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

Gautam Gowrisankaran† Ashley Langer‡ Wendan Zhang§

June 6, 2022

Abstract

Policy-making often occurs through legislation that empowers agencies to make specific rules. This system may lead to more responsive policies, but legal challenges and executive turnover may also create uncertainty. Policy uncertainty affects many important irreversible decisions such as technology adoption, entry, and exit. Because irreversible decisions require considering option value, policy uncertainty can increase costs and delay policy objectives. This paper focuses on the cost of environmental policy uncertainty, specifically uncertainty over enforcement of 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. In order to estimate annual profits, a key input into our model, we develop new approaches to estimating ramping and operation and maintenance costs. We use our estimated model to simulate how policy uncertainty affects generator exit, equilibrium costs, and pollution. We find that generators subject to MATS had substantial uncertainty over whether the policy would be enforced, with the implied probabilities of enforcement of 75% in 2013 and 45% in 2014. This uncertainty led to significant delays in both exit and the adoption of abatement technologies and decreased generator profits by \$840 million,but also decreased pollution by \$769 million to \$2.097 billion.

JEL Codes: L51, Q48, Q52
Keywords: environmental policy, air pollution, mercury, MATS, dynamic equilibrium, moment-based Markov equilibrium

---

*We thank James Archsmith, Nick Buchholz, David Evans, Matt Grennan, Derek Lemoine, Louis Preonas, Nick Ryan and seminar and conference participants at multiple institutions for helpful comments. We are grateful to Utkarsh Kumar and Lorelei Loraine for excellent research assistance.

†Columbia University, CEPR, and NBER
‡University of Arizona and NBER
§Renmin University

1 Introduction

Uncertainty over government policy affects important and irreversible decisions, such as technology adoption, entry, and exit. In many settings, the process of forming and implementing policies creates uncertainty. For instance, in the U.S., new policies often occur through broad legislation that empowers agencies to develop specific regulations to meet legislative goals. This approach allows existing policies to respond to new information and circumstances, and may lead to more responsive policies. However, developing regulations takes time and regulations may be subject to court challenges and executive leadership changes that generate policy uncertainty. This uncertainty can both increase costs and delay policy objectives.

The costs of policy uncertainty depend on at least three factors. First, they depend on the extent to which compliance requires agents to make irreversible decisions, as this requires considering option value (Teisberg, 1993; Dixit et al., 1994). Maintaining option value may increase compliance costs and decrease agents' ability to respond to other market signals. Second, the costs depend on how maintaining option value changes externalities generated by agents. For instance, the costs of policy uncertainty change if the uncertainty causes polluting firms to delay exit or technology adoption or lowers their ability to respond to market signals. Finally, the costs depend on whether the policy is eventually enforced, as even discussion of a policy can lead agents to make irreversible decisions.

This paper considers the impact of environmental policy uncertainty in the electricity sector. We estimate agents' beliefs regarding the likelihood of enforcement for the Mercury and Air Toxics Standard (MATS). MATS required that electricity generators either adopt substantial pollution abatement equipment in order to remain in the market. We model generators' technology adoption and exit decisions within a dynamic oligopoly model. Our estimates allow us to simulate how policy uncertainty affects counterfactual outcomes in the industry including generator exit, equilibrium costs, and pollution. We also compare the impact of MATS to other potential policies, such as subsidies for generator retirement. While we focus on environmental policy uncertainty, firms in many sectors—including healthcare, telecommunications, and finance—make important decisions in the face of uncertain policies.

Coal-fired electricity generating units (EGUs, or generators) are the primary emitters of air toxics from electricity generation. 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 EPA released the final MATS rule in 2012 with enforcement in 2016. The regulation has been subject to substantial judicial and administrative review, but ultimately survived. Thus, understanding the role of policy uncertainty in this context has important financial and environmental ramifications. In particular, uncertainty surrounding enforcement could lead to delayed generator exit, higher financial costs, and increased pollution.

While the federal government was formulating air toxics policies, some U.S. states mandated air toxics reductions for generators within their borders. Because these policies were either legislative or developed with input from local power producers, they were largely not subject to the same level of uncertainty. These U.S. state policies therefore allow us to understand generator behavior in the absence of policy uncertainty and ultimately to identify the level of policy uncertainty encapsulated in MATS.

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 subject to potential MATS enforcement 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.

We assume that generators compete in a moment-based Markov equilibrium (MME, Ifrach and Weintraub, 2017), where coal IPP generators make their decisions using an aggregate state that summarizes market characteristics. In our case, generators monitor fuel prices, market coal capacity, and technology adoption shares. Equilibrium effects are potentially important in our context. For instance, one generator's exit will increase rivals' profits and decrease the likelihood of their exit. Further, a generator may adopt abatement technology partially to signal a commitment to remaining in the market (Riordan, 1992; Schmidt-Dengler, 2006). Consistent with oligopoly interactions, generators in our model recognize that their decisions impact both coal capacity and the technology adoption share. However, a limitation of our approach is that we do not model ownership linkages across generators.

In general, it might be difficult to separate a higher perceived probability of future MATS enforcement from a higher exit scrap value, since both would encourage generators to exit. In our case, given our assumption that U.S. state enforcement was perceived to be certain, comparing exit rates between generators subject to U.S. state air toxics standards and those subject to MATS—after controlling for other differences—identifies these perceived probabilities.

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

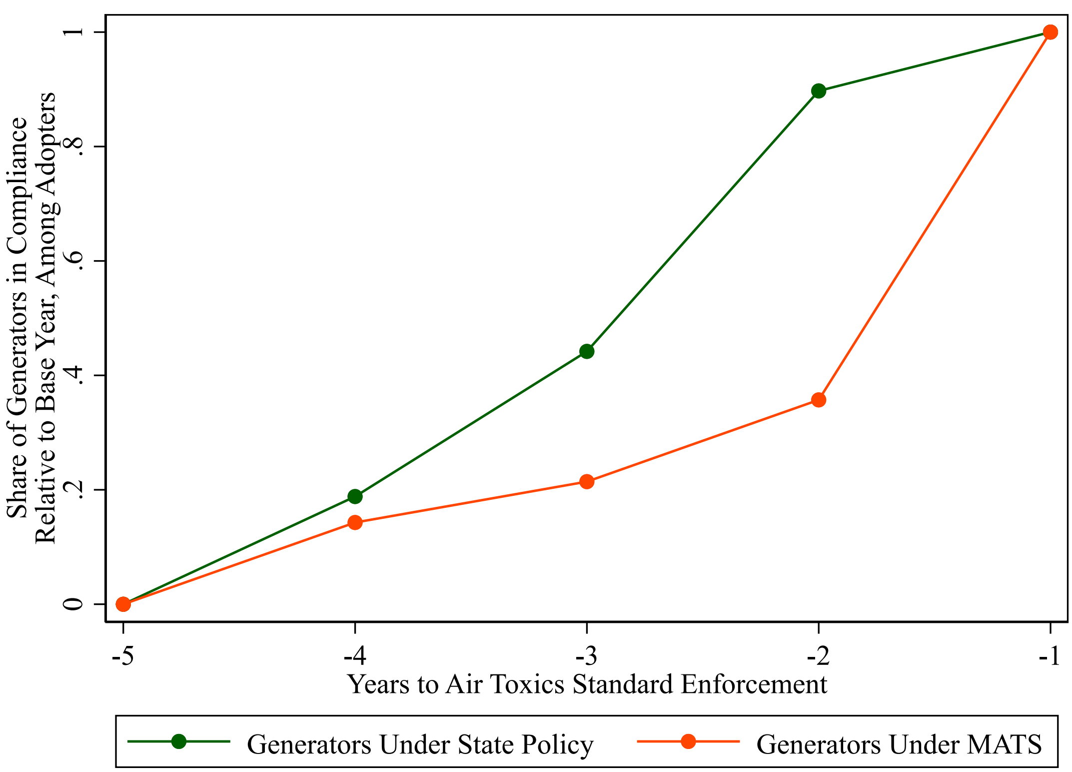

Note: Share of generators that have adopted air toxics abatement technology among generators subject to U.S. state air toxics policies, relative to the total number of generators that had not adopted abatement technology five years before air toxics standard enforcement.

Figure 1 provides intuition behind this identification argument. The figure shows the share of generators in our analysis data that will eventually adopt air toxics abatement technology in each year prior to air toxics standard enforcement. The green line shows this share for generators subject to state policies, and the orange line shows this share for generators subject to MATS. We see that four years prior to enforcement, both lines are fairly similar, but at three and two years prior to enforcement, generators subject to MATS have adopted abatement technologies at substantially lower rates than generators subject to state policies, suggesting that generators believed the probability of 2016 MATS enforcement to be substantially below one.

Our dynamic model takes as an input generators' profits in the electricity market given their state. Generators earn revenues from their electricity sales, and bear the costs of fuel, ramping (adjusting their generation level), and operations & maintenance (O&M). Fuel costs are observable in our data. Ramping costs have been shown to be important in the context of electricity generation (Cullen, 2014; Reguant, 2014; Linn and McCormack, 2019), and are particularly critical in our context because of changes in the electricity industry during our analysis period. 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 generally no longer being profitable to run during hours with low demand, when wholesale electricity prices are low. Thus, coal generators needed to cycle more often later in our sample, leading ramping costs to increase in importance. As an example, in 2008, coal generators in our data spent an average of 31.4 hours at their maximum generation level each time they ramped to maximum generation, but this dropped to 21.0 hours in 2017. If we were to ignore ramping costs, we would understate the profit reduction that has occurred over our sample, which would then understate generators' incentives to exit, and yield biased estimates of the structural parameters.

We estimate ramping and O&M costs with new, tractable methods that incorporate dynamic linkages between hours. We develop a conceptual experiment that compares sets of hours with similar future electricity prices, where generation varied in the previous hour. We implement this approach with a regression of the chosen generation level on revenues, lagged generation, and controls for future value. The difference in operating probabilities across these sets identifies ramping costs if the controls sufficiently account for the difference in continuation value across the sets. We estimate O&M costs with a related simple choice model. We use observed generation choices and our cost estimates to recover generator profits and estimate the relationship between profits and underlying dynamic states. Our cost estimation leverages the assumption that generators are price takers, but our calculation of profits—which feeds into the adoption and exit model—uses the observed behavior, which could be the result of a more general model.

Relation to the Literature: This paper builds on three main literatures. First, we build on 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) all develop measures of uncertainty, using newspaper text, SEC filing text, and options prices, respectively. Kellogg (2014), Dorsey (2019), and Handley and Li (2020) further examine the impact of uncertainty on economic outcomes such as oil drilling and firm investments. We build on this literature by recovering generators' beliefs over enforcement probability and using them to perform counterfactual simulations.

Second, we build on a literature that estimates structural models of the electricity market. Papers in this literature include Fowlie (2010), Abito et al. (2018), Linn and McCormack (2019), Scott (2021), Elliott (2022), and Gowrisankaran et al. (2022). Beyond incorporating policy uncertainty, we build on this literature by developing a dynamic oligopoly model and explicitly estimating ramping and O&M costs within this model.[1]

Finally, we build on a literature that estimates dynamic oligopoly models by approximating future states. Ifrach and Weintraub (2017) developed the concept of MME, which builds on the concept of oblivious equilibrium (Weintraub et al., 2008) but allows for the computation of industries that are not in a steady state. Recent empirical applications of these methods include Gerarden (2017), Vreugdenhil (2020), and Jeon (2022).

Summary of findings: We estimate that generator's perception of the probability of

[1] Gowrisankaran et al. (2022)—written by an overlapping set of coauthors—considers the role of U.S. state rate-of-return regulation in energy transitions, while this paper focuses on generators subject to market incentives.

MATS enforcement started at 99.9% in 2012, dropped to as low as 45.0% in 2014, and then rose to 99.9% in 2015, for an average probability of enforcement of 79.83%. Further, we estimate that exit and technology adoption are both costly, with an exit costing the generator \$229 million, technology adoption to comply with U.S. state policies costing \$147 million, and technology adoption to comply with MATS costing \$554 million. Over 30 years, our model predicts that generators subject to MATS spend \$7.24 billion on technology adoption and \$19.28 billion on exit costs.

Our counterfactual analyses investigate the impact of eliminating policy uncertainty and reducing generator exit costs. We calculate the effect of removing policy uncertainty by investigating the differences in equilibrium outcomes under an environment where policy implementation is decided in 2016 versus committing in 2012 to whether the policy will be enforced, with the mean generator perceived probability of enforcement. We find that eliminating uncertainty would increase expected generator profits by \$880 million in present discounted value. Nevertheless, eliminating uncertainty also *increases* expected pollution by about 62 million pounds of SO_2 over 30 years, valued at\$769 million to \$2.097 billion dollars. This occurs because generators are better able to align their adoption and exit decisions with market conditions such as fuel prices, which allows them to operate more—and generate more pollution—when market conditions are favorable. Removing exit costs—for instance by having the government pay for site remediation and other costs—reduces the number of generators in 2016 by 14.4% and increases generator profits by 49.8%, but this comes largely from government transfers.

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 model of operations and adoption/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,[2] under the 1990 Clean Air Act Amendments (CAAA). Air toxics are distinct from criteria air pollutants, such as NO_x and SO_2, which have been subject to long-standing CAAA regulation.

While the CAAA gave the EPA authority to regulate air toxics, it left the form of this regulation open. 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*,[3] which found that the EPA should have regulated mercury under the more stringent maximum achievable control technology (MACT) standard rather than under CAMR, 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*[4] that the EPA would have to regulate mercury and other air toxics together, rather than starting with mercury alone. 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.

To remove air toxics from emissions, generators must convert pollutants into water-soluble forms, bind them to larger particles, and precipitate the new compounds with a particulate matter catcher.[5] This basic process can be achieved with different technologies, making it difficult to determine compliance from technology adoption data. 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 the technology will

[2] https://www.epa.gov/haps/what-are-hazardous-air-pollutants.

[3] 517 F.3d 574 (D.C. Cir. 2008).

[4] 551 F. 3d 1019 (D.C. Cir. 2008), also known as the "Brick MACT" decision.

[5] Compliance may also potentially be achievable by fuel switching to cleaner coal.

be required.

EPA regulations have been vulnerable to two key sources of uncertainty. First, EPA rule-making has been subject to substantial legal challenge, up to and including Supreme Court review. Further, since the EPA is a part of the executive branch, changes in executive leadership have drastically altered the EPA's focus and regulatory philosophy. These changes in executive leadership also interact with legal challenges, e.g., when a new administration chooses not to pursue the same legal approach as the previous administration.

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. Uncertainty over MATS continued due to the 2016 election of President Donald Trump. The Trump Administration did not file the justification but left MATS in place nonetheless. In 2021, the Biden Administration did file the justification of MATS as "appropriate and necessary."

While the EPA and the courts iterated on federal air toxics regulations, some U.S. states moved ahead with their own, internal standards. Individual U.S. states have the authority to regulate emissions within their own borders. In contrast to federal regulations, U.S. state regulations were generally subject to very little uncertainty, for at least two reasons. First, in some states (e.g. CT and MD), these regulations were passed into law by state legislatures (Halloran, 2003; Pelton, 2006). Second, even in states such as IL and MA (Hawthorne and Tribune staff reporter, 2006; United Press International, 2004) where the standards were created as rules issued by the state environmental agency, the standards were generally developed in tandem with the owners of large coal generators, 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.

U.S. state air toxics regulations were weaker than MATS in some cases. In particular, the specified standards for mercury levels were sometimes higher than under MATS. The state standards mostly covered mercury rather than all air toxics. Enforcement in some cases consisted of the regulator approving generators' abatement technology adoption plans rather than emissions monitoring, as occurs under MATS.

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 emissions of SO_2 and NO_X 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.

2.2 Data Sources

We create one analysis data set at the generator-hour level and another one at the generator-year 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. Our data sets 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 coal generation for IPPs in the U.S.[6] Each observation in the CEMS data provides the heat input of the fuel used (in MMBtu), electricity production (in MWh), and criteria pollutant emissions (in pounds/MMBtu) for each generator that the EPA monitors with a CEMS, for CO_2, NO_X, and SO_2. The CEMS data further report a facility identifier and the location of each generator. As we discuss below, we use SO_2 as a proxy to measure the adoption of air toxics abatement technology and to measure annual generator emissions.

We form our analysis data by merging the hourly CEMS data with several other data

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

sources. First, the EPA also 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, in order to determine the timing of air toxics standard announcement and enforcement for each generator, we begin with a list of U.S. states with state air toxics standards as listed by the Congressional Research Service (Congressional Research Service, 2007). We then validate this list and the years of standard announcement and enforcement with a combination of newspaper articles, state environmental agency press releases, and state statutes. Table A1 in the Online Appendix lists the announcement and enforcement years for states with their own standards.

Third, we merge in annual data from the Energy Information Administration (EIA) Form 923 on whether the facility is an independent power producer (IPP) or not. This allows us to identify generators facing market incentives. Most of those generators are located in U.S. states with restructured electricity markets. We choose this definition of generators facing market incentives rather than characterizing generators by U.S state regulatory status for two reasons. First, some generators in U.S. states with regulated electricity markets may sell into restructured markets and face market incentives. Second, generators owned by load-serving entities in restructured states may face non-market incentives (Gowrisankaran et al., 2022). 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.

Fourth, we collect annual, U.S. state-level natural gas and coal prices from 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 reflects the opportunity cost of fuel faced by generators.

Fifth, we merge in hourly 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.[7] For some U.S. states, the data report prices for multiple nodes. For these states, we recover hourly electricity prices by taking the mean over the nodes. 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.

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

Seventh, we recover hourly U.S. state-level electricity load from 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 the 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 heating and cooling degree days by differencing the mean of the state-level daily minimum and maximum temperatures from 65 degrees.

2.3 Construction of Key Variables

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

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

measure air toxics compliance through a mercury CEMS.

Because the most cost-effective technologies that abate both mercury and other air toxics also reduce SO_2, one important way that the EPA determines MATS compliance is via SO_2 emissions rates. In particular, the MATS final rule (77 FR 9304) specifies that generators can comply with MATS by having SO_2 emissions rates below 0.2 lbs/MMBtu.

For most generators, we observed a large decline in SO_2 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 SO_2 emissions rate falls below 0.4, or (ii) its annual emissions rate falls by 40% or more.[8]

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

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.[9] 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, specif-

[8]In 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.

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

ically, 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.[10] 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.

2.4 Descriptive Statistics

Table 1: Generator Descriptive Statistics by Regulatory Regime

	State Policy	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 analysis sample of IPP coal generators.

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 323 IPP coal generators, of which 93 are subject to U.S. state enforcement and 239 are subject to MATS. They include a total of 2,889 generator-year observations. The table

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

further shows that 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 states. Overall, these statistics suggest that it is reasonable to assume that generators are technologically similar across the two sets of U.S. states. Generators in states subject to U.S. state policy do face lower average coal fuel prices than generators subject to MATS (\$1.70 vs \$2.12 per MMBtu), which then feeds into lower marginal fuel costs (\$17.86/MWh vs \$26.07/MWh).

Table 2: Counts of Generators Adopting Abatement Technology or Exiting

	State Policy			MATS		
Years to Enforcement	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 analysis sample of IPP coal generators.

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 the years to policy enforcement. Two key patterns emerge from these data. First, generators subject to state policies were much more likely to adopt pollution abatement technology rather than exit, while the reverse is true for generators subject to MATS. Second, in the year of policy enforcement, we see 20 generators exit, but none adopt pollution abatement technology. This leads us to assume that regulators essentially forced generators that were not in compliance with standards at the point of enforcement to exit.

Table 3 presents descriptive statistics on our hourly analysis sample. 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

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)
Heating Degree Days	13.90 (14.88)
Cooling Degree Days	2.83 (4.82)
N	26,014,143

Note: Authors' calculations based on analysis sample of IPP coal generators at hourly level.

with a mean of $39.07/MWh. Mean U.S. state electricity demand was 13.29 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.

Table 4: Change in Ramping and Natural Gas Prices Over Time

Year	Hours at Max Generation Per Ramp	Natural Gas Price Over Coal Price
2006	27.99	4.62
2007	32.96	4.39
2008	30.42	4.82
2009	23.08	2.35
2010	26.70	2.27
2011	24.68	1.94
2012	20.73	1.32
2013	22.76	1.82
2014	24.41	2.28
2015	17.57	1.41
2016	17.17	1.35
2017	19.94	1.63

Note: Authors' calculations based on analysis sample of IPP coal generators. Each observation in the second column pertains to one observed ramp to maximum generation. Each observation in the third column pertains to one observed generator-year.

Table 4 shows changes over time in the hours that coal generators spend at maximum generation each time they ramp to maximum generation and the ratio of the natural gas price to the coal price. Starting in 2009, both natural gas prices and the hours spent at maximum generation per ramp declined sharply. This change was likely caused by the rise of fracking, which led to a substantial drop in the price of natural gas. In many cases, fracking caused coal generators to be replaced by combined-cycle natural gas sources as the lowest-cost fossil-fuel electricity source. They therefore went from generating regardless of electricity prices to only generating when electricity prices were particularly high.

Figure 2: Distribution of Consecutive Hours at Maximum Generation

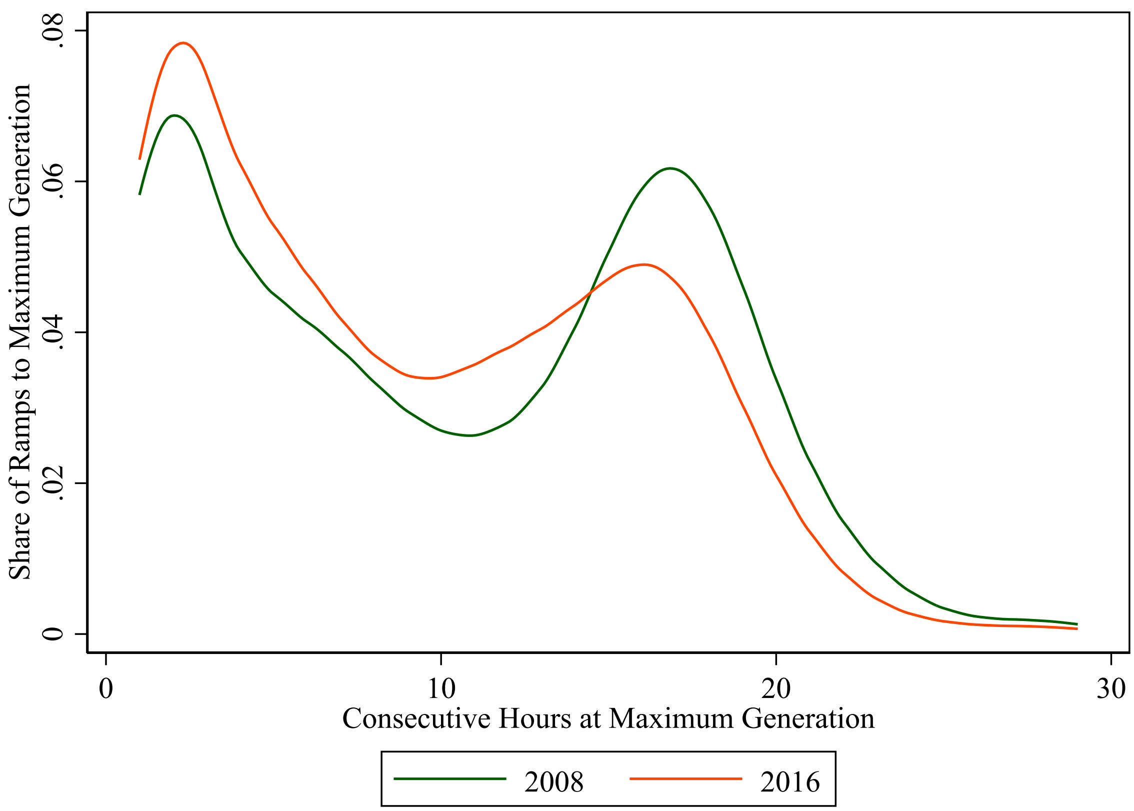

Note: Authors' calculations based on analysis sample of IPP coal generators. Each observation is one observed ramp to maximum generation. The green line displays a kernel density of the number of hours at maximum capacity for each ramp to maximum capacity in 2008. The orange line displays the same information for 2016. Both densities are truncated at 30 hours per ramp.

Figure 2 further shows the density of hours at maximum generation per ramp for 2008 and 2016. In both years, there is a peak of the distribution around two hours and another one around 18 hours. However, the 2016 distribution has more weight on the first peak, less weight on the second peak, and fewer ramps with greater than 20 hours at maximum generation. This table and figure demonstrate that an accurate measure of ramping costs is necessary to understand how the operating profits of coal generators have changed over time. Ignoring these costs might bias our estimates of abatement technology adoption and exit costs.

3 Model

We develop an infinite-horizon dynamic equilibrium model of abatement technology adoption, exit, and production for coal generator 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, $j = 1, \ldots, J_t$, 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 policies and generators obtaining information about previously announced policies. 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. 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.[11] We denote the

[11] While we allow for uncertainty about whether the policy will be enforced, we assume that the level of

full set of perceived probabilities $P_{\tau_0}, \ldots, P_1$. 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 $P_{\tau_0} = \ldots = P_1 = 1$.

Second, generators make adoption and exit decisions.[12] 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 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. Generators that have already adopted abatement technology 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 compete in hourly electricity markets and earn profits from selling electricity. 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). Ramping costs imply that production decisions across hours are dynamic, while O&M costs are essentially the remaining per-MW cost of generating in a given hour. Section 3.3 discusses hourly optimization.

Finally, any exit or abatement adoption decisions made in this year are realized. At this point, if this 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.

the standard and the date of potential enforcement is certain.

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

3.1 Annual State Space and Equilibrium

Generator j's adoption and exit decisions are a function of its annual dynamic state and its perceptions of its competitors' actions at this state. In principle, beyond the market, m, j's annual state includes its non-time-varying characteristics: heat rate, $heat_j$, capacity, K_j, and coal price, f^C; its time-varying characteristics: annual profits, Π_{jt}, $Tech_{jt}$, the belief year, τ, years to potential air toxics standard enforcement, τ',[13] and its cost shocks, $\vec{\varepsilon}_{jt}$; and the characteristics of its competitors, both IPP coal generators and other sources.

We assume that the market reflects a moment-based Markov equilibrium. In an MME, each generator includes in its state a set of aggregated market characteristics. In our case, this set includes (1) the natural gas to coal fuel price ratio, (2) the (combined IPP and non-IPP) coal capacity relative to the 95th percentile of hourly load in the market, and (3) the share of IPP coal generators that have adopted abatement technology. This approach simplifies the state space relative to a Markov-perfect equilibrium, which is important for estimation because the large number of competitors would otherwise result in a curse of dimensionality problem.

We include the fuel price ratio as a market characteristic because relative natural gas fuel price affects rivals' 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.[14] Coal capacity and adoption share approximate dynamic oligopoly incentives in this market. 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.

Consistent with dynamic oligopoly, we allow generators in our MME to understand that their actions influence the aggregate industry state. Specifically, a generator which chooses

[13] If an air toxics standard has not yet been announced or the enforcement year has already passed, then $\tau = 0$.

[14] We 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.

not to exit understands that the coal capacity next period will be higher than if it chooses to exit. Similarly, when a generator chooses to adopt abatement technology, it recognizes that this will increase the share of capacity having adopted abatement technology next period. Modeling expectations of market evolution in this way allows generators to use their adoption and exit decisions as costly preemptive signals.

We assume that generators' believe that the three aggregate market states will evolve according to 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 adoption versus non-adoption.[15] Because market fundamentals may vary across U.S. state and belief year τ, we further disaggregate our AR(1) regressions to this level.

At an MME, aggregated generators' actions must be consistent with the market evolution resulting from those actions. Specifically, an MME for a U.S. state and belief year consists of first, a set of strategies for every player (discussed below), and second, regression coefficients on industry evolution for each of the AR(1) regressions. Together, these must satisfy that 1) the state-contingent strategies reflect optimizing behavior given the regression coefficients, and 2) data simulated from these strategies yield these regression coefficients. The MME of our model therefore reflects a fixed point of these decisions and the expectations over state-contingent IPP coal capacity that these decisions imply.

3.2 Annual Dynamic Optimization

Generator j makes adoption and exit decisions based on its state, $(\Omega_{jt}, Tech_{jt}, \tau, \tau', \vec{\varepsilon}_{jt})$, where Ω_{jt} includes j's non-time-varying characteristics and the three aggregated market characteristics noted above. We now exposit the determinants of generator actions as a

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

function of this state.

When $\tau = 0$, generators face a relatively simple choice between exiting or continuing. When $\tau > 0$, 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 $\tau = 1$ or $\tau > 1$ because if the generator chooses to continue when $tau = 1$ and the air toxics standard is enforced, the generator will be forced to exit.

Consider first the case of generator j one year from enforcement ($\tau' = 1$) that has not previously adopted abatement technology.[16] This generator faces the three choices of continuing without adopting, adopting, or exiting. If the generator continues without adopting and the air toxics standards are enforced, it will be forced to exit and receive X. We can write the Bellman for this case as:

$$\begin{aligned} V(\Omega, Tech = 0, \tau, \tau' = 1, \vec{\varepsilon}) = \\ \max \big\{ \Pi(\Omega) + P_\tau X + (1 - P_\tau)\beta E[V(\Omega', 0, \tau, 0, \vec{\varepsilon'})|\Omega, \text{No Standard}] + \sigma\varepsilon_c, \\ \Pi(\Omega) - A + \beta\{P_\tau E[V(\Omega', 1, \tau, 0, \vec{\varepsilon'})|\Omega, \text{Standard}] \\ +(1 - P_\tau)E[V(\Omega', 1, \tau, 0, \vec{\varepsilon'})|\Omega, \text{No Standard}]\} + \sigma\varepsilon_a. \\ \Pi(\Omega) + X + \sigma\varepsilon_x \big\}, \end{aligned} \tag{1}$$

where Ω' is the Ω component of the state next period. Equation (1) shows that the perceived probability of air toxics standard enforcement enters into the $\tau' = 1$ value function.[17] 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

[16] Because generators believe that their enforcement beliefs will remain constant in the future, we must consider instances where belief year, τ, is greater than one even though $\tau' = 1$.

[17] We assume that Π is a function of Ω but not $Tech$ because our O&M cost estimates are similar for generators that have adopted technology.

standard is enforced. The third choice is exit. Generators that have already adopted face a choice between continuing and exiting, with no probability of being forced from the market if the standard is enforced.

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

$$V(\Omega, 0, \tau, \tau', \vec{\varepsilon}) = \max\big\{\Pi(\Omega) + \beta E[V(\Omega', 0, \tau, \tau' - 1, \vec{\varepsilon'})|\Omega] + \sigma\varepsilon_c, \qquad (2)$$
$$\Pi(\Omega) - A + \beta E[V(\Omega', 1, \tau, \tau' - 1, \vec{\varepsilon'})|\Omega] + \sigma\varepsilon_a,\ \Pi(\Omega) + X + \sigma\varepsilon_x\big\}.$$

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

3.3 Hourly Optimization

Each year, t, includes hours $h = 1, \ldots, H$, with generator j choosing a generation quantity q_{jh} in each of these hours. It chooses each q_{jh} to maximize expected annual profits, which are the sum over hours of revenues minus its three costs: fuel, ramping, and O&M. Generator j's fuel costs per MW of production are the product of the coal price, f^C, and $heat_j$.

We model ramping and O&M costs by developing and estimating a dynamic model of hourly operations. 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_jK_j for $L_j \in (0, 1)$; and (3) not generating.[18] In order to specify the generator's hourly operations Bellman equation, we assume that the generator makes an infinite horizon decision with annual discount rate β. We also assume that firms that own multiple generators make independent state-contingent profit maximization decisions for each generator.

At any hour (and suppressing the h subscript), let $q_j \in \{0, L_jK_j, K_j\}$ denote the current generation level, $\tilde{q}_j$ denote the level in the previous hour, and ω_j denote the generator's short-

[18]Dividing generation into three levels intuitively makes sense if generators take hourly electricity prices as given, bear a startup cost, have other ramping costs that are proportional to the increase in generation, and have a minimum generation level of L_jK_j.

run dynamic state, which includes information that affects current and future prices. For instance, ω_j might include load, the time of day, weather, and rival generators' availability. Further, let $r_{\tilde{q},q}$ be the cost of ramping from $\tilde{q}$ to q for $\tilde{q} < q$,[19] om be the per MWh O&M costs, and $p(\omega, q)$ to be the expected wholesale electricity price given ω and q. Combining these terms, generator j's hourly Bellman equation is:

$$v(\tilde{q}_j, \omega_j) = \max_q \left\{ \pi(\tilde{q}_j, \omega_j, q) + \beta^{1/H} E[v(q, \omega_j'|\omega_j)] \right\}, \tag{3}$$

where per-hour profit is:

$$\pi(\tilde{q}_j, \omega_j, q) \equiv q \times [p(\omega_j, q) - heat_j \times f^C - om] - \mathbb{1}\{\tilde{q}_j < q\} r_{\tilde{q}_j,q} + \sigma^g \varepsilon^g_{jq}. \tag{4}$$

Our model allows om to affect the generator's marginal costs but not fixed costs. We assume that the generation unobservable ε^g_{jq} is type 1 extreme value and $i.i.d.$ across generation levels, hours, and generators. We further assume that a generator observes its own ω_j and ε^g_{jq} at the start of each hour, but forms expectations about future values of ω_j based on the current state. Because operating profits are measured in dollars, we include a parameter, σ^g, that scales ε^g_{jq}. We sum the optimized hourly profits into annual profits, $\Pi_{jt} = \sum_h \pi(\tilde{q}_{j,h-1}, \omega_{jh}, q_{jh})$.

4 Estimation and Identification

This section discusses the estimation and identification of generation cost parameters, annual profits and pollution, and the parameters governing adoption and exit.

4.1 Generation Cost Parameters

Our model depends fundamentally on generators earning profits in hourly electricity markets. Hourly profits are equal to revenues minus fuel, ramping, and O&M costs. We calculate revenues directly from our data by multiplying wholesale electricity price, p_h, with quantity

[19] For simplicity, our base specification does not allow for "deramping" costs, although we provide additional regressions that allow for this possibility.

supplied, q_h. Our data also include generator heat rates and fuel prices, which allow us to calculate fuel costs.

We do not directly observe ramping or O&M costs, and therefore we recover them from observed behavior. We estimate these costs assuming that generators are price takers in the hourly electricity markets. As we discuss below, conditional on our cost estimates, we will use observed generation decisions, which are consistent with dynamic oligopoly interactions, to calculate profits.

We develop a simple estimator of our dynamic model in (3) that recovers ramping costs. In principle, ramping costs are identified by the extra revenue that generators expect to earn when they increase their generation level. The issue is that ramping is fundamentally dynamic: generators may increase generation in an hour in order to capture the option value of remaining at a high generation level in future hours. Our estimation approach involves finding "conceptual experiments" where the generator finds itself in different sets of situations with identical information about future prices, ω, but where generation in the previous hour varies across the sets. The difference in the probability of generation levels across these sets identifies the ramping costs.

We implement our approach with a trinomial choice model of hourly generation levels (as defined in Section 2.3) that follows from equations (3) and (4):

$$u(q_h|\tilde{q}_h, \omega_h) = \overbrace{\frac{1}{\sigma^g}[q_h p_h]}^{\text{Revenues}} - \overbrace{\frac{1}{\sigma^g} r_{\tilde{q}_h, q_h} \mathbb{1}\{\tilde{q}_h < q_h\}}^{\text{Ramping cost term}} - \overbrace{\psi x(q_h, \omega_h)}^{\text{Other costs and relative continuation value}} + \varepsilon^g_{qh}. \quad (5)$$

There are four key parameters for each generator: the standard deviation of the unobservable, σ^g, which scales revenues, the costs of ramping from off to LK ($r_{0,LK}$), off to K ($r_{0,K}$), and LK to K ($r_{LK,K}$). In this regression, the x_h variables serve as controls that capture anything that might be included in non-ramping costs and the continuation value relative to not generating, which ties this regression to our "conceptual experiments." We include fuel costs, weather, hour-of-day, and current load, interacted with indicators for the generation levels of LK and K in x_h and the year. We explicitly do not include $\tilde{q}_h$ in these controls.

Comparing (5) to equations (3) and (4), we have made two changes. First, we have collapsed the future value term in (3) and the fuel cost and operations and maintenance cost terms from (4) with $x(q_h, \omega_h)$. Second, we divided each coefficient by the standard deviation of the unobservable. This change allows us to estimate generation decisions in a multinomial choice framework, where we are now estimating a ramping cost parameter of $r_{\tilde{q}_h,q_h}/\sigma^g$. We therefore recover an estimate of ramping costs in dollars, $r_{\tilde{q}_h,q_h}$, by dividing our estimate of the ramping cost parameter by the estimated coefficient on revenues.

Identification of ramping costs hinges on the ability of x_h to appropriately capture the generator's relative state-contingent continuation values. This relies on three exclusion restrictions: that revenues, $q_h p_h$, and the lagged generation level, $\tilde{q}_h$, do not enter into x_h, and that units take wholesale electricity prices on the hourly market as given. Thus, x_h must capture the generator's expectations of relative future values sufficiently well that revenues and lagged generation do not provide additional information on these expectations. To ensure that x_h is sufficiently rich, we estimate it with a flexible functional form that captures expected future generator costs and revenues.[20]

Fundamentally, our approach is based on comparing generation across hours with identical prices and continuation values but different lagged generation levels. An alternative approach would be a conditional choice probability (CCP) estimator (e.g. Hotz and Miller, 1993; Arcidiacono and Miller, 2011), which uses a log transformation of the differences in probabilities across generation level by state. Given the infrequency with which generators change generation levels, many of these probabilities would be very close to 0, making it difficult to estimate the future values accurately. In contrast, our approach only requires that we have sufficient information to appropriately control for generators' continuation values and sufficient observations to have "matches" where continuation values are very similar, but lagged generation levels differ.

Turning to O&M costs, $om \times q$, we assume that they are proportional to the quantity

[20] We have estimated alternative specifications where we interact x_h with functions of prices over the following 20 hours. These specifications should very accurately control for continuation values, but since they include future prices, they incorporate more information than generators would actually have available. We obtain broadly similar ramping cost estimates with either specification, so use the ones based on available information.

generated, implying that they act much like an option-specific constant term in equation (5), entering via $\psi x(q_h, \omega_h)$. However, this term also includes fuel costs and the continuation value, so we cannot recover O&M costs directly from this regression. Recovering fuel costs is straightforward given that we observe heat rates and fuel prices, so the key challenge is in separating O&M costs from continuation values.

We recover O&M costs by focusing on the production decision within a set of hours where the continuation values are very similar. In particular, we estimate O&M costs with regressions using a subsample consisting of windows of hours around an actual ramp or deramp. We assume that the generator knows that it will be ramping or deramping exactly once during this window, as observed in the data. This implies that its continuation value will be the same regardless of which hour it chooses to change generation within the window, which allows us to separate O&M costs from the continuation values.

We operationalize this idea by using a window around ramping events in our hourly data that includes the 6 hours before, the observed hour, and the 5 hours after a generation change. We focus on the subsample of windows around when generators ramp from minimum to maximum generation or vice versa.[21] We further assume that the generator has perfect information about the wholesale electricity prices for every hour in this window.

Given these assumptions, the generator chooses the hour in which to change generation that maximizes the sum of profits over this window. We explain our approach for windows where generator j is ramping up to maximum generation; windows where the generator is ramping down to minimum generation are similar, but in reverse. We estimate a multinomial logit model with 12 choices, one for each hour in the window. For a window that begins in hour $\underline{h}$, we can write the relative profits in the window, w, from ramping at hour $h \in \{\underline{h}, \ldots, \underline{h}+11\}$ as:

$$\pi^w_{jh} = \sum_{\tilde{h}=h}^{\underline{h}+11} (p_{\tilde{h}} - heat_j \times f^C)(K_j - K_j L_j) - om \times (\underline{h} + 12 - h) \times (K_j - K_j L_j) + \sigma^w \varepsilon^w_{jh}, \quad (6)$$

[21] These changes are better for identifying O&M costs than hours when generators turn off or on because those decisions are more likely to be affected by unobservable factors such as required maintenance, and so may not fully be responding to immediate market price incentives.

where the parameter om gets multiplied by the total additional generation in all hours in the window starting at h. The unobservable ε^w_{jh} will capture the conditional distribution of the underlying structural residuals ε^g that stem from the decision to ramp at hour h, and have a standard deviation, σ^w.[22] As an approximation, we use a type 1 extreme value distribution for ε^w_{jh}. Because the generator bears the same ramping costs and continuation values regardless of its choice of hour in which to ramp, these terms do not enter its maximization in (6).

Our model identifies om from the increase in hourly profits in the hour the generator changes generation. In our ramping window subsample, we always observe the generator changing generation in the seventh hour of the window. Figure 3 uses this subsample to provide intuition behind the identification of our estimates. The figure shows the hourly profit per MW before O&M costs (so revenues minus fuel costs) in the hours around each ramp from minimum to maximum generation or back.[23] We indicate the fact that the generator always chooses to ramp in between hour 6 and hour 7 with a grey vertical dashed line. The O&M costs that rationalize this choice are higher than the profit bar in hour 6, but lower than the profit bar in hour 7. The mean of these estimates is approximately \$17, as indicated with the horizontal grey dotted line.

4.2 Annual Profits and Pollution

Having recovered ramping and O&M costs, we then calculate each generator's annual profits and pollution. Specifically, we write annual expected profits as:

$$\Pi_{jt} = \sum_{h=1}^{H} \big[q_{jh} p_h - q_{jh} \times heat_j \times f^C - \hat{r}_{\tilde{q}_h, q_h} \mathbb{1}\{\tilde{q}_{jh} < q_{jh}\} - \hat{om} \times q_{jh} \qquad (7)$$
$$- \hat{\sigma}^g Pr(q_{jh}|\tilde{q}_{jh}, \omega_h) \log(Pr(q_{jh}|\tilde{q}_{jh}, \omega_h)) \big],$$

where we use our estimates of $r_{\tilde{q}_h, q_h}$ and σ^g from (5) and om from (6). The second line in (7) captures the expected value of the residual, ε^g_{qh}, conditional on choice q_h. We calculate the

[22]The deramping case sums backwards over the relative profits from hour $\underline{h} + 11$ to hour h.

[23]For hours when the generator ramps from minimum to maximum generation, the hours are listed in order, while for hours when the generator ramps from maximum to minimum generation, the hours are reversed so that hour 12 occurs first and hour 1 occurs last.

Figure 3: Increase in Generating Profit Net of Operations and Maintenance Costs Relative to Hour of Ramp

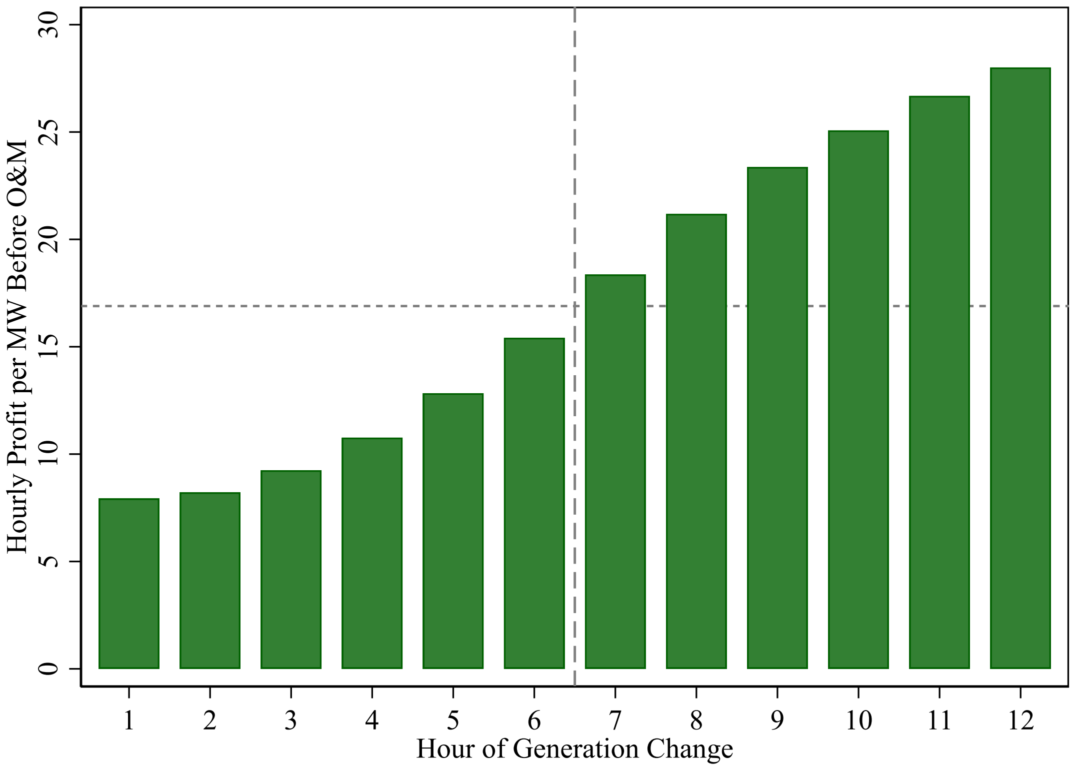

Note: Created from the sample of 12-hour windows around generator ramping events into or out of maximum generation. Hourly profits per MW before operations and maintenance are equal to electricity price minus heat rate times coal price. The horizontal axis represents the number of hours at the *lower* generation level. The vertical line indicates that generators always choose to ramp up in the 7th hour and down in the 6th hour. The horizontal line indicates one potential O&M cost that would rationalize this behavior.

estimated probability of each action at each state, $Pr(q_h|\tilde{q}_h, \omega_h)$, using our estimates from equation (5).

While (7) focuses on profits for one coal generator over one year, we use annual profits in our dynamic model of adoption and exit. Specifically, we need to predict the generator's j expected annual profits across potential states Ω. We perform this prediction by estimating Π_{jt} for each generator-year in our sample and then regressing Π_{jt} on gas fuel price, relative coal capacity, their interaction, and generator j's fixed characteristics, K_j and $heat_j$. Approximating the profit surface in this way imposes a functional form on the impact of these characteristics that allows us to interpolate—and potentially extrapolate—for our estimation and counterfactuals to states that we do not observe in the data. The identification assumption here is that this interpolation is valid.[24] We predict state-contingent annual profits in this way rather than using actions taken in the model, because our approach recovers cost estimates but not the choices in counterfactual situations.

We investigate SO_2 pollution under counterfactual policy environments. We focus on SO_2, since the value of MATS pollution reductions as calculated by the EPA (Environmental Protection Agency, 2011) is dominated by SO_2 and, consistent with our measure of MATS compliance, reductions in other pollutants will likely be approximately proportional to SO_2 reductions. For these simulations, we approximate a pollution surface with a similar regression to our profit regression, but where we use the log of annual pollution as the dependent variable.

4.3 Adoption and Exit Parameters

We estimate our adoption and exit parameters using a full solution, nested-fixed-point approach, where the unit of observation is a generator in a year.[25] Our parameters are the exit scrap value, X, abatement technology adoption cost, A, standard deviation of the unobservable, σ, and the probabilities of MATS enforcement, $P_4, \ldots, P_1$. We further allow A to vary based on whether the generator is subject to U.S. state or MATS enforcement because of the differences between U.S. state air toxics regulations and MATS discussed in Section 2.1.

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.[26]

[24] We also investigate alternative specifications of this regression, such as a logged functional form.

[25] Here, we 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.

[26] We 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 of any year.

We search over values of these structural parameters. For each candidate parameter vector, U.S. state, and belief year, we solve for the MME, which we use to simulate 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 MME. Following Ifrach and Weintraub (2017), we recover the MME as the fixed point of a process that iterates between solving individual optimization decisions and estimating the expectations regarding market evolution that are generated by these decisions.

As discussed in Section 3.1, market evolution 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.

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 from the hourly operations model. 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, 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.

5 Results and Counterfactuals

5.1 Generation Cost Parameter Results

We estimate ramping and O&M costs using our hourly generation model. We estimate ramping costs using (5), separately for generators in 5 capacity bins (0–100MW, 100–300MW, 300–500MW, 500–700MW, and >700MW). Figure 4 reports the ratios of the ramping coefficients, $\widehat{\frac{r}{\sigma^g}}$, to the operating revenue coefficient, $\widehat{\frac{1}{\sigma^g}}$, for all ramping parameters and all but the highest capacity bin.[27] We find that ramping from minimum to maximum generation (represented by red 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 4 further makes clear that ramping costs are very clearly 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.

Table 5 provides further details on the ramping cost estimates for generators between 100 and 300MW of capacity. The first column presents results without any controls for continuation value, while the second column includes our full set of controls (which is what is presented in Figure 4. Controlling for continuation value appears to be extremely important for estimating ramping accurately. With controls, operating revenue is less important by an order of magnitude, and this increases the estimated ramping costs similarly. With controls, we estimate that ramping from off to minimum generation costs these generators $132,000 while ramping from minimum to maximum generation costs $75,000. This is consistent with the literature that starting up from off is particularly costly. Further, we estimate that ramping from 0 straight to maximum generation costs these generators $233,000. Thus, ramping directly from off to maximum generation costs $26,000 more than dividing the ramp across two hours. This is again consistent with the idea that it is costly for generators to

[27] The 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 4: Variation in Ramping Cost Estimates with Capacity

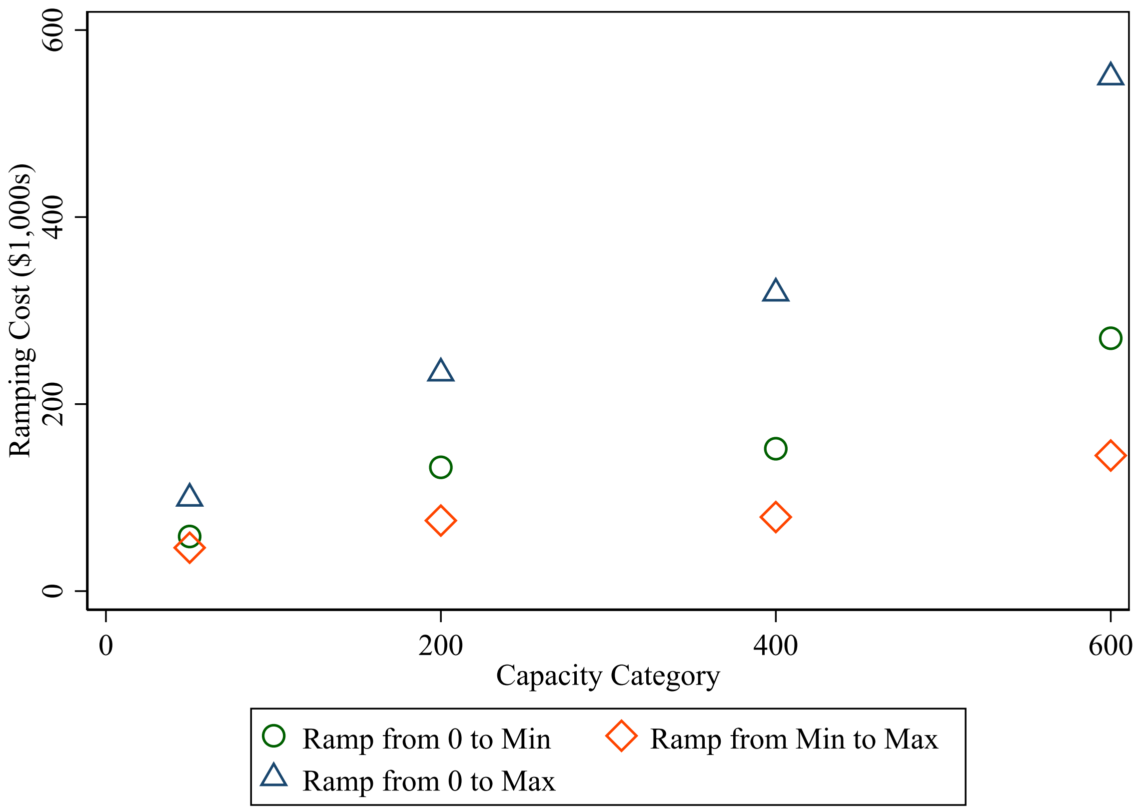

Note: Ramping costs are ratios of estimated ramping coefficient to operating revenue coefficient from separate regressions by capacity bin. Symbols are placed at the midpoints of the capacity bins.

ramp quickly to full generation.

Table 5, Column 3 allows for the possibility that deramping, or reducing generation, is also costly. The results in Column 3 show ramping costs that are substantially lower than Column 2 and deramping costs that are similar in magnitude to ramping costs. However, a comparison of Columns 2 and 3 shows that the total cost of ramping from any level to another level and then deramping back to the original level has approximately the same costs across models. Since the number of ramps will be approximately equal to the number of deramps over a year, this means that our estimates of annual generator profits will be very similar with either of these specifications of ramping costs. We choose to use the second column, without deramping costs, as our base specification.

Table 5: Estimates of Ramping Costs: 100–300MW Capacity

	No Controls	With Controls	With Controls
Operating Rev. (Million \$)	763.87*** (0.46)	76.47*** (0.84)	76.58*** (0.84)
Ramp 0 to Min	−9.03*** (0.01)	−10.13*** (0.01)	−5.21*** (0.18)
Ramp Min to Max	−6.06*** (0.00)	−5.77*** (0.00)	−2.83*** (0.19)
Ramp 0 to Max	−37.82*** (0.16)	−17.82*** (0.06)	−9.97*** (0.21)
Deramp Min to 0	–	–	−4.87*** (0.18)
Deramp Max to Min	–	–	−2.94*** (0.19)
Deramp Max to 0	–	–	−8.74*** (0.21)
N	28,686,219	28,686,219	28,686,219
Pseudo R^2	0.7612	0.8570	0.8571

Note: Regression controls include generation quantity, year relative to first year and its square, fuel price ratio and its square, relative coal capacity, and the interaction of fuel price ratio and relative coal capacity, all interacted with technology and fuel cost per MW, and weather, hour-of-day, month, and current load. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Our ramping cost estimates are in the range of previous U.S. estimates. Kumar et al. (2012) suggest that start-up costs for large coal plants may reach \$500,000, and Cullen (2014) similarly finds extremely large start-up costs. These estimates are larger than those in Reguant (2014), which finds start-up costs of €15-20k for a 150MW coal generator and approximately €30k for a 350MW coal generator in Spain.

Our identification argument for ramping costs relies on our ability to accurately capture expected relative continuation values with our controls. For our ramping cost estimates to be consistent, future prices should not predict current actions, conditional on our controls. In order to provide evidence on the suitability of our controls, we compare future mean electricity prices across generators' actual generation decisions, conditioning on the difference in our predicted net continuation values between maximum and minimum generation.

Specifically, Figure 5 conditions on three ventiles of the relative continuation value distribution and examines how the densities of the future price distribution over the subsequent 20 hours differ across chosen minimum or maximum generation.[28] The green line in each picture shows the future mean distribution of electricity prices for generators who chose to

[28]We chose these ventiles of the future price distribution because we would expect generators to be making a choice between minimum and maximum generation when the relative continuation value is fairly high.

Figure 5: Comparison of Future Electricity Prices By Generation Level

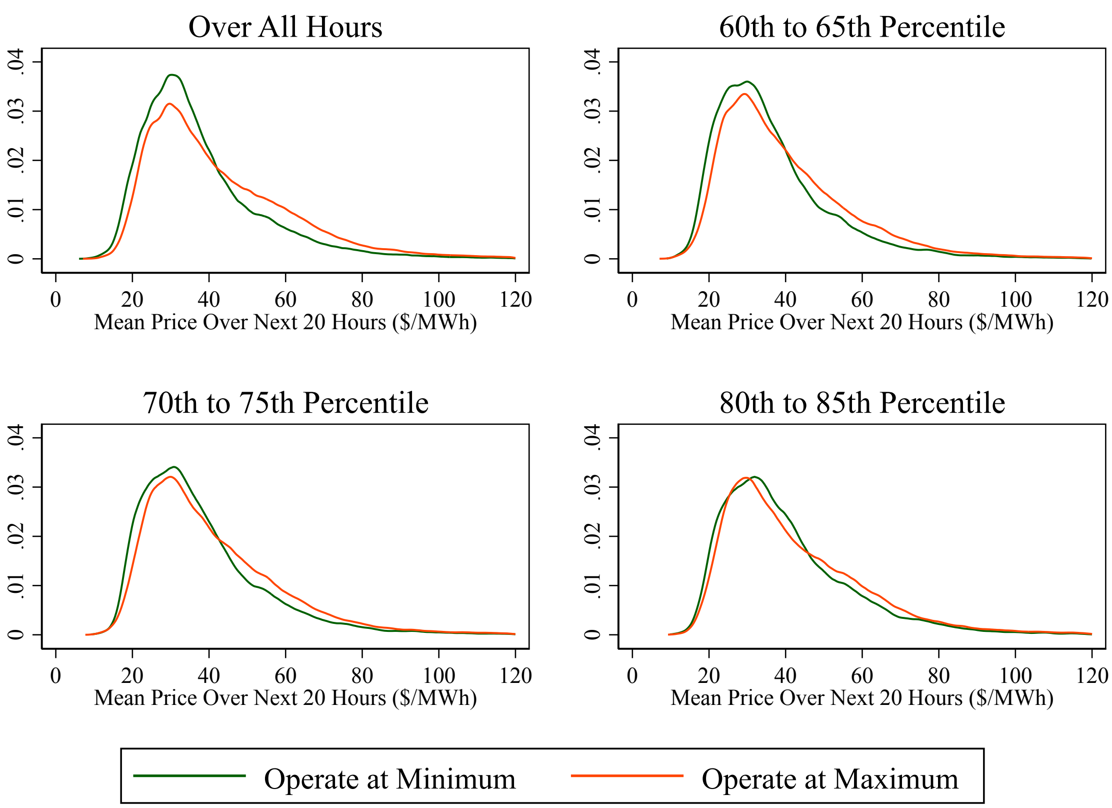

Note: Each panel plots price densities, separately by hours when the generator chooses minimum or maximum generation for generators with 100–300 MW of capacity. The top left panel shows this for all hours. The other panels show hours where the difference in other costs and relative continuation value—ψx—from operating at maximum minus minimum generation is within particular quantiles.

produce at their minimum, and the orange line shows the distribution for generators who chose to produce at their maximum. The top-left panel shows the two distributions unconditional on relative continuation value. As we would expect, when future prices are low, generators are choosing minimum generation more often. However, once we condition on relative continuation value, the lines are quite similar.

We estimate O&M costs from (6). Table 6 presents our estimates, recovered by dividing the cost per additional TW of production by the benefit per million dollars of variable profit. We estimate that generators on average pay O&M costs of $15.18 per MWh of generation, or just over $3,000 for a 200MW generator operating at full capacity. This

Table 6: Estimates From O&M Regression

Benefit per Million Dollars of Variable Profit	10.55^{***} (0.45)
Cost per TW of Additional Production	-160.18^{***} (9.07)
Observations	384,672
Pseudo R^2	0.0043

Note: Multinomial logit regression of choice of number of hours produced within a 12 hour window surrounding each ramping event into or out of maximum generation. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

estimate closely matches the raw data presented in Figure 3, and is close to the EIA's National Energy Modeling System estimates of approximately $14/MWh, which is used in Linn and McCormack (2019).

5.2 Inputs to Adoption and Exit Model

Once we have recovered estimates of ramping and O&M costs, we can calculate the generator profits we observe in the data. The 5$^{\text{th}}$ percentile of calculated profits is -$15 million and the 95$^{\text{th}}$ 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. Online Appendix Figure A1 presents a histogram of these calculated profits.

Our adoption and exit model relies on profit predictions, and our counterfactuals rely on pollution predictions, both conditional on the dynamic state. Table 7 presents the results of our regressions of calculated profits and pollution on dynamic states.

The profit regression R^2 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 (the sum of 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

Table 7: Profit and Pollution Surface Regression Results

	Annual Profit (millions of $)	Log Annual Pollution (lbs of SO_2)
Compliant		-1.285^{***} (0.054)
U.S. State Coal Capacity	-5.505^{***} (1.701)	0.335^{***} (0.064)
Gas to Coal Fuel Price Ratio	9.371^{***} (0.819)	0.798^{***} (0.042)
Interaction of Coal Capacity and Price Ratio	2.681^{***} (0.773)	-0.182^{***} (0.064)
U.S. State Coal Fuel Price ($)	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. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

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, we are less concerned that this will bias our dynamic parameter estimates.

The pollution regression has three differences relative to our profit 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 the 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 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.

Figure 6 shows that profits, as predicted with our regression of calculated profits on dynamic states, follow average calculated profits well over time. The green dashed line shows calculated profits, and the orange dash-dot line shows predicted profits (both in millions of dollars). 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, showing that exit takes time to occur.

Figure 6: Variation in Profits, Coal Capacity, and Fuel Prices Over Time

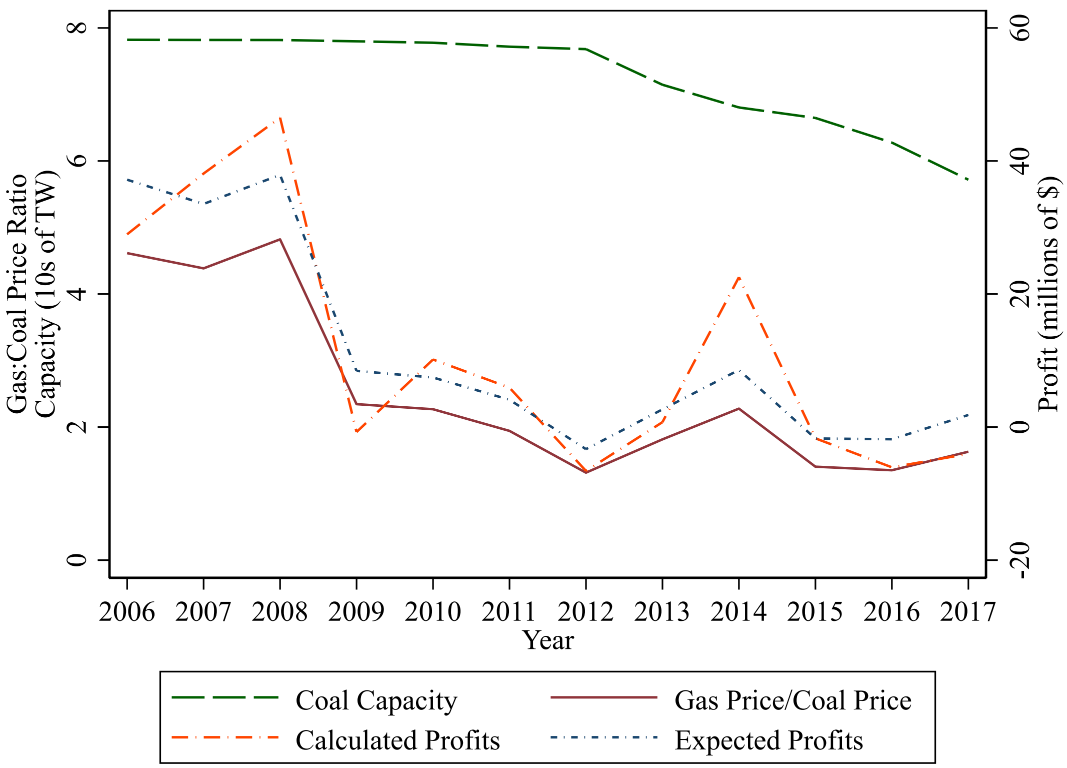

Note:

Table 8: Structural Parameter Results

	Base Specification	Same Adoption Cost State vs. MATS
Generator Costs:		
Adoption Cost (million $)	147.0	413.9
Extra MATS Adoption Cost (million $)	407.1	–
Exit Scrap Value (million $)	−196.4	−196.8
$1/\sigma$ (million $)	64.0	63.4
Predicted Enforcement Probabilities:		
Probability 2012	0.999	0.999
Probability 2013	0.746	0.506
Probability 2014	0.449	0.306
Probability 2015	0.999	0.997
Simulated Log Likelihood	−628.34	−637.88

Note: Costs for first four entries in third column are per GW of capacity.

5.3 Results of Adoption and Exit Model

Table 8 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 policies. We find that adopting air toxics abatement technology costs generators subject to state policies $147 million. Adoption costs generators subject to MATS an additional $407 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 estimates that generators need to pay $196 million in exit costs (negative scrap value) to shut down. These costs likely come from costs such as environmental site remediation that force generators to remove coal trailings from the site and otherwise remediate environmental damages from local coal electricity generation.

Finally, our base specification estimates that in 2012 generators perceived that MATS was essentially certain to be enforced.[29] This probability drops to 75% in 2013 and 45% in 2014 as court challenges to MATS arise. By 2015, however, generators realized that MATS was again very likely to be enforced in 2016, with the probability again rising to nearly one.

The second column of Table 8 requires abatement costs to be identical across generators subject to state policies and MATS. We find similar results, with adoption costs between U.S. state and federal adoption costs from the first column. Though the probability estimates in 2013 and 2014 are smaller 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.[30]

5.4 Counterfactual Results

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 9 presents counterfactual discounted costs, profits, pollution, and exit and adoption outcomes for generators subject to MATS enforcement over the 30 years from 2012-2031.

The first two columns of Table 9 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 \$7.2 billion in abatement technology costs, pay \$19.3 billion in exit costs, and earn \$46.3 billion in profits, discounted and summed over the 30 year period. Our

[29]This very high estimate may reflect the fact that we assume that U.S. state policy passage itself was certain in the announcement year. While these policies were essentially certain once announced, in some cases there was uncertainty surrounding their initial passage.

[30]We 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.

Table 9: Counterfactual Results

	Data	Estimated Model	Same Mean Probability (0.7982)	Uncertainty Resolved in 2012. Enforced with . Prob= 0.7982	No Exit Cost
Adoption Costs (Billion $)		7.22	6.94	6.52	5.01
Exit Costs (Billion $)		19.26	19.16	18.78	0.00
Total Profits (Billion $)		46.34	47.23	48.11	69.42
Pollution (Million lbs. SO_2)		868.24	880.23	942.20	740.63
Number of Generators:					
2012	191	191.0	191.0	191.0	191.0
2013	168	175.9	176.5	175.9	169.5
2014	149	162.8	163.1	161.9	151.3
2015	142	152.4	150.6	148.9	136.8
2016	121	129.5	130.7	130.0	110.8
Count Generators Adopting	14	14.2	13.5	12.74	9.8

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 2031.

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 SO_2 pollution over this period.

The third column of Table 9 takes the average estimated probability (0.7983) from our model and applies it evenly in all years, including in the enforcement year, so that MATS is only actually enforced in 79.83% of simulation draws. This case is roughly similar to the baseline in column 2—as intended—but it allows us to compare ex ante to ex post uncertainty resolution.

Specifically, column 4 of Table 9 similarly assumes a 79.83% 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 policy 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 of policy uncertainty.

With uncertainty resolved at announcement, generator profits are \$880 million higher than when there is uncertainty between MATS announcement and enforcement. Of this, \$440 million comes from lower adoption costs, and \$380 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 $\vec{\varepsilon}_{jt}$ term.

Although removing policy uncertainty would increase generator profits by \$880 million, we find that it would also increase pollution by 62 million pounds of SO_2. This occurs because generators are better able to time their exit decisions to realizations of market forces. Specifically, without policy uncertainty, coal generators remain active more when gas prices are higher, resulting in more operation and hence more pollution. We use the EPA's total pollution benefits range and SO_2 pollution reduction from MATS to calculate the range of benefits from reducing SO_2 and related pollutants by one pound as \$12.41 to \$33.83. This then implies that this pollution increase would cost \$769 million to \$2.097 billion in additional damages. In our content, this pollution increase is critical to understand the welfare impacts of removing policy uncertainty.

Finally, the last column of Table 9 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. While this counterfactual results in fewer generators in the market in 2016, the count is 14.4% lower than in the baseline estimated model in the second column, while pollution drops 14.7%. In tandem, generators' discounted profits rise dramatically, by 49.8% relative to the baseline. This column illustrates how achieving substantial reductions in coal capacity may be very costly to the government.

6 Conclusions and Further Discussion

This paper investigates the level of policy uncertainty surrounding a major U.S. environmental regulation, the Mercury and Air Toxics Standard. To do this, we estimate generators' perception of the probability that MATS will be enforced in 2016 with a dynamic oligopoly model of generator abatement technology adoption and exit behavior. These technology adoption and exit decisions are costly and irreversible, which means that recovering uncertainty from generators' choices requires estimating a structural, dynamic model. Our model is identified by the difference in abatement technology adoption and exit decisions between coal electricity generators facing MATS and the decisions of similar generators facing state air toxics standards.

In order to estimate our model, we also estimate ramping and operations and maintenance (O&M) costs using new approaches and identification arguments. Properly accounting for ramping costs is particularly important during our sample period since large declines in natural gas prices meant that coal generators switched to producing for shorter periods at maximum generation, thereby significantly increasing ramping costs.

We find that there was substantial uncertainty over whether MATS would be enforced, with generators' perceived enforcement probability falling to 45% in 2014, with an average expected enforcement probability of 80% over the 2012-2015 period. In order to understand the impact of this uncertainty on the costs and benefits of MATS, we compare the observed pattern of uncertainty to a counterfactual environment where there is a 80% chance at the moment of MATS announcement that MATS will be enforced, but this uncertainty is resolved instantly for all generators in 2012 with full commitment. We find that uncertainty increases the cost of complying with MATS by $880 billion, but increases pollution costs by $769 million to $2.097 billion.

In the context of MATS, policy uncertainty delayed abatement technology adoption and exit, increased compliance costs, and increased pollution. This occurred because policy uncertainty increased the private costs of compliance, the externalities from generators maintaining option value were negative, and the policy was eventually adopted. In other contexts, these factors may be reversed and policy uncertainty may increase welfare.

Our approach has two key limitations. First, we model each generator as an independent optimizer, not accounting for linkages across generators owned by the same firm. However, firms may coordinate adoption and exit decisions across generators due to strategic considerations. We believe that these strategic considerations are limited in our context because the market for fossil-fuel electricity is relatively unconcentrated for our sample. In particular, we find a mean HHI of 2,209 (with a standard deviation of 1,659) across the states and years that we study. Second, we estimate ramping and O&M costs with the assumption that generators are price takers in hourly electricity markets, though we calculate profits without imposing this restriction.

Overall, our analysis has provided a new way to measure policy uncertainty in an environment with substantial irreversibility and pollution externalities. MATS also provides a suitable context to analyze air toxics enforcement because the uncertainty surrounding enforcement was largely about the timing of eventual enforcement. This simplifies our model by allowing it to focus on estimating generators' perceptions about this probability. Future research could better understand how uncertainty over the probability of policy enactment, the timing of that enactment, and the characteristics of the policy itself affect economic outcomes.

References

Abito, Jose-Miguel, Christopher R Knittel, Konstantinos Metaxoglou, and André Trindade, "Coordinating Separate Markets for Externalities," Technical Report, National Bureau of Economic Research 2018.

Arcidiacono, Peter and Robert A Miller, "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity," *Econometrica*, 2011, *79* (6), 1823–1867.

Baker, Scott R, Nicholas Bloom, and Steven J Davis, "Measuring economic policy uncertainty," *The quarterly journal of economics*, 2016, *131* (4), 1593–1636.

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`.

Cullen, Joseph A, "Dynamic Response to Environmental Regulation in the Electricity Industry," 2014.

Dixit, Robert K, Avinash K Dixit, 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`.

Fowlie, 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," *Harvard Environmental Economics Program*, 2017.

Gowrisankaran, Gautam, Ashley Langer, and Mar Reguant, "Energy Transitions and Electricity Regulation," *Working Paper*, 2022.

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," Technical Report, National Bureau of Economic Research 2020.

Hawthorne, Michael and Tribune staff reporter, "State moves against mercury," *Chicago Tribune*, 2006.

Hotz, V Joseph and Robert A Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

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.

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, Agan D, and Hilleman D, "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," Technical Report, National Bureau of Economic Research 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.

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," Technical Report, Technical report. Arizona State University 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 Computation and Estimation with MME

We estimate our model and conduct counterfactuals by solving for moment-based Markov equilibria across candidate parameter vectors, U.S. states, and belief years. We recover the MME as the fixed point of an algorithm that iterates between solving individual optimization decisions and estimating market evolution regressions that are generated by these decisions.

This appendix begins by detail 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.

State Space

The state space for any generator consists of three continuous states (Ω)—the ratio of natural gas to coal fuel price, coal capacity relative to the 95$^{\text{th}}$ percentile of hourly load, and the capacity-weighted share of IPP coal generators that have adopted abatement technology—and discrete states. 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 $\{\tau, \tau', Tech\}$ based on the belief year $\tau \in \{0,1,2,3,4\}$, years to enforcement $\tau' \in \{0,1,2,3,4\}$ and 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 $\tau' = 1$ and $Tech \in \{0,1\}$, four states for $\tau = 2$ where $\tau' \in \{1,2\}$ and $Tech \in \{0,1\}$, six states for $\tau = 3$ where $\tau' \in \{1,2,3\}$ and $Tech \in \{0,1\}$ and eight states for $\tau = 4$ where $\tau' \in \{1,2,3,4\}$ 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 policy being implemented once it is announced. Thus, belief year τ is not relevant for these U.S. states. For these units, there are $1 + 2 \times$ (Enforcement Year $-$ Announced Year) discrete states.

Generators' Annual Dynamic Optimization

We solve for generators' optimal dynamic adoption and exit decisions using Bellman equations, as exposited in Section 3.2. For each generator, discrete state, bin of the continuous state, 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 $\tau' > 0$, 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 $\tau' = 0$, future adoption share is not relevant.

Market Evolution Regressions

As discussed in Section 4.3, 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 IPP and non-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 generators that have 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 policy 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.

Bootstrap Standard Errors

In order to calculate standard errors for our parameter estimates, we conduct a parametric bootstrap. This involves simulating 50 data sets created from the MME evaluated at the parameter estimates and re-estimating our model on these data sets.

We solve the MME 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, τ', given contemporaneous beliefs, τ, so that $\tau = \tau'$.

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.[31] 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.

Counterfactual Calculations

We calculate counterfactual outcomes using a similar approach to how we simulate data for our bootstrap. This involves solving the MME 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.

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

A2 Tables and Figures

Table A1: Announcement and Enforcement Dates for State Policies

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:

Figure A1: Distribution of Calculated Profits

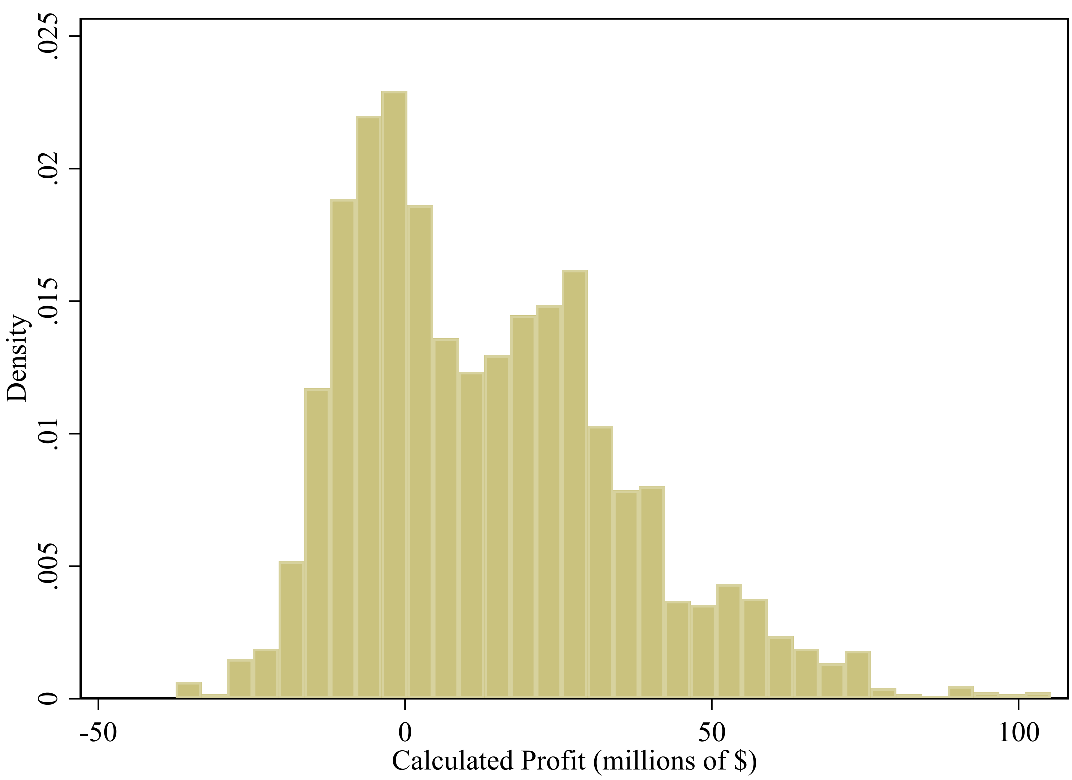

Note: Histogram of annual profits as calculated with Equation 7 and estimated ramping and O&M costs.