

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

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

August 31, 2024

Abstract

Government policy uncertainty affects irreversible decisions including technology adoption and exit. This paper quantifies uncertainty surrounding the Mercury and Air Toxics Standard (MATS). We estimate a dynamic oligopoly model for coal-fired electricity generators that recovers generators' beliefs regarding future MATS enforcement. We develop the Approximate Belief Oligopoly Equilibrium concept where players understand that their decisions impact aggregate market states. MATS enforcement created substantial uncertainty: the perceived enforcement probability dropped to 43%. Resolving uncertainty early would increase profits by \$1.39 billion but also pollution costs by \$0.652 to \$1.776 billion. Had exit been unlikely, resolving uncertainty early would have decreased pollution.

JEL Codes: L13, Q48, Q52

Keywords: environmental policy, air pollution, mercury, MATS, dynamic equilibrium, approximate belief oligopoly equilibrium

*We thank Mike Abito, James Archsmith, John Asker, Nick Buchholz, Allan Collard-Wexler, Inga Deimen, Liran Einav, Jonathan Elliott, David Evans, Sebastian Fleitas, Matt Grennan, Karam Kang, Ryan Kellogg, Robin Lee, Derek Lemoine, Alex MacKay, Bob Pindyck, Louis Preonas, Michael Riordan, Nick Ryan, Stephen Ryan, Gabriel Weintraub, anonymous referees, and seminar and conference participants at multiple institutions for helpful comments. We are grateful to Kiran Dwivedi, Konan Hara, Utkarsh Kumar, Lorelei Loraine, and Filippo Paternollo for excellent research assistance. Gowrisankaran and Langer acknowledge funding from the National Science Foundation (Grant SES-2149335). This work used High Performance Computing (HPC) resources supported by the University of Arizona TRIF, UITS, and Research, Innovation, and Impact (RII) and maintained by the UArizona Research Technologies department. Edited by John Asker.

[†]Columbia University, CEPR, and NBER

[‡]University of Arizona and NBER, alanger@gmail.com

[§]Renmin University of China

1 Introduction

Uncertainty over government policy affects important and irreversible decisions such as firm technology adoption, entry, and exit. In many settings, the process of forming and implementing policies creates uncertainty. For instance, in the U.S., the government frequently enacts new policies by passing legislation that empowers agencies to develop specific regulations. This approach may allow legislation to respond to evolving circumstances and technologies. However, developing regulations takes time, and regulations may be subject to court challenges and executive leadership changes that generate policy uncertainty. This drawn-out process is costly to firms and may affect policy goals, since firms often must make irreversible decisions before knowing whether a policy will be enacted.

This paper measures firms' perceptions of policy uncertainty and simulates how resolving policy uncertainty earlier would affect outcomes. We consider a major pollution regulation in the electricity sector where enforcement was uncertain: the Mercury and Air Toxics Standard (MATS). We estimate players' beliefs regarding the likelihood of MATS ultimately being enforced by modeling generators' technology adoption and exit decisions within a dynamic oligopoly framework. We then use our estimates to simulate how the timing of uncertainty resolution affects counterfactual outcomes in the industry, including generator exit, equilibrium costs, and pollution. Because firms in many sectors—including healthcare, telecommunications, and finance—are oligopolists making irreversible decisions in the face of similarly uncertain policies, our methods and approaches are more broadly applicable.

To illustrate the importance of the timing of uncertainty resolution, consider a coal generator facing a one-time decision of whether to exit the market, in which case it foregoes \$150 million in expected discounted profits. Absent environmental regulation, this generator will not exit. Suppose the Environmental Protection Agency (EPA) issues a rule that the generator needs to install pollution abatement technology—at a cost of \$200 million—to remain active. Parties challenge the rule in court, and observers expect both that the probability the Supreme Court ultimately upholds the rule is 50% and that the Court will not rule until after the generator must make its exit decision (which we call *ex post* uncertainty resolution). Since the generator's net expected profits will now be \$150 million minus a 50%

chance of \$200 million (and a 50% chance of 0), it expects \$50 million in net profits and remains in the market. Consider a policy change that allows the Court to rule before the exit decision (which we call *ex ante* uncertainty resolution). In this case, the generator exits 50% of the time—when the Court upholds the rule—implying that *ex ante* uncertainty resolution increases the probability of exit.

The result that resolving uncertainty early increases the likelihood of the irreversible decision does not always hold. For instance, in contrast to the above, a generator whose expected discounted profits are \$50 million instead of \$150 million will exit with *ex post* uncertainty resolution (since its net expected profits are $-\$50$ million). This generator will again exit 50% of the time with *ex ante* uncertainty resolution, meaning that *ex ante* uncertainty resolution *decreases* the probability of exit here.

A long-standing literature has highlighted that the impact of the *level* of uncertainty on irreversible decisions depends on the curvature of the profit function (???). We look at a related problem: how would resolving uncertainty sooner affect outcomes? Curvature is also key in our application, which generalizes the above example to generators facing a continuous distribution of expected future profits over market outcomes (in addition to policy uncertainty). For generators which are close to the exit margin—similar to the \$50 million expected profit case above—the CDF of expected profits is likely concave, implying that the exit probability will decrease more if the Court decides not to enforce than it will increase if the Court decides to enforce. In this case, as in the example with \$50 million expected profits above, *ex ante* uncertainty resolution will decrease exit in expectation. The opposite will occur for generators which are far from the exit margin.

This result is related to the broader real options literature which emphasizes that uncertainty can increase the option value of delay (????). Our example above isolates the impact of curvature rather than option value, since the generator does not have the option to delay exit in order to gain information. However, as in the real options literature, *ex ante* uncertainty resolution will allow the generator to make its exit choice in response to the actual Court decision rather than its expectation, which will tend to increase expected profits.

Turning to our empirical setting, coal generators (formally, coal-fired electricity generating units) are the primary emitters of air toxics from electricity production. These pollutants,

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

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

Equilibrium effects are potentially important in our context. For instance, one generator’s exit will increase rivals’ profits and decrease their likelihoods of exit. Further, a generator may adopt abatement technology partially to signal a commitment to rivals to remaining in the market (??). Yet the potentially large number of generators in each market results in a curse of dimensionality that makes estimation and the computation of equilibria difficult.

We therefore develop an equilibrium concept called Approximate Belief Oligopoly Equilibrium (ABOE). In an ABOE, all players are oligopolists that compete in a Markov Perfect Equilibrium and recognize that their actions affect market states. However, rather than

keeping track of each competitor’s status, players keep track of market states that provide information about the competitive landscape. Each generator’s beliefs about how these market states evolve—conditional on its own actions—are consistent with competitors’ equilibrium actions as approximated by conditional moments. Thus, in our setting, generators use approximately correct beliefs to make technology adoption and exit decisions.

In our application, the aggregate market states are the coal capacity relative to peak demand, natural gas to coal fuel price ratio, and abatement technology adoption share. Coal capacity and adoption share capture information about expected future profits that occur in the oligopoly context. They also model generators’ potential preemption strategies—of adopting abatement technology and not exiting—for which it is crucial that they understand how their actions affect the market state evolution, as modeled by the ABOE. The fuel price ratio is a major determinant of the expected costs and revenues from electricity sales. We model generators’ beliefs regarding these three state transitions as simple statistical processes that approximate the transitions resulting from rivals’ unobservable cost shocks. Generators understand that transitions of coal capacity and adoption share vary with their own actions.

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

Relationship to the Literature: This paper builds on three main literatures beyond the papers discussed above. First, we extend a recent literature that measures economic and policy uncertainty and evaluates the impact of this uncertainty on economic outcomes. ?, ?, ?, and ? develop measures of uncertainty using financial indicators, newspaper text, SEC filing text, and options prices, respectively. Among other papers, ?, ?, ?, and ? examine the impact of uncertainty on investment, trade flows, and home prices. We add to this literature

by recovering generators’ beliefs over enforcement probabilities and using the estimates to perform counterfactual simulations on the timing of uncertainty resolution.

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

Finally, we extend the literature on dynamic oligopoly models with aggregate market states and approximately correct beliefs. ABOE differs from approximate equilibrium techniques such as Moment-based Markov Equilibrium (MME, Ifrach and Weintraub, 2017) by treating all players as oligopolists who understand that they affect the market states. In an MME, the researcher classifies each actor as either fringe or dominant, actors keep track of each dominant actor’s exact status, and actors may move between these groups.¹ MME and ABOE build on Oblivious Equilibrium (Weintraub et al., 2008) by allowing for aggregate shocks and equilibrium computation when an industry is not in a steady state. These approaches relate to the game theoretic literature that discusses equilibrium convergence when players have approximately correct beliefs (??). Another empirically tractable approach to handling large state spaces and limitations on players’ abilities to process information is Experience Based Equilibrium (EBE, Fershtman and Pakes, 2012; ?). In an EBE, firms learn about their conditional payoffs in different states—which are a function of their rivals’ strategies—from repeated play. MATS only occurs once implying that actors do not visit states repeatedly in our setting, limiting our ability to specify an EBE.

Summary of findings: We estimate that generators’ perception of the probability of MATS enforcement started at 100% in 2012, dropped to as low as 42.8% in 2014, and then rose to almost 100% in 2015, for an average probability of enforcement of 79.13%. Further, we estimate that exit and technology adoption are both costly, with an exit costing the generator

¹Recent empirical MME applications include ?, ?, ?, and ?.

\$327 million, technology adoption to comply with U.S. state standards costing \$312 million, and technology adoption to comply with MATS costing \$883 million. Our model predicts that generators subject to MATS would spend \$11.8 billion on technology adoption and \$30.5 billion on total exit costs, discounted to 2012.

Our counterfactual analyses investigate the impact of resolving policy uncertainty earlier and reducing generator exit costs. We calculate the effect of changing the timing of policy uncertainty resolution by investigating the differences in equilibrium outcomes under a setting where policy implementation is decided in 2016 versus one where there is a commitment in 2012 to whether the standard will be enforced in 2016, both with the mean generator perceived probability of enforcement. Resolving uncertainty immediately would increase expected generator profits by \$1.39 billion in present discounted value. Yet, it also *increases* expected pollution by about 53 million pounds of SO₂ over 30 years, valued between \$652 million and \$1.776 billion dollars, in part by decreasing generator exit. As in the example above where generators expect lower profits, many coal generators were close to the exit margin during the MATS enforcement period, which is why we find that resolving policy uncertainty early leads to less exit and more pollution.

Removing exit costs—for instance by having the government pay for site remediation—reduces the number of generators in 2016 by 15.5% and increases generator profits by 59.6%. Thus subsidizing exit costs increases exit, but requires substantial government transfers to coal generators.

The remainder of this paper is organized as follows. Section ?? discusses the institutional framework, data, and construction of key variables. Section ?? specifies our structural model of technology adoption and exit. Section ?? explains our approach to estimation and identification. Section ?? presents our results and counterfactuals. Finally, Section ?? 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). 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 a maximum achievable control technology (MACT) standard, instead of CAMR which was a voluntary cap-and-trade regulation (?). Although the rule was vacated, our data show that some generators did install mercury abatement technologies during the CAMR period.

At approximately the same time that *New Jersey v. EPA* vacated CAMR, the courts found in *Sierra Club v. EPA* that the EPA would have to regulate mercury and other air toxics together, rather than starting with mercury alone.⁴ In response to these decisions, the EPA finalized MATS in 2012, after releasing earlier versions of the proposed rule in 2011. The final MATS rule required generators to comply with MATS by 2015, but extensions to 2016 were built into the rule and were widely granted.

The investments necessary to achieve compliance with MATS are irreversible and costly, implying that generators may not want to adopt these technologies unless they are fairly certain that compliance will be required. MATS compliance technologies convert pollutants into water-soluble forms, bind them to larger particles, and precipitate the new compounds with a particulate matter catcher.⁵ This basic process can be achieved with different technologies, making it difficult to determine compliance from technology adoption data alone.

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

²<https://www.epa.gov/haps/what-are-hazardous-air-pollutants>.

³517 F.3d 574 (D.C. Cir. 2008).

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

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

administration may change legal approaches.

In the context of MATS, uncertainty arose from both of these sources. The final rule was challenged by several U.S. states' attorneys general. The result of these challenges was that, in 2015, the Supreme Court remanded MATS to the EPA for additional justification that MATS was "appropriate and necessary." However, the order left MATS in place, which effectively meant that generators needed to comply by the 2016 deadline. In 2017, the incoming Trump administration did not file the justification but left MATS in place nonetheless.⁶

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

To define our sample of U.S. states with air toxics standards, we start with the Congressional Research Service report, and then we use a combination of newspaper articles, state environmental agency press releases, state statutes, and our data to create our list of U.S. states with standards and their announcement and enforcement years. On-line Appendix ?? provides details surrounding U.S. state policies, our criteria for classifying the policy of each U.S. state, and a comparison of U.S. state standards to MATS. Table ?? in On-line Appendix ?? lists announcement and enforcement years for these U.S. states.

In contrast to federal regulations, these U.S. state standards were generally subject to very little uncertainty once implemented for at least two reasons. First, in some states (e.g. CT and MD), these standards were passed into law by state legislatures (Pelton, 2006). Second, even in states such as IL and MA where the standards were created as rules issued by the state environmental agency, they were generally developed in tandem with the owners of large coal generators (Hawthorne and Tribune staff reporter, 2006; United Press International, 2004),

⁶In 2021, the Biden Administration did file the justification of MATS.

which led to substantially fewer and weaker judicial challenges. For this reason, we can use the decisions of generators subject to U.S. state standards to identify the costs of generator exit and compliance in the absence of policy uncertainty.

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

Supporting our assumptions on the timing, salience, and importance of MATS to generators not subject to U.S. state standards, many of these generators responded immediately to the 2012 MATS announcement by reporting to the Energy Information Administration (EIA) that they planned to retire. Specifically, at the beginning of 2012, 8.6% of generators subject to MATS newly reported that they would retire between 2012 and 2015. The analogous 2012 increase for generators subject to U.S. state standards was only 1.4%.⁷

Finally, emissions from coal generators were also subject to other pollution regulations during our analysis period. The most important of these is the Cross-State Air Pollution Rule (CSAPR), which regulated SO₂ and NO_x emissions from generators in certain eastern U.S. 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

While our primary analysis data set is at the generator-year level, we calculate annual generator profits using data at the generator-hour level. Fossil fuel power plants are generally made up of a collection of generators that may have different costs, capacities, and abatement technologies. Because of the differences across generators within a plant, we focus on

⁷Calculations from EIA form 860 based on coal generators in Eastern Interconnection.

decisions at the generator, and not plant, level.⁸ Our data sources include information on each generator’s hourly production, costs, emissions, market conditions, demand, prices, and abatement technology adoption.

Our primary data source is the EPA’s Continuous Emissions Monitoring System (CEMS) database. These data are at the generator-hour level. Our sample covers 2006 to 2017 for U.S. states in the Eastern Interconnection, which includes the vast majority of IPP coal generators in the U.S.⁹ Each observation in the CEMS data provides the heat input of the fuel used (in MMBtu), electricity production (in MWh), and CO₂, NO_x, and SO₂ emissions (in pounds/MMBtu) for each generator that the EPA monitors with a CEMS. The CEMS data further report a facility identifier and the location of each generator. As discussed in On-line Appendix ??, we use SO₂ as a proxy to measure the adoption of air toxics abatement technology and to measure annual generator emissions.

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

Second, we merge in annual data from the EIA Form 923, at the facility (or plant) level. We use these data to classify a coal generator as an IPP if the facility to which it belongs is an IPP—as reported by the EIA—at any point in our sample. Form 923 also includes an operator identifier, which we use to understand the potential for generators to jointly optimize (?).

Third, we create and merge in a data set of 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 measures the opportunity cost of fuel faced by generators.

Fourth, we merge in hourly wholesale electricity prices by U.S. state. We obtain prices for

⁸We define generators using the EPA’s definition of units, which is based on emissions release points.

⁹We drop all coal generators in two states. New Jersey adopted a standard in 2003 but likely granted exceptions to some generators, making its policy difficult to classify. All Oklahoma generators appear to enter after MATS was announced.

nodes in each Regional Transmission Organization (RTO) or Independent System Operator (ISO) in the Eastern Interconnection.¹⁰ For some U.S. states, the data report prices for multiple nodes. For these states, we take the mean over the nodes for each hour. For other states, e.g. Georgia, there is no reported electricity price. In these cases, we assign the price from the node that is geographically closest to the state.

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

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

Finally, we use county-level weather data from PRISM.¹¹ 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 ?, we recover daily heating degrees, which measure the amount by which population-weighted state average daily temperature exceeds 65 degrees (if at all), and analogously cooling degrees.

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

2.3 Descriptive Statistics

Table ?? presents generator-level descriptive statistics on our analysis data separately for generators subject to U.S. state air toxics standards and MATS. Our analysis data contain 310 IPP coal generators, of which 93 are subject to U.S. state standards and 217 are subject to MATS. They include a total of 2,791 generator-year observations. The statistics on generator mean capacity levels and heat rates show that there is substantial overlap in the technologies

¹⁰Specifically, we retrieve electricity prices from the New England ISO, New York ISO, PJM, Midcontinent ISO, and the Southern Power Pool.

¹¹The PRISM data can be accessed at <https://zenodo.org/records/10625288>.

of generators across the two sets of U.S. states, though the balance across sets is not perfect. Generators in U.S. states with air toxics standards do face significantly lower average coal fuel prices than generators subject to MATS (\$1.77 vs. \$2.20 per MMBtu), which then feeds into lower marginal fuel costs (\$18.48 vs. \$27.29 per MWh). We account for these differences in our estimation strategy by allowing expected generator profits to differ based on capacity, heat rate, and U.S. state fixed effects.

Table 1: Generator Descriptive Statistics by Regulatory Regime

	U.S. State Standards	MATS
Capacity (MW)	279.69 (213.78)	246.01 (286.83)
Heat Rate (MMBtu/MWh)	10.51 (2.25)	11.99 (4.71)
Coal Fuel Price (\$/MMBtu)	1.77 (0.17)	2.20 (0.39)
Marginal Fuel Costs (\$/MWh)	18.48 (3.93)	27.29 (14.27)
Generators	93	217
Generator-years	841	1950

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

Table 2: Counts of Generators Adopting Abatement Technology or Exiting

Years to Enforcement	U.S. State Standard			MATS		
	Adoptions	Exits	Share Complied	Adoptions	Exits	Share Complied
4	12	0	0.34	2	23	0.30
3	5	1	0.51	1	18	0.53
2	9	0	0.77	2	7	0.64
1	4	4	1.00	9	21	1.00
Total	30	5		14	69	

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

Table ?? reports the number of generators that adopt air toxics abatement technology or exit the market by years to enforcement. For both generators subject to U.S. state standards and those subject to MATS, there is substantial technology adoption and exit before the policy deadline. These “early” actions may occur for a variety of reasons, including that: (1) unrelated maintenance issues may require a generator to stop generating for some time,

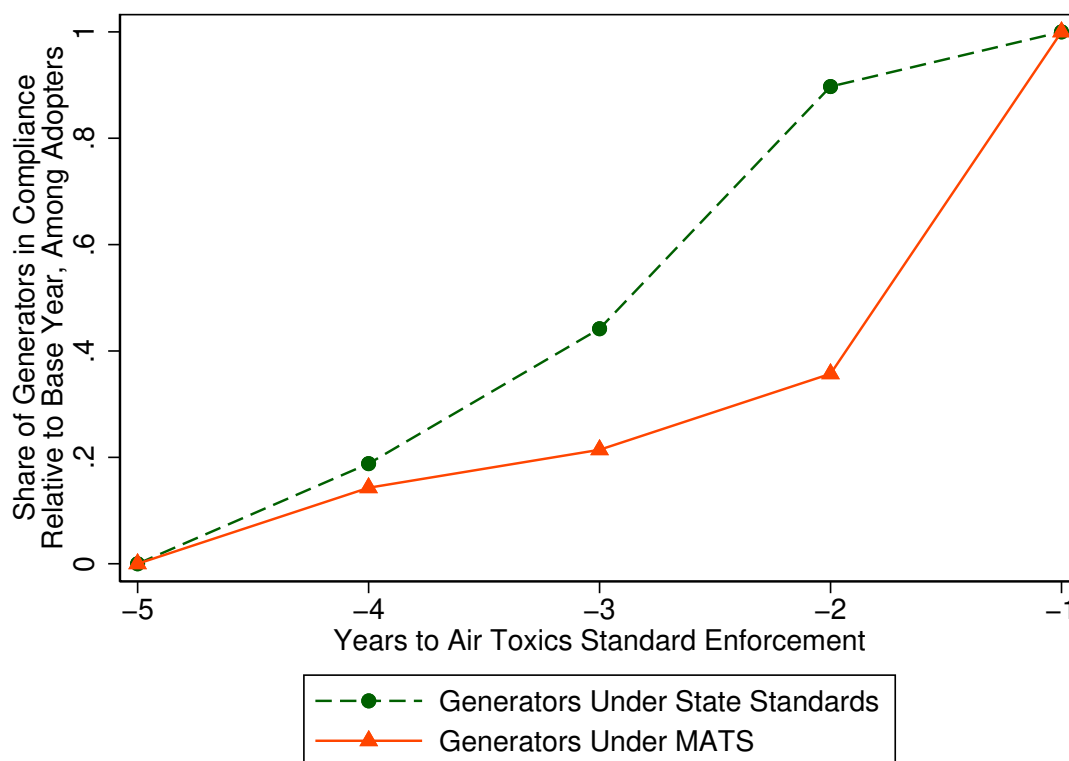
making installing abatement equipment or exiting less costly, (2) idiosyncratic shocks to load may affect both adoption and exit decisions, and (3) technology adoption may signal a commitment to rivals that the generator plans to stay in the market, increasing rivals' propensities to exit. We model the first two forces with idiosyncratic cost shocks, and the third via generators' dynamic decision-making process.

Table ?? further shows that generators subject to U.S. state standards were much more likely to adopt pollution abatement technology than exit, while the reverse is true for generators subject to MATS. The fact that natural gas fuel prices were substantially lower during the MATS enforcement period than earlier led to lower coal generator profits, making exit the preferred option for many. We model this causal pathway by allowing generator profits—and hence the continuation value from adoption—to depend on fuel prices. This pattern may also be partially caused by MATS being stricter than U.S. state standards, which we model by allowing for variation in the compliance costs across the two regimes.

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

Finally, Table ?? presents descriptive statistics on the hourly generation data we use to calculate annual generator profits. Over our sample period, IPP coal generators operated at maximum generation for 45% of hours, minimum generation for 24% of hours, and were off for 31% of hours, for a mean hourly generation level of 177.02 MWh. Generators in our sample faced wholesale electricity market prices with a mean of \$39.38/MWh. Mean hourly

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



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

Table 3: Descriptive Statistics of Hourly Generation

Share of Hours at Maximum Generation	0.45 (0.50)
Share of Hours at Minimum Generation	0.24 (0.43)
Share of Hours at Zero Generation	0.31 (0.46)
Generation (MWh)	177.02 (250.96)
Electricity Price (\$/MWh)	39.38 (25.14)
U.S. State Electricity Demand (GWh)	13.47 (5.37)
Daily Heating Degrees	13.91 (14.91)
Daily Cooling Degrees	2.83 (4.82)
N	25,194,607

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

U.S. state electricity demand was 13.47 GWh, which was driven by a mean of 13.91 heating degree days and 2.83 cooling degree days. All of these variables have substantial variation both over time and across generators.

3 Model

We develop an infinite-horizon dynamic equilibrium model of abatement technology adoption and exit for coal independent power producers (IPPs), which we refer to as “generators” for brevity. Each year, t , there is a set of generators that are currently operating who value the future with discount factor $\beta = 0.95$. Each generator, $j = 1, \dots, J_t$, has a time-invariant heat rate, capacity, and an indicator for whether it has adopted air toxics abatement technology. 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 ?? discusses the state space and equilibrium.

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

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

Second, generators make adoption and exit decisions.¹³ Generators that have not yet adopted abatement technology must decide whether to adopt and pay an adoption cost $A - \varepsilon_{jat}$, continue operating without adopting and receive a payment ε_{jct} , or exit and earn a scrap value $X + \varepsilon_{jxt}$. Technology adoption does not affect generator profits directly, but serves a preemption role by ensuring compliance with standards and credibly committing to rivals that the generator is less likely to exit in the future, thereby potentially encouraging rivals to exit. Generators that have already adopted or are in years before air toxics standard announcement only choose between continuing to operate or exiting.

The cost shocks to generator j , $\vec{\varepsilon}_{jt} \equiv (\varepsilon_{jat}, \varepsilon_{jct}, \varepsilon_{jxt})$, are type 1 extreme value *i.i.d.* across options, years, and generators, and are generator j 's private information at the decision point. These shocks arise as generators may have idiosyncratic maintenance needs or contracts that cause variation in the costs of adopting or exiting. Section ?? details long-run dynamic optimization.

Third, conditional on the technologies and capacities of generators and other sources, generators earn annual profits, Π , from competing in hourly electricity markets. Their annual revenues from selling electricity are the sum of the hourly wholesale electricity market prices times the quantities supplied. Generators bear three types of costs: fuel, ramping, and operation & maintenance (O&M); ramping costs make this problem dynamic. Section ?? describes profit maximization within a year.

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

A limitation of our analysis is that we model each generator as an independent optimizer. We believe that coordination in exit and adoption decisions for generators is limited because most IPPs operate in a single market, most IPPs have few power plants, and most plants have few generators. On-line Appendix ?? provides statistics regarding IPP market structure.

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

3.1 Equilibrium and State Space

As discussed above, equilibrium effects are potentially important in our setting, but keeping track of all rivals' choices results in a curse of dimensionality that substantially complicates equilibrium computation and estimation. Accordingly, we develop and apply a concept called Approximate Belief Oligopoly Equilibrium (ABOE). In an ABOE, each player forms perceptions of market evolution conditional on its actions, and these perceptions are consistent with equilibrium play when fit to a simple functional form. Each player chooses dynamic best responses based on these perceptions, which are an approximation of actual state transitions. Therefore, an ABOE consists of sets of strategies for every player, aggregate market states, and moments on state transitions conditional on that player's actions such that: 1) the state-contingent strategies reflect optimizing behavior given the aggregate states and moments, and 2) simulated state transitions that condition on the actions of the player but otherwise use the candidate ABOE strategies are accurately approximated by these moments.

Our approach is similar to Moment-based Markov Equilibrium (MME). In an MME, players use aggregate states and base their expectations of future states on moments. However, in an MME, each player at any point in time is either dominant or fringe, and only dominant players understand that they can affect the market state. An ABOE differs from an MME in that, in an ABOE, every player recognizes that it can influence the aggregate state. We account for this awareness by allowing the expected distribution of state transitions to condition on a player's own actions.

In environments with aggregated state spaces such as ABOE and MME, optimal strategies are not necessarily Markovian because players may know more about their rivals than is encapsulated by the chosen aggregate states (Ifrach and Weintraub, 2017). In this way, equilibrium methods with aggregated state spaces may result in problems of serially correlated unobservable states (as identified by ?). However, strategies will be Markovian if we further assume that expected state-contingent profits and state transitions are a function only of the aggregate market state and not of any additional information that players may possess. In practice, this assumption requires that the aggregate market state and moments on state transitions capture the critical information that players use to determine strategies.

One could test this assumption by adding further states to the market state space and examining the stability of the estimates (as in Krusell and Smith, 1998). On-line Appendix ?? presents details on different approximate equilibrium concepts used in the literature and differences between them, and Section ?? examines the robustness of our annual profit results to alternative market states.

We specify three aggregate market state variables: (1) coal share, the (combined IPP and non-IPP) coal capacity divided by the 95th percentile of hourly load in the market, (2) the gas-to-coal fuel price ratio, and (3) the share of IPP coal capacity that has adopted abatement technology. Additionally, generator j 's state includes its non-time-varying and time-varying characteristics. The non-time varying characteristics include heat rate, $heat_j$, capacity, K_j , and fixed market characteristics such as the coal price, f^C . The time-varying characteristics include an indicator for having adopted abatement technology, $Tech_{jt}$, the belief year, τ , years to potential air toxics standard enforcement, $\tilde{\tau}$,¹⁴ and its cost shocks, $\vec{\varepsilon}_{jt}$.

A simple conceptual model motivates the use of the first two aggregate market states. Suppose that electricity is generated with either coal or natural gas, the marginal costs of each are determined only by their fuel prices, generators bid their costs in wholesale markets, and natural gas capacity responds quickly to fuel prices. When natural gas fuel prices are high, coal capacity would be dispatched first. Thus, coal generation would be used up to the minimum of coal capacity and load. Hourly wholesale electricity prices would reflect coal marginal costs in hours with low load and gas marginal costs in peak hours.

In contrast, when natural gas fuel prices are low, natural gas would be dispatched first and coal generation would then only be used when load is sufficiently high. Coal generator profits, in either case, could be computed as the sum over hours of coal usage times the wholesale electricity prices net of coal marginal cost.

In this simplified model, profits can then be determined by fuel prices, coal capacity relative to load, and generators' own characteristics.¹⁵ Importantly, the model does not require including natural gas capacity as a state, but instead captures the impact of this

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

¹⁵We include adoption share as an aggregate state to capture generators' preemption motive, which it does by affecting transitions rather than profits directly.

variable on coal generators' profit via the fuel price ratio, though this is a limitation to the extent that natural gas capacity may not respond immediately to changes in the fuel price ratio. While our actual model is much richer than this simplified conceptual model, the simplified model leads us to use coal share and the fuel price ratio to capture the first-order economic forces. Finally, we focus on natural gas as the sole competitor to coal since generation from other fossil fuels (e.g. distillate fuel oil) is more expensive than both coal and natural gas generation, and renewable generation was fairly low during our sample period.

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

We assume that generators believe that the conditional moments of the three aggregate market state transitions reflect simple AR(1) processes.¹⁶ We model coal and adoption shares as autoregressive since they are unchanged until generators take actions, and we follow the literature in modeling fuel prices as evolving autoregressively (?). These processes approximate the combination of exogenous market-level unobservables and the structural unobservables, $\vec{\varepsilon}$, for all generators. Because each generator believes that the market technology adoption share will depend on whether it adopts, we specify different AR(1) regressions for adoption share conditional on that generator's adoption versus non-adoption choice.¹⁷ Because market fundamentals may vary across U.S. state and belief year τ , we further disaggregate our AR(1) regressions to this level. Following the definition of an ABOE, for every U.S. state and belief year, we compute the equilibrium as a fixed point of dynamic optimization decisions and the AR(1) regression coefficients.

¹⁶We specify separate AR(1) regressions with normally distributed residuals for each aggregate state.

¹⁷We only specify one AR(1) regression for coal share since generators that exit no longer value the market state.

3.2 Generator Dynamic Optimization

Generator j makes adoption and exit decisions based on its state, which includes the three aggregated market characteristics noted above, denoted Ω_t , its own time-varying characteristics, $(Tech_{jt}, \tau, \tilde{\tau}, \vec{\varepsilon}_{jt})$, and its non-time-varying characteristics. When $\tilde{\tau} = 0$, generators face a relatively simple choice between exiting or continuing. When $\tilde{\tau} > 0$, generators that have not yet adopted abatement technology face an additional decision of whether to adopt the technology. In this case, the value of continuing depends fundamentally on whether $\tilde{\tau} = 1$ or $\tilde{\tau} > 1$ because in years with $\tilde{\tau} > 1$, there will be no enforcement decision at the end of the year and so the generator will get to continue even if it has not adopted.

Consider first the dynamic decision in a year with $\tilde{\tau} = 1$ (one year from enforcement) in the case where the generator has not previously adopted abatement technology.¹⁸ The generator faces the three choices of continuing without adopting, adopting, or exiting. If it continues without adopting and the air toxics standard is enforced, it will be forced to exit and receive X . We can write the Bellman equation for this case as:

$$\begin{aligned}
 V_j(\Omega, Tech = 0, \tau, \tilde{\tau} = 1, \vec{\varepsilon}_j) = & \quad (1) \\
 \max \{ & \Pi_j(\Omega) + P_\tau X + (1 - P_\tau)\beta E[V_j(\Omega', 0, \tau, 0, \vec{\varepsilon}_j^T) | \Omega, \text{No Standard}] + \sigma \varepsilon_{jc}, \\
 & \Pi_j(\Omega) - A + \beta \{ P_\tau E[V_j(\Omega', 1, \tau, 0, \vec{\varepsilon}_j^T) | \Omega, \text{Standard}] \\
 & + (1 - P_\tau) E[V_j(\Omega', 1, \tau, 0, \vec{\varepsilon}_j^T) | \Omega, \text{No Standard}] \} + \sigma \varepsilon_{ja}, \\
 & \Pi_j(\Omega) + X + \sigma \varepsilon_{jx} \},
 \end{aligned}$$

where a prime indicates next period's value of the variable.

Equation (??) shows that the perceived probability of air toxics standard enforcement enters into the $\tilde{\tau} = 1$ Bellman equation.¹⁹ In (??), 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

¹⁸Because generators believe that their enforcement beliefs will remain constant in the future, we must consider instances where belief year, $\tau > 1$, even though $\tilde{\tau} = 1$.

¹⁹We assume that Π is a function of Ω but not $Tech$, because neither our O&M cost estimates nor observed heat rates vary substantially with technology adoption.

$1 - P_\tau$, the air toxics standard will not be enforced and the generator will never be forced to comply. The second choice is to invest in abatement technology this year, which we indicate by updating $Tech$ to 1 in the future state. In this case, the generator is not forced to exit regardless of whether the air toxics standard is enforced. For both of these choices, the continuation values are dependent upon whether the standard is enforced because enforcement will potentially change the set of rivals remaining in the market. The third choice is exit. Generators that have already adopted face a choice between continuing and exiting that mirrors equation (??), except with no adoption choice and no probability of being forced from the market if the standard is enforced.

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

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

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

3.3 Annual Profits

Generators earn annual profits, $\Pi_j(\Omega)$, by competing in hourly electricity markets over the year. The Bellman equations depend fundamentally on these profits, which vary across the aggregate and individual generator states. As we discussed in Section ??, the forces of competition will mean that two of the three aggregate states—coal share and fuel price ratio—directly affect profits through their impact on market prices. The third aggregate state—technology adoption share—will not directly affect profits, but will affect future state transitions. This is because, in our ABOE, generators that adopt technology recognize that this will increase future adoption share, which may cause generators that have not adopted to be more likely to exit.

Following ?, we assume that hourly profits are revenues net of fuel, ramping, and oper-

ations and maintenance costs. Each year, t , includes hours $h = 1, \dots, H$, with generator j choosing a generation quantity q_{jh} in each of these hours. We specify annual profits as:

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

where p_h is the per-MWh electricity price, f^C is the coal fuel price (so that $\text{heat}_j \times f^C$ are fuel costs per MWh), om is the per-MWh O&M cost, \tilde{q}_{jh} is the generation quantity in the previous hour, $r_{j,\tilde{q},q}$ is the cost to generator j of ramping from \tilde{q} to q , and $\sigma^g \varepsilon_{jhq}^g$ is the idiosyncratic ramping cost shock, with ε_{jhq}^g distributed type 1 extreme value. We model generators as having three choices each hour: (1) generation at capacity, K_j ; (2) minimum generation $L_j K_j$ for $L_j \in (0, 1)$; and (3) not generating (?).

Because of ramping costs, generators' hourly quantity choices affect future profits, thereby making their optimal decisions dynamic, which complicates the model. Nonetheless, we include ramping costs in profits because they have been shown to be important in the context of electricity generation (???) and are particularly critical during our analysis period because of changes in the electricity industry (?). Specifically, the advent of hydraulic fracturing ("fracking") led to sharp declines in the price of natural gas fuel starting around 2009. This, in turn, led to coal generators frequently having fuel and O&M costs above wholesale electricity prices, which increased generator cycling and hence the importance of ramping costs. As an example, in 2008, coal generators in our sample averaged 30.1 hours at their maximum generation level each time they ramped to maximum generation, but this dropped to 19.9 hours in 2017.

4 Identification and Estimation

Our model relies on structural parameters that indicate the probabilities of MATS enforcement, P_τ , exit scrap value, X , abatement technology adoption cost, A , and the standard deviation of the annual cost shock, σ , as well as generators' annual profits from hourly electricity markets. Annual profits are a function of additional estimated structural parameters, specifically, ramping costs, $r_{\tilde{q},q}$, O&M costs, om , and the standard deviation of the ramping

cost shock, σ^g .

While many dynamic models base identification solely on the dynamic decisions (?), we follow ? and ? and identify and estimate state-contingent annual profits from generator actions and observed data within a year rather than from their dynamic discrete choices, and then treat them as observable in the long-run dynamic estimation. We extend this literature by approximating state transitions with an ABOE and approximating profits with a flexible functional form across states. This section focuses first on identification and estimation of the long-run dynamic parameters, which rely upon the assumption that we observe annual profits at every state. We then turn to the parameters that underlie annual profits.

4.1 Adoption and Exit Parameters

Similar to a difference-in-difference model, identification of our long-run dynamic parameters is based on a comparison of actions between generators subject to MATS and those subject to U.S. state air toxics standards, all else equal. Because of the differences between U.S. state standards and MATS discussed in Section ??, we allow A to vary across the two groups. We assume that generators that comply with their U.S. state’s standard do not face additional costs from MATS compliance.²⁰

As in the difference-in-difference literature, identification of these parameters relies on the assumption that generators subject to U.S. state standards would act similarly to those subject to MATS conditional on observables, if the U.S. states had not passed their own standards.²¹ Formally, we assume that mean exit costs are the same for all generators and there is no correlation between unobservable cost shocks, whether a generator is subject to MATS, and the time to enforcement. Thus, identification of the MATS enforcement probabilities reflects the intuition underlying Figure ??: to the extent that generators subject to MATS delay exit and adoption relative to generators subject to U.S. state standards all

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

²¹For instance, many environmental economics papers use this type of difference-in-difference approach to exploit variation in Clean Air Act Amendment attainment status across counties (e.g., ??).

else equal, our model will recover lower estimates of their enforcement probabilities. Since we assume that we observe annual profits at every state, we do not include a parameter on profits in the dynamic decision, which allows us to identify the scale parameter, σ , in dollars.

Accordingly, we need to account for similar threats to identification as that literature. Most importantly, almost all U.S. states' enforcement windows occurred before MATS. As we discussed in Section ??, this is a period when natural gas fuel prices were lower than in the MATS enforcement window, which might have led to more exit and less adoption. Our structural model exploits economic theory to capture the fact that these differences impact decision-making solely through current and future expected profits and their effect on equilibrium.

There are, however, two key differences between our setting and the difference-in-difference literature. First, we observe two related generator decisions, adoption and exit. Our approach allows us to model these two decisions jointly. The value of X and the values of A across the two groups are identified by the state-contingent rates at which generators choose each action given their expected profits.

Second, we do not include year fixed effects, which would control for macroeconomic shocks. Since enforcement windows were the same for all generators subject to MATS and largely did not overlap across the two groups, year fixed effects would not be well identified. Moreover, counterfactuals would be difficult to implement in a model with year fixed effects, since they would require assumptions on the values of those fixed effects out of sample. However, the fact that enforcement windows varied across U.S. states with standards limits the bias from macroeconomic shocks on the exit and adoption cost parameters.

Another limitation regarding our identification strategy is that we do not include annual fixed costs of operation, since these are difficult to identify separately from exit costs (?). Thus, our estimates of X will capture the present discounted value of fixed costs of operation net of any capacity payments. Although variation in fixed costs across U.S. states could lead to concerns about bias, our flexible profit surface regression—discussed in Section ?? below—mitigates these concerns.

Turning to estimation, we use a full-solution simulated maximum likelihood nested-fixed-

point approach where the unit of observation is a generator in a year.²² We search over values of the long-run dynamic structural parameters, limiting the probabilities of MATS enforcement to lie between 0 and 1. For each candidate parameter vector, U.S. state, and belief year, we solve for an ABOE, which we use to calculate a likelihood. The likelihood for any generator and year is the equilibrium probability of the observed action, evaluated at the ABOE. For generators in the years between air toxics standard announcement and enforcement which have not yet adopted, potential actions include continuing to operate, exiting, and adopting. For generators in other years, there are two potential actions: continuing to operate and exiting.²³

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

4.2 Annual Profit and Pollution Surfaces

Identification and estimation of the ramping and O&M cost parameters follows ?. When estimating ramping costs, we group generators into capacity bins and perform the estimation separately for each bin. These estimated parameters allow us to calculate profits at any

²²We 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 .

²³We 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.

observed state as:

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

where we indicate our estimated parameters with hats.²⁴ The second line in (4) captures the expected value of the ramping cost shock, ε_{jhq}^g , conditional on choice q_{jh} . We calculate the probability of each action at each hour, $Pr(q_{jh}|\tilde{q}_{jh})$, from our cost estimation.

As discussed above, the estimation of our long-run dynamic parameters leverages the assumption that we observe profits at every long-run dynamic state, $\Pi_j(\Omega)$. But, as is common in dynamic models, our data contain only a subset of potential states. Since evaluating equilibrium dynamic oligopoly behavior in wholesale electricity markets for counterfactual states is both complicated and requires many additional assumptions, we instead take a more tractable approach. Specifically, we approximate profits across states with a profit surface generated from a flexible regression of profits on functions of the states. This approach is broadly consistent with multiple models of competition, including Cournot and dynamic pricing, but does not impose a particular model. We then use the predicted value of profits at each state in our estimation of our long-run dynamic parameters.

This approach requires that the profit surface accurately captures how expected profits would vary as a generator moves between states, as well as the differences in expected profits between generators subject to U.S. state standards and MATS. In order to use variation in predicted profits that stems from variation in actions across generators rather than unobservable differences in the underlying characteristics across generators or markets, we include U.S. state fixed effects in our profit surface. Within a U.S. state, there is both variation over time (e.g. in the fuel price ratio) and across generators (e.g. in capacity). Our functional form exploits both of these sources of variation and the interaction between them to account for the fact that the impact on profits of these time-varying regressors may depend on the characteristics of the generator. This allows us to capture complex oligopoly interactions without fully specifying the intensity of competition across U.S. states and years. Finally,

²⁴We add the “j” subscript to ramping costs to indicate their variation across capacity bins.

since we are interested in how profits vary across model states, we weight observations in the profit surface regression by the inverse of the number of generators in each U.S. state-year. This means that each U.S. state and year contributes equally to the estimates, ensuring that our coefficients are not overly affected by U.S. states and years that have many generators. We thus estimate the profit surface with weighted ordinary least squares.

For our counterfactuals, we also want to understand how SO₂ pollution would vary across policy environments. We focus on SO₂, since the value of MATS pollution reductions as calculated by the EPA (?) is dominated by SO₂ and, consistent with our measure of MATS compliance, reductions in other pollutants will likely be approximately proportional to SO₂ reductions. We estimate a pollution surface with a similar weighted ordinary least squares specification as our profit regression, but where we include an indicator for the generator having adopted abatement technology, log the dependent variable, annual pollution, and the independent variables other than adoption, and only include generator-years with non-zero pollution.

5 Results and Counterfactuals


5.1 Cost, Profit, and Pollution Estimates

We estimate ramping and O&M costs following ?, separately for generators in 5 capacity bins (0–100MW, 100–300MW, 300–500MW, 500–700MW, and >700MW). Figure ?? reports the ratios of the ramping coefficients, $\frac{\hat{r}}{\hat{\sigma g}}$, to the operating revenue coefficient, $\frac{\hat{1}}{\hat{\sigma g}}$, for all ramping parameters and all but the highest capacity bin.²⁵ We find that ramping from minimum to maximum generation (represented by orange diamonds) is the least costly, ramping from off to minimum generation (green circles) is more costly, and ramping directly from off to maximum generation in one hour (blue triangles) is the most costly. The fact that ramping from zero to minimum is more costly than from minimum to maximum likely reflects startup costs, which are captured by our ramping cost estimates. Figure ?? further makes clear that ramping costs are strongly increasing in generator capacity. In fact, the largest generators

²⁵The highest capacity bin has ramping estimates that follow the same pattern as the smaller bins, but have a substantially larger scale. We do not show the results for this bin for scaling reasons.

have extremely high ramping costs, with ramping from off to minimum generation costing approximately \$1.05 million.

Figure 2: Variation in Ramping Cost Estimates with Capacity



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

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

Turning to operations and maintenance costs, we find that generators on average pay

O&M costs of \$15.55 per MWh of generation, corresponding to a little over \$3,000 for a 200MW generator operating at full capacity. This estimate is close to the EIA’s National Energy Modeling System estimates of approximately \$14/MWh, which is used in ?.

Once we have recovered estimates of ramping and O&M costs, we calculate the generator profits for generator-years in our analysis data. The 5th percentile of calculated profits is -\$7.9 million and the 95th is \$99 million, with 62% of annual profits above zero. The sizable share of calculated profits that are below zero suggests that there may be substantial option value to remaining in operation or sizable exit costs. Figure ?? in On-line Appendix ?? presents a histogram of these calculated profits.

Table 4: Profit and Pollution Surface Regression Results

	Annual Profit (millions of \$)	Log Annual Pollution (lbs. of SO ₂)
Compliant		-0.758*** (0.063)
U.S. State Coal Share	-21.984*** (6.266)	2.155*** (0.275)
U.S. State Coal Share Squared	10.777*** (1.953)	-0.129* (0.076)
Gas-to-Coal Fuel Price Ratio	5.181*** (0.873)	0.467* (0.250)
Gas-to-Coal Fuel Price Ratio Squared	-0.625*** (0.097)	0.145** (0.070)
Coal Share × Fuel Price Ratio	-1.141*** (0.327)	0.127** (0.055)
Heat Rate (MMBtu/MW)	-0.313*** (0.113)	0.582** (0.234)
Capacity (MW)	-0.035*** (0.005)	-0.389 (0.244)
Capacity Squared (MW ²)	-2.3e-05*** (4.0e-06)	0.112*** (0.021)
Capacity × U.S. State Coal Share	-0.007*** (0.002)	-0.032 (0.026)
Capacity × Fuel Price Ratio	0.048*** (0.001)	-0.005 (0.039)
U.S. State FE	Y	Y
Observations	2899	2718
R^2	0.7791	0.7001

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

Turning to our estimated profit and pollution surfaces, Table ?? presents the results of our regressions of calculated profits and pollution on functions of model states. The profit regression R^2 is relatively high at 0.7791, implying that the profit surface captures much of the

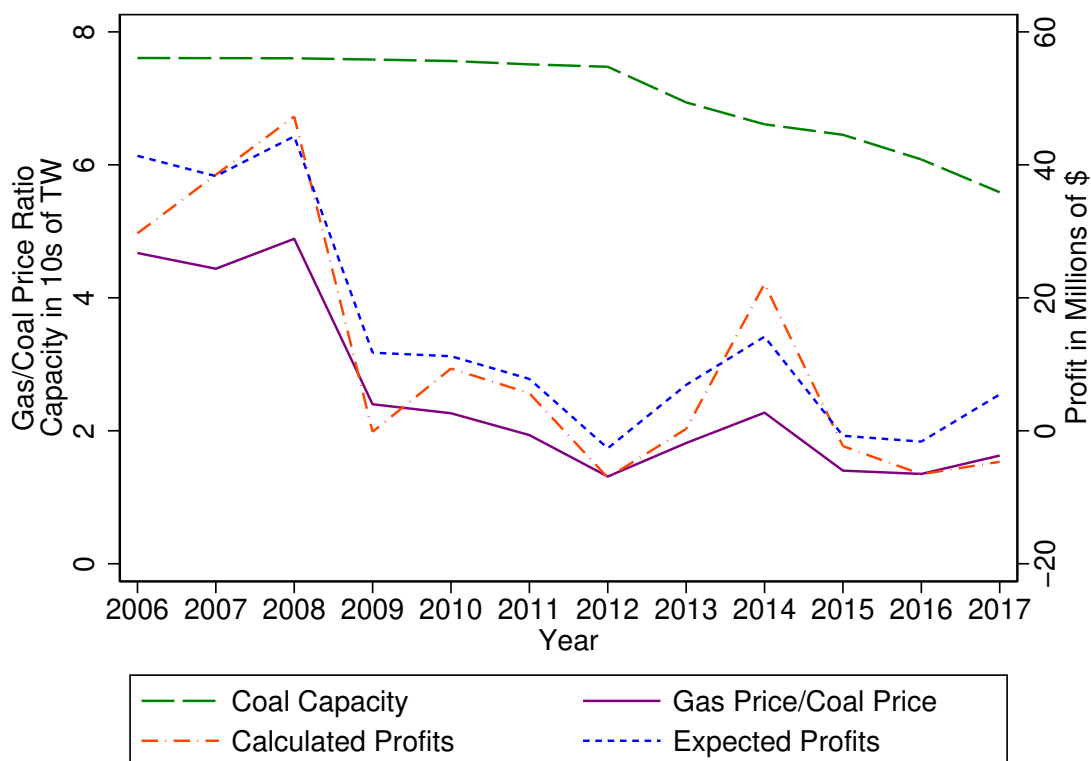
variation in profits across states. The estimated coefficients show that generators with higher heat rates—that burn coal less efficiently—have lower profits. For 92.7% of observations in our data, increasing capacity leads to higher profits.

Our profit surface is also broadly consistent with the implications of oligopoly interactions. Specifically, we find that increasing coal share reduces profits for 91.7% of observations, and increasing the fuel price ratio increases profits for 97.6% of observations. Further, the coefficients on the interactions between these variables and generator capacity indicate that larger generators’ profits react more strongly to these market forces.

For an ABOE to well-approximate a Markov Perfect Equilibrium, the profit surface should capture the critical determinants of profits. To examine the robustness of our profit surface, Table ?? in On-line Appendix ?? presents additional specifications. Column (1) reproduces our base results, while column (2) further explores the impact of market forces on profits, by adding the HHI based on capacities of IPP coal operators and the share of market capacity controlled by the generator’s operator. We find that the coefficient on HHI is both statistically and economically insignificant. The coefficient on the fraction of coal capacity operated by the same operator within the U.S. state is positive and statistically significant. However, including this variable only increases the R^2 of the profit surface regression from 0.7791 to 0.7802, suggesting that our base profit surface captures how market forces affect generator profits reasonably well. Column (3) adds an indicator for the generator being in compliance to our base specification. We find that the coefficient is statistically and economically insignificant. Finally, column (4) removes U.S. state fixed effects from our base specification. The R^2 drops to 0.6071 implying that our remaining regressors explain a substantial portion of the variation in profits.

Figure ?? shows that expected mean profits—as predicted with our regression of calculated profits on model states—follow calculated actual profits well over time. The solid purple line shows the gas-to-coal price ratio, which indicates that when gas prices dropped, coal generator profits also dropped substantially. The green long-dashed line shows that coal capacity started falling approximately two years after the fall in gas prices, with gradual exit continuing throughout the rest of the sample. Our structural model will attribute this pattern to idiosyncratic shocks to the value of exit.

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



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

The pollution surface regression, presented in column 2 of Table ?? has a relatively high R^2 of 0.7001, and all of the coefficients have the expected signs. Specifically, higher U.S. state gas-to-coal price ratio and generator heat rate and capacity are all associated with higher pollution for all observations in our data. We would expect coal generators to run more when gas prices are higher and hence generate more pollution. Similarly, higher heat rate and capacity generators will both use more coal for a given market state and hence generate more pollution, all else equal.

5.2 Results of Adoption and Exit Model

Table ?? displays our structural results for two potential specifications. The first column presents our base specification, where generators subject to MATS are allowed to have dif-

Table 5: Structural Parameter Results

	Base Specification	Same Adoption Cost State vs. MATS
Predicted Enforcement Probabilities:		
Probability 2012	1.000*** (0.082)	0.989*** (0.058)
Probability 2013	0.738*** (0.150)	0.623*** (0.128)
Probability 2014	0.428** (0.179)	0.297** (0.130)
Probability 2015	1.000*** (0.072)	0.999*** (0.072)
Generator Costs:		
Adoption Cost (million \$)	312.2*** (113.1)	694.5*** (90.7)
Extra MATS Adoption Cost (million \$)	570.6*** (72.9)	–
Exit Scrap Value (million \$)	–326.5*** (75.5)	–313.8*** (68.6)
1/ σ (million \$)	40.3*** (4.4)	41.2*** (5.4)
Simulated Log Likelihood	–616.97	–624.76

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

ferent adoption costs than generators subject to U.S. state standards. We find that, in 2012, generators perceived that MATS was essentially certain to be enforced.

This probability drops to 74% in 2013 and to 43% in 2014 when the U.S Supreme Court agreed to hear arguments in the MATS case (*Michigan v. EPA*). By 2015, however, generators realized that MATS was again very likely to be enforced in 2016, with the probability again rising to nearly one. As shown previously in Table ??, generators subject to MATS disproportionately exited or adopted in both 2012 and 2015 compared to 2013 and 2014. This is reflected in our perceived probability estimates. The high 2012 probability may reflect the fact that we assume that the U.S. state enforcement probabilities were 1 in all years, but these enforcement probabilities may only have been very close to 1 in the years *after* announcement, with substantial uncertainty in the year of announcement; see On-line Appendix ??. We therefore interpret our 2012 probability estimate with caution. The remaining estimates, however, suggest that generators' actions revealed substantial uncertainty surrounding 2016 MATS enforcement.

Beyond generators' perceived probabilities of enforcement, we find that adopting air toxics abatement technology costs \$312 million for generators subject to U.S. state standards. Adoption costs an additional \$571 million for generators subject to MATS, 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 was approximately \$17 million. We further estimate that generators needed to pay \$327 million in exit costs (negative scrap value) to shut down. One study has found decommissioning costs for coal generators of up to \$466,000 per MW, which is similar but somewhat smaller than our estimates (?). Our estimates are therefore consistent with substantial site remediation costs, and true exit costs may be even higher if our estimated exit costs include avoided fixed operating costs (?).

The second column of Table ?? requires abatement costs to be identical across generators subject to U.S. state standards and MATS. We find similar results, with adoption costs between U.S. state and MATS adoption costs from the first column. Though the probability estimates in 2013 and 2014 are smaller than in our baseline, the temporal patterns are quite similar, starting at approximately one in 2012, falling in 2013 and again in 2014, and then recovering to nearly one.²⁶

Turning to model fit, columns 1 and 2 of Table ?? 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 14.5 generators would adopt abatement technology while 14 adopted in practice. We underpredict the exit rate by about 7 percentage points, which may be due to the estimation sample being different from our counterfactual sample.

5.3 Counterfactual Results

We use the structural parameter estimates from our base specification to simulate a series of counterfactuals. In each of the counterfactuals, we solve for a new ABOE by re-solving for the

²⁶We 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 6: Model Fit and Counterfactual Results

	Data	Estimated Model	Same Mean Prob.: 0.7913	Uncertainty Resolved in 2012 w/ Mean Prob.	No Exit Cost
	(1)	(2)	(3)	(4)	(5)
Adoption Costs (Bill. \$)		11.77	11.28	10.62	7.93
Exit Costs (Bill. \$)		30.47	30.27	29.59	0.00
Total Profits (Bill. \$)		64.42	65.96	67.35	102.79
Pollution (Mill. lbs. SO ₂)		785.72	793.75	846.26	575.10
Number of Generators:					
2012	183	183.0	183.0	183.0	183.0
2013	160	169.8	170.4	170.6	163.4
2014	142	158.1	158.3	158.8	146.4
2015	135	148.7	146.7	147.5	132.7
2016	114	125.9	127.2	131.4	106.4
Count Adopting	14	14.5	13.8	12.98	9.7

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 expected outcomes with uncertainty completely resolved in 2012, with the mean estimated probability across years. Column 5 sets exit costs to 0. First four rows of results report the total discounted profits or costs from 2012 through 2041.

fixed point between generators' expectations of the evolution of the market state and their adoption and exit decisions. Columns 2 through 5 of Table ?? present the estimated model and counterfactual discounted costs, profits, pollution, and exit and adoption outcomes for generators subject to MATS enforcement over the 30 years from 2012-2041.²⁷

Column 2 of Table ?? shows that, in our estimated model, generators will pay \$11.8 billion in abatement technology costs, pay \$30.5 billion in exit costs, and earn \$64.4 billion in profits, discounted and summed over the 30 year period. Our estimates of generator abatement costs are similar to the EPA's ex ante estimates of compliance costs, which were \$9.6 billion (?).

²⁷Our counterfactuals change MATS enforcement for IPP coal generators only. We hold the evolution process for non-IPP coal share constant since those generators face very different incentives (?).

Generators will also produce 786 million discounted pounds of SO₂ pollution over this period.

Column 3 of Table ?? takes the average estimated probability (0.7913) from our model and applies it evenly in all years, including in the enforcement year, so that MATS is enforced in 79.13% of simulation draws. We include this case because it allows us to consider dynamics in a simplified baseline setting where uncertainty is resolved ex post, but the probability of enforcement is unchanged for four years. This case, which we use to compare to column 4, is roughly similar to column 2, as intended. However, unlike the baseline model, MATS is only enforced 79.13% instead of 100% of the time.

Column 4 of Table ?? considers ex ante uncertainty resolution, again assuming a 79.13% probability of MATS enforcement, but with enforcement now decided randomly at the moment that MATS is announced in 2012. This means that the *level* of the standard is identical to the counterfactual in column 3 in expectation, but that there is no uncertainty after announcement over whether MATS will be enforced. Thus the comparison between columns 3 and 4 provides a clear description of the costs and benefits of resolving policy uncertainty earlier, keeping the expected enforcement probability the same.

With uncertainty resolved at announcement, generator profits are \$1.39 billion higher than when the uncertainty is resolved in 2016. Of this, \$660 million comes from lower adoption costs, and \$680 million comes from lower exit costs. The remainder of the savings accrues 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.

While resolving uncertainty in 2012 increases expected profits by 2.1%, it also increases the number of generators remaining in 2016 by 3.3%, and increases pollution by 6.6%. The EPA's regulatory impact analysis for MATS implies a social cost of SO₂ emissions between \$12.41 and \$33.83 per pound.²⁸ Using this estimate of damages, we find that resolving uncertainty ex ante would result in damages of \$652 million to \$1.776 billion via the 52.5 million pound increase in pollution. This pollution increase is therefore critical to understanding the welfare impacts of removing policy uncertainty.

²⁸? finds MATS co-benefits (which are the predominant source of benefits) of between \$33 and \$90 billion and SO₂ emissions reductions of 1.33 million tons. Dividing the benefits by the reductions yields benefits per pound of \$12.41–\$33.83.

Finally, column 5 of Table ?? 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 total adoption costs, since some generators choose to exit rather than adopt pollution abatement technologies. This counterfactual results in 15.5% fewer generators in the market in 2016 and a drop in pollution of 26.8% relative to the baseline in the second column. In tandem, generators' discounted profits rise dramatically, by 59.6%. This column illustrates how achieving substantial reductions in coal capacity may be very costly to the government.

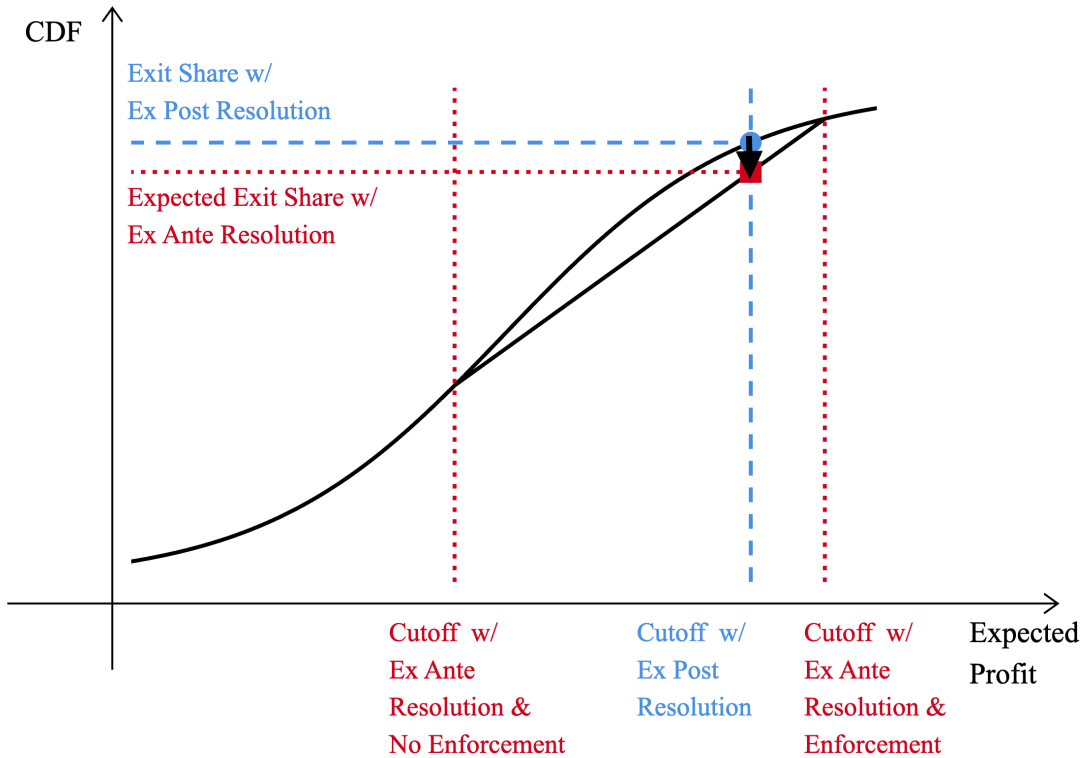
5.4 Understanding the Impacts of Early Uncertainty Resolution

In order to better understand our counterfactual findings regarding ex ante uncertainty resolution, Figure ?? focuses on a conceptual cumulative distribution function of a generator's random shocks between 2012 and 2016 that affect its exit decision. This builds on the Supreme Court example from the introduction, but differs in two ways. First, we use our average estimated MATS enforcement probability, 79.13%, rather than 50%. Second, we generalize the example to a continuous distribution of expected profit to capture market uncertainty such as changing gas fuel prices. We focus on exit instead of adoption because exit was the primary means of compliance for generators subject to MATS (as shown in Table ??).

The blue dashed vertical line in Figure ?? represents the minimum profit shock that would lead the generator to remain in the market given the current policy with uncertainty resolution *after* the generator makes its exit decision (ex post resolution). The blue horizontal dashed line indicates the corresponding probability of exit with ex post uncertainty resolution.

If, however, uncertainty were resolved *before* the generator needs to make its exit decision (ex ante resolution), then the generator would have two cutoffs: one if it knows that it needs to comply and another if it knows that it does not. We represent these two cutoffs with red dotted vertical lines. The dotted line on the right represents the cutoff with enforcement, and the line on the left represents the cutoff without enforcement. We draw these lines so that the blue line is approximately 80% of the way towards the right line. The red dotted horizontal

Figure 4: Conceptual Expected Profit Cumulative Distribution Function



Note: Conceptual cumulative distribution function for a single generator’s expected profits in a given year. Cutoffs indicate the minimum expected profit necessary for the generator to remain in the market. Thus, the CDF represents the share of expected profit draws that are below a given cutoff and lead to generator exit.

line represents the expected exit probability when viewed from before the uncertainty is resolved, knowing that the uncertainty will be resolved ex ante. This probability is the weighted average of the exit probabilities from the two dotted vertical lines.

Our counterfactuals show that the exit probability with ex ante uncertainty resolution is lower than with ex post uncertainty resolution (represented by the downward arrow from the blue dot to the red square). As we can see in Figure ??, this occurs when exit is likely. Many generators being close to the exit margin is consistent with the very low natural gas prices during this period, which reduced annual profits for many coal generators to close to or below 0, as shown in Figure ?? in On-line Appendix ??. Alternatively, if generators had seen substantially greater value to remaining in the market, then resolving uncertainty ex ante would *increase* the probability of exit. This dichotomy mirrors the example in Section ??, but generalizes it to a continuous distribution of expected profits.

Regardless of the likelihood of exit, earlier uncertainty resolution will increase generator profits by allowing them to better respond to random shocks. With a single-peaked distribution of exit shocks (i.e., S-shaped exit shock CDF), it will also attenuate extreme outcomes in expectation. When exit probabilities are high (on the concave portion of the profit CDF) it will decrease them and when they are low (on the convex portion) it will increase them.

6 Conclusions

This paper makes three main contributions. First, we provide a new approach to measuring the level of policy uncertainty and use this approach to recover coal electricity generators' perceptions of policy uncertainty surrounding a major U.S. environmental regulation, the Mercury and Air Toxics Standard. We identify the level of policy uncertainty using the difference in abatement technology adoption and exit decisions between generators facing MATS and those facing U.S. state air toxics standards, all else equal. Second, to implement our approach, we estimate a dynamic model of generators' technology adoption and exit decisions for which we develop a new equilibrium concept that we call Approximate Belief Oligopoly Equilibrium (ABOE). This concept allows us to capture equilibrium effects and avoid the curse of dimensionality that would stem from keeping track of the actions of the many generators within a market. Finally, we use our estimated model to investigate how resolving uncertainty earlier would affect outcomes.

We find that there was substantial uncertainty over whether MATS would be enforced, with generators' perceived enforcement probability falling to 43% in 2014 when the U.S. Supreme Court agreed to hear arguments in the MATS case. The average expected enforcement probability was approximately 80% over the 2012-2015 period. In order to understand the impact of the timing of uncertainty resolution, we compare the observed pattern of uncertainty to a counterfactual environment where there is the same approximately 80% chance at the moment of MATS announcement that MATS would be enforced in 2016, but this uncertainty is resolved instantly in 2012. We find that resolving uncertainty earlier decreases the cost of complying with MATS by \$1.39 billion but increases pollution damages by \$652 million to \$1.776 billion as it causes more generators to remain in the market in expectation.

Our new equilibrium concept, ABOE, contributes to a growing literature that uses approximations to make the estimation of dynamic games more tractable. Even with approximate equilibrium approaches, estimation of dynamic games still requires limiting the dimension of the underlying state space so that the approximations are credible. For instance, we make some assumptions—notably that generators do not coordinate when making decisions—that would be hard to generalize because doing so would require a much more complicated underlying state and action space that would then be hard to credibly approximate. In contrast, assumptions that do not expand the underlying state space—such as allowing technology adoption costs to vary by generator age or allowing profits to depend on technology adoption—could be implemented if we had sufficient data to identify these interactions. Future research on dynamic approximation methods could explore the selection process for aggregate states; the variation in estimates across different approximate equilibrium concepts; and the convergence properties of approximate equilibrium solutions to Markov Perfect Equilibria as the state space grows.

Beyond our contribution to dynamic equilibrium estimation, this paper highlights the importance of the curvature of expected profits in determining the impact of uncertainty in the context of changing the timing of uncertainty resolution. In our setting, resolving uncertainty early, all else equal, would have increased generator profits but decreased exit for generators close to that margin. We find that resolving uncertainty early would have decreased exit and increased pollution overall, because many coal generators were close to the exit margin during the MATS enforcement window. This result holds more generally in the case of negative externalities, where policies may cause firms to take irreversible actions such as exit that reduce the externality. Thus, with negative externalities and when many firms are close to the exit margin, delay in uncertainty resolution *helps* policy goals. Delay in uncertainty resolution will instead *hurt* policy goals in settings where the irreversible outcome is unlikely or where a positive externality exists.

More generally, recent U.S. Supreme Court cases—including the 2022 *West Virginia et al. v. Environmental Protection Agency* case—have called into question the power of agencies to interpret legislation in formulating rules, long established under the Chevron doctrine (?). By affecting the share of policies for which enforcement depends upon court rulings, these

decisions affect policy uncertainty and have increased interest in its impact. Future research could investigate how actual policy choices surrounding the relative power of legislatures, administrative agencies, and the courts affect firm profits and policy goals through their impact on both the level of policy uncertainty and the timing of uncertainty resolution.

Data Availability

We provide data and code for replicating the results in this article at ? in the Harvard Dataverse, XXX.

On-Line Appendix

A1 Evidence on U.S. State Air Toxics Standards

In the years before the EPA released the final MATS rule, a number of U.S. states drafted and passed their own rules and legislation in order to tackle mercury emissions. This section first documents how we determine U.S. state announcement and enforcement dates. Second, we show supporting evidence for our assumption that U.S. state standards were certain once announced. Finally, we explain why we believe that compliance with U.S. state-level standards may have been less onerous than compliance with MATS.²⁹

A1.1 Implementation of U.S. State Standards

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

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

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

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

²⁹We thank Jacob Felton, Amy Mann, Rory Davis, Alison Ray, Joanne Morin, Bruce Monson, and Anne Jackson from state agencies in Connecticut, Delaware, Illinois, Maryland, Massachusetts, and Minnesota for assistance in understanding U.S. state regulations.

(Zibel, 2006). The enforcement date was January 1, 2010 (Maryland Department of the Environment, n.d.).

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

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

New Hampshire passed the Clean Power Mercury Reductions bill in 2005. The law specified mercury emissions reductions by 2013 (Platts, 2005).

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

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

In some cases, we assume that generators were subject to MATS instead of U.S. state standards, even though we found evidence of U.S. state policies. Florida developed a cap-and-trade policy based on CAMR rather than a technology standard (Congressional Research Service, 2007). Michigan's regulation was a backstop to MATS and would only come into force if MATS was invalidated (Michigan Department of Environment, Great Lakes, and Energy, n.d.). New York proposed a relatively weak rule in 2006 with a 50% decline in mercury emissions by 2010 and a 90% decline by 2015 (New York State Department of Environmental Conservation, 2006). Our data show that many plants in New York were not compliant until 2016 under our criterion for compliance with U.S. state regulations, which is based on SO₂ emissions rates. North Carolina's enforcement date was in 2018, substantially after MATS (Congressional Research Service, 2007). As discussed in Section ??, Pennsylvania's standard was overturned by the state's Supreme Court (Mustian and Demase, 2009).

Finally, New Jersey adopted mercury regulations in 2003 with enforcement scheduled for

2007 (Gurney, 2003). However, two New Jersey generators had high SO₂ emissions rates after 2007 that were inconsistent with them being compliant. Given the ambiguity surrounding regulations for New Jersey, we dropped New Jersey generators from our sample.

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

We assume that generators believed that U.S. state standards, once announced, would be enforced with certainty. In three U.S. states (Connecticut, Maryland, and Minnesota), the standard was passed as legislation. This reduces the potential for legal challenges as states have authority to pass legislation that regulates pollutants for sources located within their borders. In other U.S. states, the standards were implemented via rule-making. In most cases, standards were generally developed in conjunction with industry groups. For instance, in Connecticut, the Senate chair of the environment committee commented that “[i]ndustry took a look at that battle and realized it was in their interest to collaborate and compromise” (Halloran, 2003). In Illinois, regulators included alternative emission standards at the behest of local industry, which limited emissions from power plants as a whole rather than from individual generating units (Illinois Pollution Control Board, 2006).

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

Despite the fact that U.S. state standards were essentially certain once passed, there was uncertainty surrounding their initial announcements. Maryland’s Healthy Air Act provides perhaps the most dramatic example of this uncertainty. Republican Governor Robert Erlich and the Democrat-controlled legislature drafted competing bills to address air pollution (Pelton, 2006). The eventual bill that the legislature passed was closer to the Democrats’ version. Erlich’s staff allegedly locked his office door at 4:30 on a Friday afternoon at the end of the legislative session to avoid him being presented with a set of bills that included the

Healthy Air Act. Legislative aides slid the bill under the governor's door. Maryland's Attorney General determined that such action constituted presentment (Marimow and Mosk, 2006), and Erlich ultimately signed the bill (Zibel, 2006). Further illustrating that there was uncertainty surrounding the initial announcements, some U.S. states such as Georgia proposed mercury standards that were never put in place (Cash, 2006). We do not model this source of uncertainty which, as we noted in Section ??, may explain why we find that generators subject to federal standards perceived the probability of MATS enforcement to be 1 in 2012.

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

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

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

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

reduction (Minnesota State Legislature, 2006). New Hampshire specified an 80% reduction in mercury emissions (Platts, 2005). Wisconsin specified the minimum of a limit of 0.0080 lbs/GWh and a 90% reduction (Legislative Reference Bureau Wisconsin, 2006).

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

A2 Data Details

A2.1 Construction of Key Variables

Although compliance is a central variable in our analysis, adoption of technologies that lead to compliance is only imperfectly reported in our data. Specifically, while the EPA reports information on pollution abatement technology by generator, the included descriptions are not detailed enough to evaluate whether the technology complies with MATS. For instance, we observe that some generators use “flue-gas desulfurization” but do not observe the vintage of that technology or the efficacy of the particular equipment installed. Nonetheless, we investigated using this variable to define generator compliance but found it to be unreliable. Specifically, some generators subject to MATS that did not report being compliant according to this variable continued to operate post-enforcement and also had large declines in SO₂ emissions, suggesting they had made other, unreported, technology investments.

Because the most cost-effective technologies that abate both mercury and other air toxics also reduce SO₂, one important way that the EPA determines MATS compliance is via SO₂ emissions rates. In particular, the MATS final rule (77 FR 9304) specifies that generators can comply with MATS by having SO₂ emissions rates below 0.2 lbs/MMBtu. For most generators, we observed a large decline in SO₂ emissions rates in a particular year before air toxics enforcement. For generators subject to MATS, these declines frequently reduced

annual average emissions rates below 0.2 lbs/MMBtu, although in a number of cases the post-decline emissions rates were between 0.2 and 0.4, with some variation across years. In contrast, pre-decline rates were typically well above 1. For this reason, we defined a generator as having adopted MATS abatement technology in the first year when (i) its 3-year forward moving average SO₂ emissions rate falls below 0.4, or (ii) its annual emissions rate falls by 40% or more.³⁰

Turning to U.S. states with their own standard, many 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 and these U.S. states ended up measuring compliance with a combination of technology reporting, periodic stack tests, and other emissions reporting.³¹ Given this complexity, we used a similar method to determine compliance with U.S. state standards as for MATS, with one difference: we used a cutoff of 0.7 here rather than the 0.4 cutoff we used for MATS. This is because we found that the post-decline emissions rates were typically below 0.7 but often above 0.4 lbs/MMBtu. This is also consistent with the evidence in Section ?? that state air toxics standards may be less stringent than MATS.

Exit decisions are also central to our model because generators may respond to air toxics standards by exiting the market. For our generator-year data set, we defined a generator to have exited after the last year in which we observe it generating with coal in the CEMS data.³² For our generator-hour data set, we defined 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 used the end of its CEMS appearance as the exit hour.

Beyond adoption and exit, we also need to define generator fixed characteristics, specifically, minimum and maximum generation levels conditional on generating, capacity, and

³⁰In 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.

³¹MATS—which was implemented after state air toxics standards—did not attempt to measure air toxics compliance through a mercury CEMS.

³²Thus, conversions from coal to natural gas—as analyzed by ? in response to MATS—will appear as exits in our data.

heat rate. Similarly to ?, we defined a generator’s maximum generation level as the 95th percentile of its observed hourly generation conditional on operating in the CEMS data. We also used this as the generator’s capacity.³³ We defined a generator’s minimum generation level as its modal generation level between the 5th and 60th percentile of capacity. We then binned hours with positive generation into minimum and maximum generation levels based on whichever level is nearer.

Finally, we calculated 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 defined each generator’s time-invariant heat rate as the mean hourly heat rate across hours in the maximum generation bin.

A2.2 Generator Market Structure

For tractability, we assume that strategic decisions are made at the generator level, rather than at the IPP (firm) or power plant levels. As discussed in Section ??, we believe that this assumption is reasonable because most IPPs have few plants and most plants have few generators. Specifically, operators on average across years have 1.61 plants (1.49 within a U.S. state), with 73.7% of operators controlling a single plant. Plants have on average 2.34 generators, with 67.2% of plants having two or fewer generators.

We also find that most generators operate predominantly in one U.S. state: the average number of U.S. states per operator in a year is 1.08. This leads us to believe that modeling joint ownership across markets is relatively unimportant.

Finally, even the largest IPPs do not have a national presence in coal generation. The largest operator—as measured by number of generators—controlled a maximum of 24 generators (at the beginning of our sample), of which 21 were in Ohio. By the end of the sample, this operator controlled only 10 generators, 7 in Ohio and 3 in Pennsylvania. Even in Ohio, this operator had only 27.4% of the IPP coal capacity, across years in our sample.

³³We choose the 95th percentile because while generators can generate above listed capacity, this extra generation is extremely costly in the long-run.

A3 Different Approximate Equilibrium Concepts

This paper develops and uses a concept called Approximate Belief Oligopoly Equilibrium (ABOE), which builds on earlier equilibrium concepts that break the curse of dimensionality with different assumptions. This appendix provides a taxonomy of the major differences between some of these different approaches. We identify four characteristics of these equilibrium approaches:

1. **Consideration of only the steady state:** One way to break the curse of dimensionality is to only consider markets that are in a long-run steady state. This then means that a player needs to only keep track of its own characteristics and not market characteristics. Oblivious Equilibrium (OE, Weintraub et al., 2008) uses this approach, by letting players react only to the steady state distribution of rivals' states.
2. **Aggregate market states that summarize relevant information:** This alternative approach allows for the analysis of markets not in steady state. Krusell and Smith (1998) and Lee (2013) were among the initial adopters of this assumption in empirical settings. It is used by Moment-based Markov Equilibrium (MME, Ifrach and Weintraub, 2017) and ABOE. Experience-Based Equilibrium (EBE, Fershtman and Pakes, 2012) extends the state space beyond just payoff-relevant state variables and hence does not rely on aggregation.
3. **Moments that approximate beliefs about transitions:** Even with the use of aggregate moments to simplify the state, state transitions may be complicated enough that transition moments are helpful in reducing computational costs. Krusell and Smith (1998), MME, and ABOE use moments in this way, with our paper specifying the moments as the coefficients of AR(1) processes. Lee (2013) has a smaller number of aggregate market states and therefore can calculate state transitions without approximations. EBE also does not rely on approximations (as above).
4. **Players that do not recognize their impact on the aggregate state:** In situations with many players, authors have simplified the state space by assuming that

players do not recognize that their own actions will affect the market state. This is the approach taken by Krusell and Smith (1998) and Lee (2013). MME relaxes this assumption by specifying a set of oligopolists who recognize their exact impact on aggregate states but are the only players to recognize that they can affect next period's aggregate states. The set of oligopolists can vary endogenously over time. An ABOE instead treats all players symmetrically. In an ABOE, all players recognize that their actions will affect next period's aggregate state, but they also have only approximately correct beliefs regarding this impact.

A4 Equilibrium Computation and Estimation

We estimate our model and conduct counterfactuals by solving for equilibria across candidate parameter vectors, U.S. states, and belief years. We recover the equilibrium as the fixed point of an algorithm that iterates between solving individual optimization decisions and estimating market evolution regressions on data that are generated by simulations that use these decisions. We do not offer a formal proof of the existence of equilibrium, although we believe that this would be an application of Brouwer's Fixed Point Theorem because each generator's reaction function is continuous in the reaction functions of other generators. Further, while we do not provide a proof of convergence of our algorithm, we obtained convergence for all parameters that we tried.

Another concern is whether our equilibrium is unique. Multiple equilibria often exist in exit models. For instance, at some states, it may be possible to sustain one generator but not two generators in equilibrium, so there may be equilibria where generator 1 exits and others where generator 2 exits. The private information in our model limits the extent of multiple equilibria but does not eliminate them as a concern. Intuitively, the more important the private information, the less likely the model is to have multiple equilibria, which we believe is the case in our setting. Further, we start our equilibrium calculations at the data values, which will select an equilibrium that is "close" to the data if there are multiple ones.

This appendix next details our model's state space and explains 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.

A4.1 State Space

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

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

Units subject to U.S. state standards are certain about the standard being implemented once it is announced. Thus, belief year τ is not relevant for these U.S. states. For these units, there are $1 + 2 \times (\text{Enforcement Year} - \text{Announced Year})$ discrete states.

A4.2 Details of Generators' Dynamic Optimization

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

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

A4.3 Market Evolution Regressions

As discussed in Section ??, 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, coal share is the sum of the non-IPP and IPP coal capacity divided by load. We specify an exogenous AR(1) regression for non-IPP coal share, where it depends on its lag, the lagged fuel price ratio, the interaction between these two variables, and a constant. We estimate one non-IPP coal share regression across our entire sample. We cannot estimate

separate non-IPP regressions by state because our assumption that this process is exogenous implies that we cannot simulate additional data beyond what are observed, and we do not have sufficient observations to credibly identify transitions by U.S. state.

We believe that an exogenous evolution here is a reasonable approximation for two reasons. First, while there is overlap in some U.S. states, restructured U.S. states tend to have mostly IPP coal generation and regulated U.S. states tend to have mostly non-IPP coal generation. Second, non-IPP coal generators are largely subject to rate-of-return regulation which causes their incentives to be different from the incentives of IPP coal generators (?). We examined the most prominent exception to this pattern in our data, Louisiana, which was evenly split between IPP and non-IPP coal generators in 2006. By the end of our sample, no non-IPP coal generators in Louisiana had exited, while there was substantial exit of IPP coal generators. This reinforces the idea that rate-of-return regulation provides utilities very different incentives with respect to installing abatement technology or exiting.

The IPP coal share evolves endogenously in the model. For each U.S. state and year, we generate the following year's IPP coal share by simulating generators' decisions given their optimizing behavior as calculated from the Bellman equations. To simulate the next period's coal share—which is the dependent variable in the first AR(1) process—we add a draw of IPP coal share (derived from the model) to a draw of non-IPP coal share (derived from our non-IPP regression). Generators only value knowing future coal share in the case where they do not exit. Thus, we approximate this endogeneity of future market structure by randomly selecting one generator, and forcing it to remain active, regardless of its simulated strategy as predicted by the model. It would be possible to estimate a separate AR(1) regression for each generator, where we force each generator to remain active in its regression, but we chose this random selection approach instead for tractability.

When outside the enforcement window, coal share 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.

Second, 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. Similar to the non-IPP coal share, because we are identifying this regression from data—rather than model

simulations—we estimate one regression across our entire sample.

Third, the share of IPP coal capacity that has adopted abatement technology evolves endogenously to the model, similarly to the IPP coal share. 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 with IPP coal share, 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 share, 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 share 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 share regression, we use as the dependent variable the expectation of the IPP coal share in the next year—given optimizing generator policies—plus a simulated draw from the (exogenous) non-IPP coal share process. Similarly, for the adoption regression, the dependent variable is the expectation of the adoption share in the next year, again given optimizing policies. We use expectations rather than simulation draws in order to reduce the variance of our dependent variables.

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

We let the probability of enforcement at the end of the final year before enforcement be

equal to P_τ . We therefore simulate the realization of enforcement as a correlated shock to all units in the state.

A4.4 Bootstrapped Standard Errors

In order to calculate standard errors for our parameter estimates, we conduct a parametric bootstrap. This involves simulating 100 data sets created from the equilibrium evaluated at the parameter estimates and re-estimating our model on these data sets.

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

In our simulations, we assume that the exogenous processes—fuel price ratio and non-IPP coal share—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.³⁴ In each subsequent year, we update these variables by drawing from their evolution distributions.

We also begin the endogenous market state variables—IPP coal share 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.

A4.5 Counterfactual Calculations

We calculate counterfactual outcomes using a similar approach to how we simulate data for our bootstrap. This involves solving the ABOE of the model and then simulating data.

³⁴Florida and Michigan do not report any IPP coal generators until 2008. We therefore start our bootstrapped data sets in 2008 for these states.

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 ??.


A5 Extra Tables and Figures

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

State	Announced	Enforced
Connecticut	2003	2009
Massachusetts	2004	2009
New Hampshire	2005	2013
Illinois	2006	2010
Delaware	2006	2010
Maryland	2006	2011
Minnesota	2006	2011
Wisconsin	2008	2016

Note: Announcement and enforcement years are based on sources discussed in On-line Appendix ??.

Figure A1: Distribution of Calculated Profits



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

Table A2: Profit Surface Robustness

	Annual Profit (millions of \$)			
	(1)	(2)	(3)	(4)
U.S. State Coal Share	-21.984*** (6.266)	-18.202*** (6.418)	-21.650*** (6.280)	-16.980*** (2.044)
U.S. State Coal Share Squared	10.777*** (1.953)	9.975*** (1.983)	10.778*** (1.954)	6.585*** (0.573)
Gas-to-Coal Fuel Price Ratio	5.181*** (0.873)	5.356*** (0.875)	5.283*** (0.882)	8.679*** (1.069)
Gas-to-Coal Fuel Price Ratio Squared	-0.625*** (0.097)	-0.628*** (0.097)	-0.633*** (0.097)	-1.051*** (0.124)
Coal Share \times Fuel Price Ratio	-1.141*** (0.327)	-1.237*** (0.328)	-1.142*** (0.327)	-2.030*** (0.400)
Heat Rate (MMBtu/MW)	-0.313*** (0.113)	-0.262** (0.114)	-0.320*** (0.114)	-0.779*** (0.130)
Capacity (MW)	-0.035*** (0.005)	-0.040*** (0.006)	-0.035*** (0.005)	-0.107*** (0.006)
Capacity Squared (MW ²)	-2.3e-05*** (4.0e-06)	-2.0e-05*** (4.1e-06)	-2.3e-05*** (4.0e-06)	1.3e-05*** (4.7e-06)
Capacity \times U.S. State Coal Share	-0.007*** (0.002)	-0.007*** (0.002)	-0.007*** (0.002)	0.005** (0.002)
Capacity \times Fuel Price Ratio	0.048*** (0.001)	0.048*** (0.001)	0.049*** (0.001)	0.051*** (0.001)
IPP Coal Operator HHI		-1.861 (3.257)		
Operator Capacity Share		7.416*** (2.291)		
Compliant			0.782 (0.963)	
U.S. State FE	Y	Y	Y	N
Observations	2899	2899	2899	2899
R^2	0.7791	0.7802	0.7792	0.6071

Note: Regression of calculated profits from observed data on dynamic model states. We exclude observations for generators that did not produce during the year. Standard errors are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Appendix References

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

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

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

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

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

Fershtman, Chaim and Ariel Pakes, “Dynamic Games with Asymmetric Information: A Framework for Empirical Work,” *The Quarterly Journal of Economics*, 2012, 127 (4), 1611–1661.

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

Gurney, Kaitlin, “State Toughens Mercury Limits; Power and Industrial Plants Must Cut Levels Sharply Over 3 Years. U.S. Actions Spurred the Move,” *The Philadelphia Inquirer*, 11 December 2003.

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

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

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

Ifrach, Bar and Gabriel Y. Weintraub, “A Framework for Dynamic Oligopoly in Concentrated Industries,” *The Review of Economic Studies*, 2017, 84 (3), 1106–1150.

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

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

Krusell, Per and Anthony A Smith Jr, “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 1998, 106 (5), 867–896.

Lee, Robin S., “Vertical Integration and Exclusivity in Platform and Two-Sided Markets,” *American Economic Review*, 2013, 103 (7), 2960–3000.

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

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

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

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

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

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

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

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

New York State Department of Environmental Conservation, “Hearing Report on the Proposed New York Section 111(d) State Plan for Implementation of Coal-Fired Electric Steam Generating Unit Mercury Emission Guidelines,” 14 November 2006.

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

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

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

Platts, “New Hampshire Accord Would Reduce Mercury, Sulfur Emissions,” *Coal Outlook*, 14 November 2005.

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

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

Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy, “Markov Perfect Industry Dynamics with Many Firms,” *Econometrica*, 2008, 76 (6), 1375–1411.

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

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