, and adoption/exit/continue choice, we simulate the expected future value by taking the mean over the values resulting from each of 200 co-prime Halton draw vectors. We transform each vector into normal residuals using the AR(1) regression mean squared errors and calculate the value function of each resulting state. Since these resulting states will potentially lie between grid points, we approximate their values by linearly interpolating across the nearest two grid points in each of the three dimensions (resulting in an interpolation over 8 grid points).

When $\tilde{\tau} > 0$, the distribution of future values depends on whether the generator chooses to adopt or continue. Specifically, in making the choice to adopt new abatement technology, generators recognize that the adoption share will include one additional adopter next period. In the choice to continue instead of adopting, this adopter will not be present. For generators that have already adopted, their choice to continue does not affect the adoption share evolution. They therefore rely upon coefficients from three different adoption share regressions in their choice-specific value function calculations, as we discuss below. When $\tilde{\tau} = 0$, future adoption share is not relevant.

A3.3 Market Evolution Regressions

As discussed in Section 4.1, market evolution is governed by three continuous states that evolve according to AR(1) processes. We assume that the residuals of these state evolution regressions are *i.i.d.*, allowing us to simulate them with co-prime Halton vectors as discussed above. We discuss each of these evolutions in turn.

First, since the fuel price ratio evolves exogenously to the model, we estimate a simple AR(1) regression of fuel price ratio on its lag and a constant term. Because we are identifying

this regression from data—rather than model simulations—we estimate one regression across our entire sample.

Second, coal capacity is the sum of the non-IPP and IPP coal capacity (relative to load). We specify an exogenous AR(1) regression for non-IPP coal capacity, where it depends on its lag, the lagged fuel price ratio, the interaction between these two variables, and a constant. Similar to the fuel price ratio, we estimate one non-IPP coal capacity regression across our entire sample. The IPP coal capacity evolves endogenously in the model. For each U.S. state and year, we simulate the next year’s IPP coal capacity by simulating generators’ decisions given their optimizing behavior as calculated from the Bellman equations. Generators only value knowing future coal capacity in the case where they do not exit. Our simulations approximate this endogeneity of future market structure by randomly selecting one generator to remain active, regardless of its simulated strategy. To simulate the overall coal capacity, we add a draw of IPP coal capacity (derived from the model) to a draw of non-IPP coal capacity (derived from our non-IPP regression). When outside the enforcement window, coal capacity depends on its lag, the lagged fuel price ratio, their interaction, and a constant. When inside the enforcement window, we add the lagged adoption share and years to air toxics standard enforcement as additional regressors.

Third, the share of IPP coal capacity that has adopted abatement technology also evolves endogenously to the model. During the enforcement window, we model the share that has adopted as being a function of its lag, the lagged fuel price ratio, the interaction, the years to air toxics standard enforcement, and a constant. As noted above, generators recognize that their choices will affect this evolution. Accordingly, we estimate three versions of this regression, corresponding to the choices of adoption, continue (when not yet adopted), and continue (when previously adopted). As with IPP coal capacity, we approximate this effect in the first two cases by randomly selecting one generator that had not already adopted and requiring that it adopt or not. When the generator has already adopted, we do not need to randomly select the actions of any generator.

Because we run the coal capacity and adoption share regressions on simulated data, we choose the number of observations and values of the regressors. The number of observations

is the product of the number of simulation draws and simulation years. We start with 1000 simulation draws and 14 simulation years, increasing the number of simulation draws in case of convergence difficulties. We start the regressors at the values in the first year of our sample, 2006. For the coal capacity regression, we use as the dependent variable the expectation of the IPP coal capacity in the next year—given optimizing generator policies—plus a simulated draw from the (exogenous) non-IPP coal capacity process. Similarly, for the adoption regression, the dependent variable is the expectation of the adoption share in the next year, again given optimizing policies. We use expectations rather than simulation draws in order to reduce the variance of our dependent variables.

In contrast, to construct the following year’s regressors, we need to simulate choices given optimizing generator adoption and exit decisions. We do this using a simulation draw from each AR(1) process. We deviate from this updating process in one case in order to obtain sufficient variation in the adoption share variable. Specifically, in the year of standard announcement, we start half of the simulations with each generator having adopted with 0.25 probability and the other half with each generator having adopted with 0.75 probability.

We let the probability of enforcement at the end of the final year before enforcement be equal to P_τ . We therefore simulate the realization of enforcement as a correlated shock to all units in the state.

A3.4 Bootstrapped Standard Errors

In order to calculate standard errors for our parameter estimates, we conduct a parametric bootstrap. This involves simulating 25 data sets created from the equilibrium evaluated at the parameter estimates and re-estimating our model on these data sets.

We solve the equilibrium following the same nested fixed point approach as in our estimation. For the same U.S. states and years as our base data, we then simulate generator adoption and exit and fuel price ratio and non-IPP coal capacity evolution. We assume that generators make decisions in each year to enforcement, $\tilde{\tau}$, given contemporaneous beliefs, τ , so that $\tau = \tilde{\tau}$.

In our simulations, we assume that the exogenous processes—fuel price ratio and non-IPP

coal capacity—evolve according to the AR(1) processes estimated from the observed data. We begin our simulations with these variables set to their actual 2006 values.²⁹ In each subsequent year, we update these variables by drawing from their evolution distributions.

We also begin the endogenous market state variables—IPP coal capacity and adoption share—at their observed 2006 values. We simulate the evolutions of these variables by aggregating simulated draws from generators’ equilibrium strategies.

For convenience, we modify our bootstrap sample from our analysis data in two ways. First, while a few IPP coal generators enter the sample after 2006, we assume that they are present starting in 2006. Second, we limit the realizations of the exogenous processes to not go below 0.01.

A3.5 Counterfactual Calculations

We calculate counterfactual outcomes using a similar approach to how we simulate data for our bootstrap. This involves solving the ABOE of the model and then simulating data. Our approach differs in four ways. First, in each counterfactual, we use different values of the structural parameters. Second, we limit our analysis to only those generators that are subject to MATS enforcement. Third, our counterfactual analysis covers a different time period than the bootstrap. In particular, we begin our analyses in 2012 when MATS was announced and simulate forward 30 years, in order to understand the long-run effects of alternative policy environments. Finally, in order to understand the effects of counterfactual policy environments on pollution outcomes, we calculate the expected pollution outcomes for each of our counterfactual simulations, using our pollution surface introduced in Section 4.2.

A4 Extra Tables and Figures

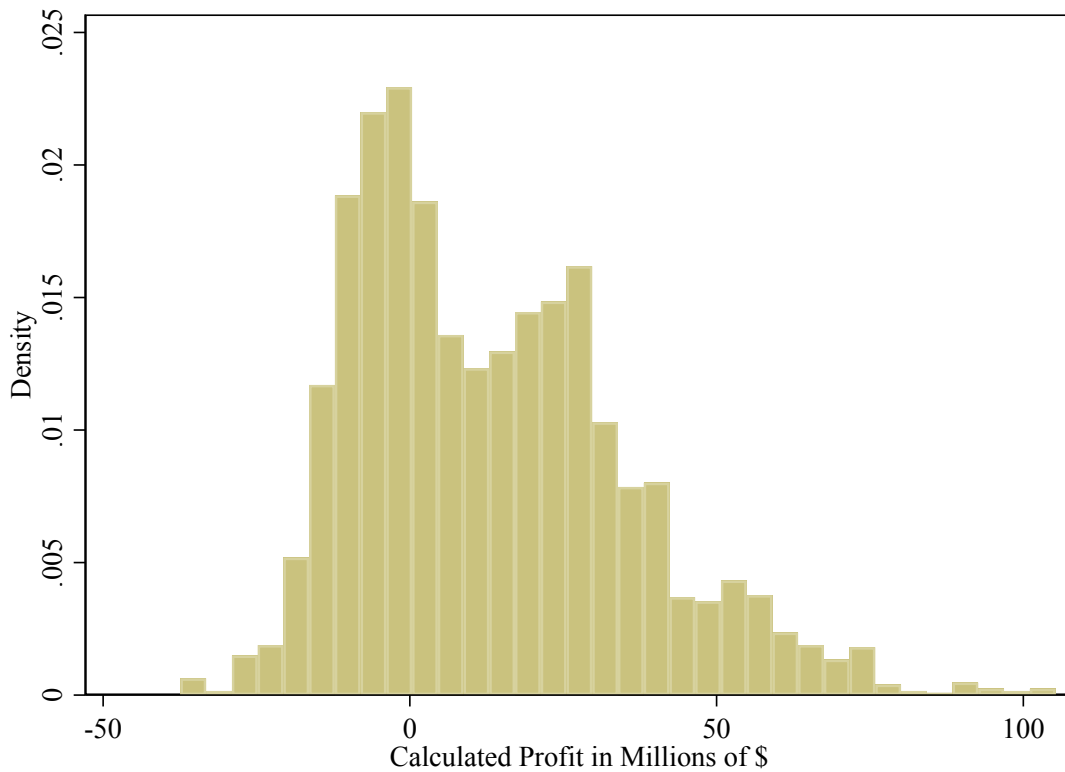
²⁹Florida and Michigan do not report any IPP coal generators until 2008. We therefore start our bootstrapped data sets in 2008 for these states.

Table A1: Announcement and Enforcement Years for U.S. State Standards

State	Announced	Enforced
Connecticut	2003	2009
Massachusetts	2004	2009
Maryland	2006	2011
Illinois	2006	2010
Delaware	2006	2010
Minnesota	2006	2011
Wisconsin	2008	2016

Note: Announcement and enforcement years are based on sources discussed in Appendix A1.

Figure A1: Distribution of Calculated Profits



Note: Histogram of annual profits as calculated with equation (4) and estimated ramping and O&M costs.

Appendix References

Bauer, Scott, “Judge Allows Wis. Rules on Mercury to Move Ahead,” *Associated Press*, 2008.

Cash, Cathy, “Some Coal-Burning States Working on Tougher Mercury Limits; Industry Warns of Shutdowns,” *Electric Utility Week*, 2006.

Congressional Research Service, “Mercury Emissions from Electric Power Plants: States Are Setting Stricter Limits,” 2007. Available at https://www.everycrsreport.com/files/20070222_RL33535_9490786af904352a3618fa139135a3471dab06fb.pdf.

Connecticut General Assembly, “Public Act No. 03-72, Connecticut, House Bill no. 6048,” June 2003.

EPA, “40 CFR §§ 63.9980:63.10042,” 2012.

Gunn, Dorothy, “Board Adopts a Final Notice Proposal Concerning the Control of Mercury Emissions from Coal-fired Electric Generating Units,” December 2006.

Halloran, Elizabeth, “State Takes Action on Mercury,” *Hartford Courant*, July 2003.

Hawthorne, Michael and Tribune staff reporter, “State Moves Against Mercury,” *Chicago Tribune*, 2006.

Hughes, John, “Approving Electric Generating Unit Multi-Pollutant Proposed Regulation,” December 2006.

Illinois Pollution Control Board, “Title 35, Ill. Adm. Code § 225,” December 2006.

Jones, Daniel, “Mad Hatters’ Legacy,” *Hartford Courant*, September 2002.

Legislative Reference Bureau Wisconsin, “Wis. Adm. Code § NR-446: Control of Mercury Emissions,” 2006.

Marimow, Ann and Matthew Mosk, “The Mystery of the Locked Door,” *Washington Post*, April 2006.

Maryland Department of the Environment, “The Maryland Healthy Air Act.”

Maryland General Assembly, “Maryland Code, Environment § 2: Healthy Air Act,”
April 2005.

McIntire, Richard, “Governor Unveils Bold Air Quality Improvement Plan – Maryland’s
Clean Power Rule,” December 2005.

Michigan Department of Environment, Great Lakes, and Energy, “Mercury (Hg),”
Online.

Minnesota State Legislature, “Chapter 201, Minn. Session Laws 2006: Mercury Emissions
Reduction Act of 2006,” May 2006.

Mustian, Mark and Lawrence Demase, “Pennsylvania Supreme Court Upholds Com-
monwealth Court decision, Overturning Pennsylvania Mercury Rule,” *Lexology*, December
2009.

Office of the Registrar of Regulations, “Title 7, Del. Adm. Code § 1146,” December
2006.

Pelton, Tom, “State Loosens Pollution Rule,” *Baltimore Sun*, 2006.

Pioneer Press, “Wisconsin Mercury Rule Challenged in Court,” *Pioneer Press*, 2008.

The Massachusetts Department of Environmental Protection, “Chapter 190, Mass.
Acts of 2006: Massachusetts Mercury Management Act,” May 2006.

United Press International, “Mass.: Toughest Mercury Emissions Limits,” *UPI*, 2004.

Wisconsin Department of Natural Resources, “Air Toxics and Mercury.”

Zibel, Alan, “Ehrlich Signs Stem Cell, Pollution Bills,” *Baltimore Business Journal*, 2006.