Escalation of Scrutiny: The Gains from Dynamic Enforcement of Environmental Regulations *

Wesley Blundell† Gautam Gowrisankaran‡ Ashley Langer§

May 4, 2018

Abstract

In the United States, federal and state governments spent over $20 billion in 2014 on the enforcement of environmental regulations and laws, including the Clean Air Act and Amendments and the Clean Water Act. The Environmental Protection Agency uses a dynamic approach to enforcement, including by designating repeat offenders as *high priority violators* and targeting them with elevated scrutiny and penalties, which may allow it to mitigate the costs of enforcement and improve compliance. We estimate the benefits of dynamic monitoring and enforcement to ensure compliance with air pollution standards. We develop and estimate a single-agent, dynamic, discrete-time model of a plant faced with a regulator. The central decision that a plant faces is whether and when to invest in pollution abatement technologies. The regulator enforces environmental laws with three principal actions: inspecting plants, recording violations, and issuing fines. Plants face a disutility from investment and from regulatory enforcement. We use a fixed grid approach to estimate specifications where the disutility parameters are heterogeneous across plants. We estimate that investment and the designation as a high priority violator are both very costly to plants. Reductions in the non-linearity of fines would have large adverse impacts on the fraction of plants that are high priority violators and would increase both pollution and the costs of enforcing environmental regulations.

Keywords: environmental regulation, dynamic estimation, Clean Air Act

---

*We thank Christian Belzil, Wayne Gray, Alex Hollingsworth, Anatolii Kokoza, Derek Lemoine, Ivan Rudik, Jay Shimshack, Kyle Wilson, Tiemen Woutersen, Chenyu Yang, and seminar participants at numerous institutions for helpful comments.

†California State University, East Bay

‡University of Arizona, HEC Montreal, and NBER

§University of Arizona

1 Introduction

In the United States, federal and state governments spent nearly $21 billion in 2014 on the enforcement of environmental regulations and laws, including the Clean Air Act and Amendments and the Clean Water Act. These regulations impact nearly every industrial facility in the U.S., leading plants to invest substantial amounts of time and money in environmental compliance. The massive expenditures by regulators and plants in enforcing and complying with environmental legislation makes it critical to understand the efficiency of regulatory monitoring and enforcement mechanisms for pollution control.

Moreover, there are numerous complications to designing effective mechanisms for monitoring and compliance with environmental regulations: perfect monitoring of emissions is impossible; plants' investments in pollution abatement are costly, discrete, and take time to implement; penalties for non-compliance are limited by bankruptcy laws and political pressure; and plants are heterogeneous in their costs of pollution mitigation and regulatory compliance. These complications have led environmental regulators to create systems whereby repeat offenders are punished much more severely than are one-time violators. The U.S. Environmental Protection Agency (EPA) designates persistent offenders as *high priority violators*, and targets them with elevated scrutiny and penalties. Figure 1 shows the EPA's inspection rate, level of fines, and violation rate, displayed separately based on whether the plant is in compliance, a regular (not-high-priority) violator, or a high priority violator. In each case, the level of scrutiny increases dramatically with the regulatory state.[1]

This dynamic approach to enforcement potentially allows regulators to mitigate some of the above complications. For instance, with imperfect monitoring, repeated findings of violation add certainty to the noisy signal from any one inspection. In addition, the threat of high fines may result in high compliance rates without the political cost of actually imposing high fines frequently. Finally, dynamic enforcement may lead to improved compliance among plants with high mitigation costs without excessively penalizing low-cost plants.

[1] The higher fines for plants in high priority violator status in Figure 1 could be due to dynamic enforcement or to those plants violating environmental norms more frequently. Our analysis allows for both of these explanations.

Figure 1: EPA Clean Air Act and Amendment enforcement activities by environmental regulatory status, 2007-2013.

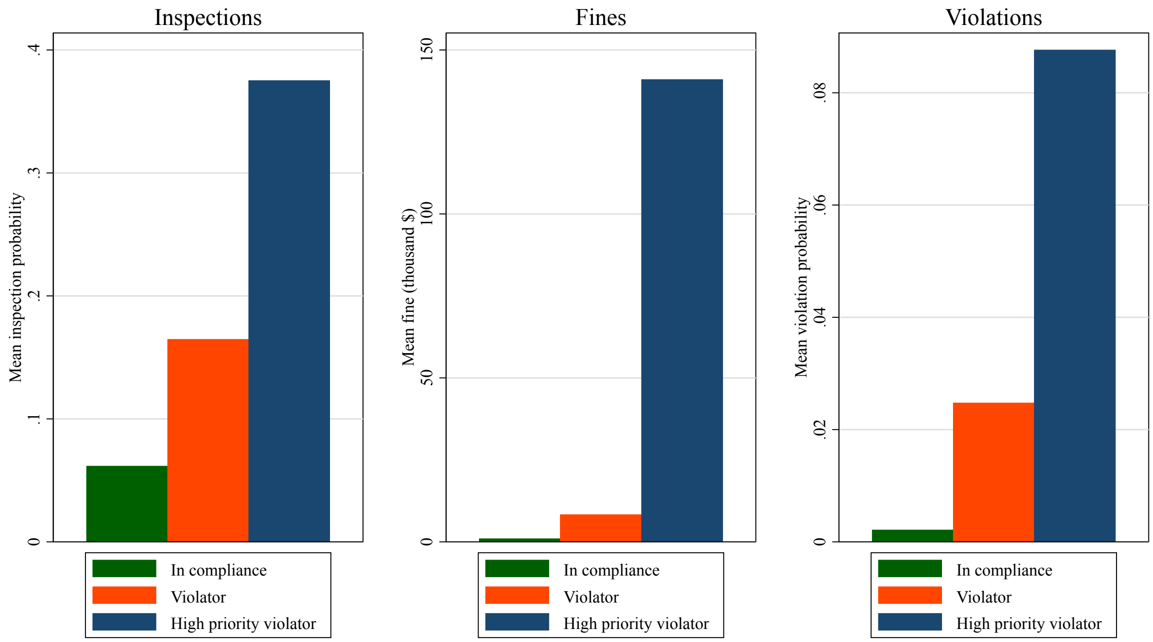

This paper addresses two broad research questions on environmental monitoring and enforcement of air pollution laws. First, we document empirically the use of dynamic enforcement of environmental regulations by the EPA for industrial facilities covered by the U.S. Clean Air Act and Amendments (CAAA). Second, we evaluate the value of the current dynamic enforcement approach and alternative enforcement regimes in ensuring plants' compliance with the CAAA.

We believe that our modeling and estimation framework may be generally useful in evaluating environmental regulations. Dynamic enforcement is widely documented for many environmental regulations in the U.S., including the CAAA (Evans, 2016; Blundell, 2017) and the Clean Water Act (Earnhart, 2004; Shimshack and Ward, 2005). It is also widely used for environmental regulation in other countries, for example, for petroleum storage in Canada (Eckert, 2004); air pollution in Norway (Telle, 2013); soil, water, and air pollution in Belgium (Blondiau et al., 2015); and waste management in Japan (Shinkuma and Managi, 2012). Dynamic enforcement is also widely used (and studied) beyond environmental regulation, e.g., in tax compliance and auditing (Landsberger and Meilijson, 1982) and worker health and safety regulation (Ko et al., 2010). Thus, a framework to understand the value added from dynamic enforcement of regulations is useful in a variety of contexts.

The above papers have not considered dynamic enforcement of regulations with a model that formally accounts for the dynamic, optimizing behavior of plants. Yet, evaluating the value of dynamic enforcement without a dynamic model is inherently problematic. As noted by a recent survey article in this area, a static model of enforcement would tend to predict compliance rates that are well below those observed in the real world, because it would not account for the future penalties from current non-compliance (Shimshack, 2014). Shimshack further notes that to account for these effects, one would need to consider a dynamic model, such as the theoretical model proposed by Harrington (1988). Because of the difficulty in estimating such models, Shimshack views environmental monitoring and enforcement as "both understudied and controversial," (p. 3).

In order to understand the impact of different regulatory structures on environmental compliance, we develop and estimate a single-agent, dynamic, discrete-time model of a plant faced with a regulator, in the spirit of Harrington (1988). The central decision that a plant faces is whether and when to invest in pollution abatement technologies. At any point in time, the plant has a regulatory state, which lists the plant's EPA region and industrial sector, whether it is a regular or high priority violator, and a summary of its history of prior violations and investments.

Each period, the regulator first decides whether or not to inspect the plant. It has a policy rule that indicates the probability of an inspection in the period based on the plant's regulatory state. The inspection reveals a signal about the plant's environmental compliance. The signal indicates the presence of a reportable violation and whether the plant is in compliance, a regular violator, or a high priority violator, given current environmental standards and the plant's previous state. The regulator also has a second policy rule, which indicates how much to fine the plant, if at all, based on the signal and the plant's regulatory state. The plant obtains disutility from regulatory enforcement actions and designation as a high priority violator. It knows the regulator's stochastic and dynamic policy rule and has the same information as the regulator regarding the distribution of the environmental compliance signal.

Following the regulator's action, if the plant is a regular violator or high priority violator, it makes a binary choice of whether to invest in environmental mitigation. The plant faces a cost of investment, including an idiosyncratic component. Though costly, investment increases the probability that the plant will return to compliance (thereby reducing its regulatory burden) over the subsequent periods. Effectively, each period, the plant first bears the costs of the regulatory actions of inspections, violations, and fines, and designation as a high priority violator. Then, if it is not in compliance, it invests if its cost of investment is less than its expected future benefit in the form of future lower environmental scrutiny.

Our model and estimation framework is broadly similar to the dynamic model of investment developed by Rust (1987). As in Rust (1987), the renewal action in our model is investment. To fit our institutional features and data, we model regulatory actions and outcomes as more stochastic than in Rust (1987), with investment only probabilistically returning plants to compliance and transitions driven in part by the regulator's actions. We estimate models of regulatory actions and transitions conditional on the regulatory state and investment decision—with simple probit and tobit regressions that account for the limited dependent nature of these actions—and then use the fitted values from these regressions as inputs to our dynamic optimizing model of plant investment decisions.

Our study make uses of extensive data that include information on virtually all industrial facilities covered by the CAAA. Our data report inspections, fines, violations, compliance status, and investment decisions for a seven-year long panel with over 5 million plant-quarter observations. These extensive data allow us to develop a framework that appropriately accounts for plants' dynamic incentives to invest in pollution abatement and for the heterogeneity in plants' costs of investment and regulation. Investments in our data mostly correspond to codes for the resolution of substantial environmental problems reported by the EPA.

We estimate two main econometric specifications. The specifications differ in their assumptions regarding the heterogeneity in the disutility parameters across plants and in their estimation approach. First, we estimate a specification where the disutility parameters are the same across plants. For this specification, we estimate the disutility parameters using a maximum likelihood nested fixed point estimator where the dependent variable is investment.

Second, we estimate a specification where the disutility parameters are heterogeneous across plants. Our heterogeneity specifications allow for each plant to have a fixed vector of disutility parameters that are drawn from a finite grid of potential disutility parameters. The structural parameters that we estimate in this specification are the population weights of observing each of the potential disutility parameters. By choosing a finite grid, this specification essentially allows us to estimate a non-parametric distribution of the cost of investments and regulatory actions and outcomes across plants. The generalized method of moments (GMM) estimator that we use here is computationally very tractable.

Identification of our model is somewhat different from much of the environmental regulation literature. As in that literature more generally (surveyed in Gray and Shimshack (2011) and elsewhere), our fundamental research objective is to understand the extent to which government enforcement activities lead to reductions in pollution. Our approach uses structural estimation to estimate the cost to the plant of investments in pollution mitigation technologies and the benefits from these investments in terms of reduced regulatory scrutiny. Using these estimates, we then conduct counterfactual simulations to evaluate how investments and pollution would change under different regulatory regimes. Heuristically, to understand the costs and benefits of investment, one could estimate a discrete choice specification where plant investment is regressed on a constant (interpreted as the cost of investment) and the decreased regulatory scrutiny that would accrue from such investment. Importantly, one could not use the actual future scrutiny as the regressor, as it would be endogenous, but would need to simulate how the future scrutiny changed from investment. Our first specification essentially performs this regression, also allowing for plants to have the option value of investment in future periods. It will hence be identified by the extent to which investment lowers regulatory scrutiny among plants with observably similar characteristics and the extent to which plants respond to this decrease through investment. Variation in the regulatory regimes across EPA regions and industrial sectors adds substantial identifying variation here. Our second specification recovers the heterogeneity in utility parameters by using the equilibrium distribution of plant states and correlations in investment across time for plants. The panel nature of our data adds to identification here.

Summary of results. The estimates from our first specification (where constrain the disutility to be the same across plants), show that plants have a substantial disutility from environmental enforcement. Inspections, violations, fines, and classification as a high priority violator all carry substantial and statistically significant costs. We estimate that investment is very costly to plants, with a new investment being equivalent to a $25 million fine (holding constant other regulatory actions). This finding contrasts with the fact that average fines are low and compliance rates are high. It occurs because the dynamic implications of investment are important. We also find that plants bear a substantial cost of being a high priority violator—in addition to the fines and regulatory actions that they receive while in this status—with the mean cost to a plant from being a high priority violator equivalent to a fine of $350,000 per quarter. This cost may occur because being a high priority violator may affect the plant's reputation and relationship with the surrounding community.

Our estimates from our second specification, which allows for random coefficients with a GMM estimator shows that there is substantial heterogeneity across plants in their disutility from environmental enforcement, with the mean of the coefficients being similar to the first specification. In particular, we find that there is a substantial mass of plants which find investment and fines very costly; another mass which finds both less costly relative to the first group; and a third mass which finds fines relatively more costly than in the other two groups.

Using our estimated parameters, we construct counterfactual estimates of how plants' investment decisions and regulatory status would change if the regulatory structure were different. In particular, we focus on (1) the non-linearity of penalties across regulatory states, (2) the regulatory compliance costs, and (3) the stigma from being a high priority violator. We examine how changes in these variables would affect plant utility and investment decisions, equilibrium plant compliance rates, regulatory actions, and overall pollution levels. We report counterfactuals for both specifications. We find that if we were to eliminate the non-linearity of fines while keeping total equilibrium fines the same as in the baseline, we would have far more plants in HPV status. This result is particularly dramatic for the random coefficients estimator. We also find that if we eliminated the cost of inspections, violations, and HPV status, the HPV rate would also increase substantially. Altogether, our counterfactuals show that dynamic regulation has a big impact in increasing environmental compliance. This effect is magnified when we account for the heterogeneity in the utility functions across plants with our GMM random coefficients estimator.

Relation to literature. This paper relates most closely to three distinct literatures. First, we build on a substantial theoretical literature on the role of dynamic enforcement of environmental regulations, by providing empirical findings. As noted above, this literature started with Harrington (1988) which builds on the classic work of Becker (1968) on the economics of crime. Papers in this literature evaluate the conditions for which varying the rate of inspection between groups can decrease social costs (Friesen, 2003; Harford, 1991) and show that dynamic enforcement may fail to decrease social costs under regulatory dealing (Heyes and Rickman, 1999) and asymmetric information (Raymond, 1999).

Second, we build on an emerging literature that structurally estimates firm behavior in the presence of environmental regulatory policies (Timmins, 2002; Ryan, 2012; Duflo et al., 2016; Fowlie et al., 2016; Kang and Silveira, 2018). Kang and Silveira (2018) estimate a model of adverse selection and heterogeneous punishment for the water pollution in California. Unlike our work, this paper focuses on information but not on dynamics. Duflo et al. (2016) estimate a dynamic model of environmental regulatory enforcement for plants in India. Our model differs from Duflo et al. in a number of ways that relate to the stochasticity that is present in the CAAA context, where inspections and other regulatory actions do not occur with certainty and where investment does not remediate violations with certainty.

Third, we extend the estimating framework for dynamic discrete choice models with random coefficients (Arcidiacono and Miller, 2011; Fox et al., 2011; Gowrisankaran and Rysman, 2012; Connault, 2016; Fox et al., 2016; Nevo et al., 2016). In this dimension, our paper is most similar to Fox et al. (2011, 2016); Nevo et al. (2016) in that it uses the same fixed grid GMM approach and the same basic computational techniques. We further develop this approach to allow for dynamic models with infinite horizons and to allow for moments that will further capture the presence of heterogeneity.

The remainder of the paper is organized as follows. Section 2 documents the regulatory context and its implications regarding plant incentives for compliance. Section 3 details our data and provides reduced-form evidence that motivates the choices in our structural model. Section 4 describes our model of environmental regulation and plant investment in pollution abatement. Section 5 details our estimation strategy and identification. Section 6 presents our results and counterfactuals. Section 7 concludes.

2 Regulatory Framework

2.1 Background

We examine the dynamic enforcement of environmental regulations, focusing on the enforcement of the Clean Air Act and Amendments (CAAA). Before turning to a description of our data and empirical approach, it is important to understand how the federal government enforces the CAAA and the implications of this enforcement regime for plants' incentives to invest in pollution abatement.

Congress passed the Clean Air Act in 1963 in an effort to improve air quality nationally. The scope of regulation was expanded with amendments in 1965 and 1970, at which point President Nixon created the Environmental Protection Agency (EPA) to enforce the Clean Air Act and other environmental legislation. The Clean Air Act was again amended in 1977 and 1990. Our data are from after 1990 and hence we consider the Clean Air Act and Amendments.

Many of the decisions of how exactly to structure this enforcement was left to the EPA, which created a system of inspections, violations, fines, and other requirements (e.g. self-reporting paperwork). This enforcement structure aims to ensure that plants are complying with the CAAA emissions and technology standards and to move plants that are out of compliance back into compliance via plant investments in improved processes or technology.

While the structure of CAAA enforcement is dictated by the EPA, much of the actual enforcement activity is carried out by regional- and state-level environmental protection agencies.[2]

In particular, the EPA divides the country into 10 geographic EPA regions. Significant portions of the EPA's operations are conducted through these regional offices. For instance, regional EPA offices conduct inspections and/or issue sanctions when a state's enforcement is below required levels. They also assist states with major cases. Another function of the regional offices is to identify and synthesize the concerns of their states in order to form a "regional view" that will be factored in the EPA's decision making process. Overall, the regional offices represent the smallest level that captures both the interpretation of federal policy and geographic preferences for enforcement. For these reasons, our empirical analyses use the EPA region for identifying variation.

Under this enforcement system, plants that are in compliance can expect to be inspected regularly and must self-report certain potential violations to regulators. Thus, plants that are in compliance could enter "violator" status either because the regulator physically comes to the plant to conduct an inspection and discovers a violation or because the plant self-reports that, for instance, a piece of machinery broke and the plant was briefly out of compliance until it was repaired. Being a violator subjects the plant to additional inspections, which could possibly uncover additional violations, and potential fines. Plants can accumulate multiple violations within violator status and will only return to compliance once those violations have been resolved. The cost to the plant of being a violator therefore comes not only from the investment cost required to resolve outstanding violations, but also from an increased level of regulatory oversight that may be costly to the plant in terms of monetary costs such as fines and non-monetary costs such as the disruption of facing frequent inspections.

Further, if a plant is substantially or persistently out of compliance, then the plant may be designated a high priority violator (HPV). HPVs are subjected to more frequent inspections (which can lead to uncovering additional violations) and higher fines and other penalties

[2] While these agencies are often called something other than an "EPA" (e.g., the Florida Department of Environmental Protection), we will refer to them as state and/or regional EPAs for consistency. State and regional EPAs are required to maintain a minimum level of enforcement but can exceed this threshold (Shimshack, 2014).

than "regular" violators, a fact which we also demonstrate empirically. Further, HPVs that do not resolve all of their violations in a timely manner were, during the time frame of our analysis, placed on the EPA "Watchlist."

The Watchlist was originally intended to increase oversight by the EPA, but was also reported in the press and could lead to increased attention on the plant from local politicians and civilian environmental protection groups and can be costly to the plant in terms of reputation or increased local regulations (Evans, 2016). Once plants are designated HPVs, they can only exit HPV status by resolving *all* of their outstanding violations, regardless of whether those violations would independently elevate the plant to HPV status. The combination of increased inspections, violations, fines, and general regulatory oversight means that HPV status is—and is intended to be—substantially costly for plants.

2.2 Implications of Current Regulatory Structure

Before turning to our data, it is worth understanding why the dynamic regulatory enforcement favored by the EPA may actually lead to higher compliance in this context than other approaches. In a world where both plants and the regulator can perfectly observe non-compliance, the plant can instantly correct it with a costly investment, and penalties are unbounded, the regulator's optimal approach may be to issue a severe penalty whenever a violation is observed. Assuming that the penalty is high enough, plants will always invest in compliance as soon as a violation occurs and no penalties will ever be issued.

However, enforcement of the CAAA violates this simple model in a number of ways. First, since the EPA is not perfectly able to observe violations of the CAAA and plants may not even be aware of violations until they are revealed by EPA, severely punishing plants when violations are uncovered may not actually improve compliance. Instead, the EPA issues only minor penalties for an initial violation and then only increases penalties and oversights as additional violations accumulate. This approach allows for the possibility that the EPA's inspection process may contain noise, and apparent violations may not actually be present or may not be severe. Additionally, this approach recognizes that plants may not have been aware of the violation before the EPA's inspection, by providing the plant an opportunity to correct the problem before substantial penalties are imposed.

Second, partially because of bankruptcy laws and partially because of political pressure, penalties are limited. In particular, there may be a cost to the EPA of issuing extremely severe penalties since driving plants out of business for small infractions would undermine political support for the CAAA in particular and the EPA in general. Thus, there is an advantage to the EPA of obtaining compliance without issuing numerous large penalties. This is exactly what this dynamic enforcement scheme achieves: plants may make costly investments in pollution abatement even if the current penalty is lower than the investment cost if the investment decreases the regulatory burden that the plant will face in the future (Harrington, 1988).

Finally, the combination of limits on the total magnitude of penalties and heterogeneous investment costs across plants may further add value to dynamic enforcement. If plants have heterogeneous investment costs, then a dynamic enforcement approach will have low-cost plants investing in pollution abatement when they are merely regular violators and the penalties are relatively low. Higher cost plants will potentially wait until they become HPVs—and are faced with higher penalties—before investing in pollution abatement. This non-linear penalty approach will therefore allow the regulator to achieve compliance by high cost plants without over-penalizing low cost plants (Harford, 1991). This ability to induce compliance among high cost plants when total penalties are limited may be particularly important if high-investment-cost plants also emit a disproportionate share of total pollutants (Heyes, 2000).

3 Data and Reduced Form Evidence

3.1 Description of Data

Our study primarily uses two CAAA monitoring and enforcement databases. The first is the Environmental Compliance History Online (ECHO) database, which records information on investment and regulatory compliance. The second is the National Emissions Inventory (NEI) database, which records emissions information. Both databases are at the plant level, maintained by the EPA, and publicly available. We now detail our use of both databases.

The Environmental Compliance History Online (ECHO) database

The ECHO database is divided into a number of component datasets. We principally use five ECHO components: (1) the *Facility Registry Service* dataset, (2) the *Air Facility System Actions* dataset, (3) the *Air Program Historical Compliance* dataset, (4) the *High Priority Violator History* dataset, and (5) fine data scraped from the EPA website's "Enforcement Case Search" page.[3] We discuss each of these components in turn.

First, the *Facility Registry Service* dataset is a master list of plants. For our purposes, it provides address information and the six-digit North American Industry Classification System (NAICS) industrial sector for the plant. Most of our analyses control for EPA region and for industrial sector with the first two-digits of the NAICS code, which we derive from this dataset.

Second, the *Air Facility System Actions* dataset (or *Actions* dataset for short) records the history of regulatory actions taken by state, regional, and federal environmental regulators, from the Q4:2006 through the Q4:2014.[4] We use this dataset to create our base list of inspections, violations, fines, and investments. Each record in this dataset details an action, such as an inspection, a notice of violation, a fine, or the review of an environmental mitigation investment. Plant are identified by a field called the AFS ID. Each record lists a calendar date and provides information on the related *EPA program*,[5] and the penalty amount, when the action is a fine. For each plant, we combine EPA actions across all EPA programs in order to capture completely its regulatory enforcement status. We deflate the penalty amount by the non-energy current price index and record amounts in 2007 dollars. We define a plant investment if either we observe a code indicating the resolution of an environmental issue or

[3] See `https://echo.epa.gov/facilities/enforcement-case-search`.

[4] The EPA transitioned to a new reporting system after 2014.

[5] The CAAA includes many different statutes that address different dimensions of air pollution. The EPA enforces different statutes through different programs.

a compliance certification under the *New Source Review* under Title V occurred.[6] Since this dataset is subject to federal minimum data requirements, we believe it provides a relatively complete description of the regulatory action history for each plant.

Third, the *Air Program Historical Compliance* dataset records the historical compliance status for each plant and EPA program at the AFS ID and quarter level. These data derive from a combination of self-reports by plants and regulator inputs. We follow the literature (Laplante and Rilstone, 1996; Shimshack and Ward, 2005) in treating the self-reported data as accurate.[7] We use this dataset to determine whether a plant is in compliance or a violator in any quarter. We assume that a plant must be in compliance with every CAAA program in order to be considered in compliance in our analysis. This dataset provides a more direct measure of violator status than does the *Actions* dataset since the *Actions* dataset does not always indicate when a regular violation is resolved. Since this dataset is at the plant / quarter level, we aggregate EPA actions to this level and use this as the time period for our analysis.

Fourth, the *High Priority Violator History* dataset records the dates at which a plant receives or resolves a high priority violation. We use this dataset to record the quarter of entry and exit from HPV status. Analogous to the *Air Program Historical Compliance* dataset, this dataset provides the most direct measure of HPV status. Because HPV status is triggered by substantial or persistent violation of the CAAA, we also assume that the plant needs to make an investment in pollution mitigation to leave HPV status.

Finally, we augment our fine data with scraped environmental case data. We retrieved approximately 22,000 EPA enforcement cases from their enforcement web site. These data provide a more complete picture of penalties faced by plants. For instance, they include the cost of compliance evaluations, which are part of the penalty amount but not in the *Actions* dataset; they include information from cases that involve both air and water pollution; and they include cases that should have been entered in the *Actions* dataset but were not. We

[6] According to the ECHO data dictionary, the *New Source Review* is a permitting program that "assures that state-of-the-art control technology is installed at new or existing plants that are undergoing a major modification."

[7] The literature makes this assumption because the expected penalty from purposefully deceiving regulators is far greater than the penalty for an emissions violation.

use as our fine amount the maximum amount across the two datasets. The plant identifier in this dataset is the FRS number and not the AFS ID. The EPA provides a crosswalk from AFS IDs and FRS numbers. While AFS ID and FRS numbers both are specific to a plant, we found some instances where multiple AFS IDs correspond to a single FRS number (and no instances of the opposite). Accordingly, we aggregate our data from the AFS ID to the FRS number. Thus, our unit of analysis is the FRS number / quarter.

Our main analysis data merge together the above five datasets from ECHO. For our analysis sample, we keep quarters from Q1:2007 until Q3:2013. The *Actions* dataset starts shortly before the beginning of this period but we start our data in 2007 to be able to use lagged values of variables. Although this data source supposedly continued through 2014, we noticed fewer reported cases after Q3:2013, which we believe are due to early transitions to the new database. This motivates our choice to end our analysis data in Q3:2013.

We make three main adjustments to our analysis data. First, in some cases, we observe a violation at some quarter t in the *Actions* dataset but the plant is not reported to be a violator in the *Historical Compliance* dataset at quarter t or $t+1$ and did not receive a fine at quarter t. We believe that these violations likely reflect minor issues that are dissimilar to other violations, and hence we exclude them from our analysis. Second, in some cases, we observe a violation at some quarter t in the *Actions* dataset and the plant is reported to be a violator at quarter $t+1$ but not at quarter t. In this case, we assume that the reporting that indicated that the plant was in compliance at quarter t was erroneous, and hence we record the plant as being in violator status at quarter t. Finally, while most of our investments are for plants that are not in compliance, in some cases, we observe investments for plants that are in compliance. We assume that these investments are not for environmental mitigation and hence do not count them as actual investments in our estimation.

Table 1 provides summary statistics on our main analysis data. Our analysis data contain 5,141,180 plant / quarter observations covering 216,118 unique plants, of which 164,988 are present in every quarter of our sample period. As is well-documented in the literature (Evans, 2016), compliance is high: over our entire time frame, 96.4 percent of observations indicate compliance. Compliance is also high when considering individual plants: 89.7 percent of

Table 1: Summary Statistics on Estimation Sample

Status:	Compliance	Regular violator	High priority violator
Regulatory actions:			
Inspection (%)	6.16	16.47	37.52
Fine amount (thousands of $)	0.96	8.30	141.00
	(46.73)	(145.93)	(621.47)
$\mathbb{1}\{\text{Fine} > 0\}$ (%)	0.13	2.26	12.96
Regulatory outcomes:			
Violation (%)	0.21	2.47	8.76
Entrance into HPV status (%)	0.08	1.03	0.00
Plant actions:			
Investment (%)	0.00	3.35	16.85
Investment (from resolution code) (%)	0.00	3.15	15.58
Investment (from NSR) (%)	0.00	0.20	0.43
Investment (from HPV exit) (%)	0.00	0.00	0.91
Dropped investment in compliance (%)	0.15	0.00	0.00
Plant / quarter observations	4,904,291	135,289	49,348

Note: authors' calculations based on estimation sample. Regulatory actions and outcomes are based on lagged status. Plant actions are based on current status.

plants are never out of compliance, while 7.4 percent of plants have at least one quarter in which they have a violation but are never in HPV status, and only 2.8 percent of plants have at least one quarter in which they are in HPV status.

Consistent with Figure 1, plants in compliance are inspected at much lower rates (6.2%) than plants in regular violator status (16.5%) and plants in HPV status (37.5%). Similarly, fines are much higher for violators and even higher for HPVs. Violating plants are more likely to incur further violations. Violating plants are also much more likely to enter HPV status than are plants in compliance.

We find that investment occurs in 3.4% of quarters when a plant is a violator and in 16.9% of quarters when a plant is a HPV. We derive the vast majority of these investments from codes that indicate the resolution of an environmental problem. We derive a much smaller set of investments from New Source Reviews and from exiting high priority violation status. Finally, we observe codes that are indicative of investment in 0.15% of plant / quarters in compliance, but do not count these as investment as noted above.

The National Emissions Inventory (NEI) database

We use the NEI database to evaluate pollution emissions across regulatory states. These data are only available every three years. We use the 2008 and 2011 NEI data, which are during our main analysis data sample period. The plants recorded in the NEI data also do not link perfectly with the plants in the base analysis data. For these and other reasons, we do not use these data in our base estimation. However, we do use them to calculate the mean pollution by regulatory state, which we then use to evaluate the likely pollution levels generated under counterfactual enforcement policies.

We merge the NEI database with the ECHO database using a multi-step procedure. First, we merge the datasets using an incomplete linkage file provided by the EPA. Second, we merge plants based on whether there is an exact match on plant name, street address, zip code, 6-digit NAICS industry, city, and state. Third, we merge plants based on whether there is an exact match on subsets of the above variables (plant name, street address, and city first, then plant name, industry, and city, then plant name and zip code).[8] Our match rate of 60% is similar to other papers that use the NEI data, i.e. Lyubich et al. (2018) report a match rate of 77.4% between the NEI and the Annual Survey of Manufacturers.

Table 2 provides summary statistics on the reported criteria air pollution for our analysis data, by industrial sector.[9] We report in this table the ten industrial sectors with the most number of plant / quarter observations in our analysis data. There is substantial variation in the pollution levels across sectors. The most polluting sector in our data is the manufacturing of wood, petroleum, and pharmaceutical products. Plants in compliance here have about 228 tons per quarter of emissions while plants in HPV status have about 1,960 tons per quarter of emissions. The second highest is the mining and extraction sector. Across the sectors, plants in violator status emit more pollution that plants in compliance. This effect is particularly pronounced for plants in HPV status.

[8]We are in the process of increasing our match rate by also performing an exact matching based on city and state, along with a fuzzy match on plant name, plant street address and zip code.

[9]Criteria air pollutants include sulfur dioxide (SO_2), nitrogen dioxide (NO_2), lead, carbon monoxide (CO), ozone (O_3), and particulate matter (PM).

Table 2: Summary Statistics on Mean Criteria Air Pollution Levels

Sector	Observations in analysis data	Mean level in compliance	Mean level as regular violator	Mean level as HPV
Mining & extraction	837,825	122.9	390.4	1420.2
Manufacturing: wood, petroleum, pharma	818,036	228.1	666.0	1960.2
Other	667,147	37.6	59.9	281.6
Manufacturing: metal	666,108	76.6	116.5	706.7
Other services	519,050	1.6	4.6	32.2
Wholesale trade	234,302	13.3	12.9	49.1
Retail trade	204,476	1.4	4.6	3.0
Manufacturing: food, textiles	186,721	100.3	373.6	247.4
Transportation	176,988	175.1	261.1	316.6
Educational services	155,811	76.4	128.0	157.9

Note: table reports summary statistics on total criteria air pollution levels in tons for plant / quarter observations in our analysis data, matched to the NEI data based on EPA region, industrial sector, and status as being in compliance, a regular violator, or a HPV.

3.2 Empirical Foundations of the Structural Model

In our dynamic model of plant behavior, the plant's decisions are a function of its regulatory state. In principle, the regulatory state lists the plant's history of prior violations and investments and its EPA region and industrial sector. In practice, we need to summarize this information for tractability. We now provide evidence that motivate our state space and other modeling choices.

Investment

We first investigate the role of current and past investment in affecting violator status. Table 3 provides a regression of whether a plant returns to compliance in a period (from regular or high priority violator status) on current investment, and four quarter lags of investment.

We find that investment in the previous quarter is a very strong predictor of a return to compliance, increasing the probability of a return by 42 percentage points. Investment two quarters ago is a weaker, though still statistically significant and positive predictor. In contrast, current investment, and further lags of investment are all negative predictors.[10]

[10] The negative coefficient on current investment may be due to plants investing when their regulatory state gets worse, in the sense of incurring more likely penalties from not investing.

Based on these regressions, our state space allows for two lags of investment to affect the regulatory state. We also assume that current investment does not have any impact on helping a plant return to compliance in the current period, only in the subsequent two periods. Finally, the lack of a current effect of investments motivates our timing assumption that investment is made at the end of each period, after the regulator's actions and regulatory outcomes.

Table 3: Investment and Resolution of Violations

Dependent variable: return to compliance		
Current investment	-0.112^{***}	(0.002)
One quarter lag of investment	0.417^{***}	(0.006)
Two quarters lag of investment	0.105^{***}	(0.006)
Three quarters lag of investment	-0.010^{**}	(0.005)
Four quarters lag of investment	-0.043^{***}	(0.005)
Number of observations	184,637	

Note: regressions include fixed effects for 2-digit NAICS industrial sector and EPA region. Regression uses the estimation sample restricted to plants not in compliance in the previous quarter. Standard errors, which are clustered at the plant level, are in parentheses. ***, **, and * indicates statistical significance at the 1%, 5%, and 10% levels, respectively.

Focusing now on investment in the previous quarter, Figure 2 shows in more depth the frequency with which this investment resulted in a return to compliance. If the plant starts the period in HPV status and did not invest in the previous quarter then it will, with certainty, finish the quarter in HPV status. If the plant did invest, there is still a 24% chance that it will finish the period in HPV status, but there is now a 50% chance that the plant will transition to compliance and a 26% chance that the plant will transition to regular violator status. Lagged investment similarly increases the rate at which the plant transitions from regular violator status to compliance, although some plants do transition from regular violator status to compliance even without investment. Thus, overall, investment will increase the probability that a plant returns to compliance, but does not result in compliance with certainty.

Figure 2: Effect of investment on regulatory state

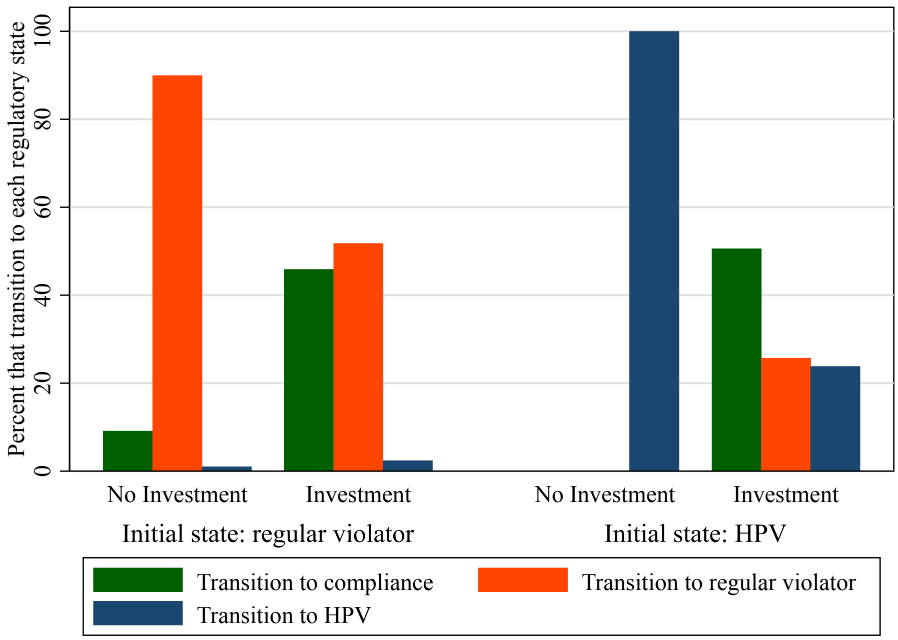

Note: authors' calculations based on estimation sample. Initial state and investment are from previous quarter.

Depreciated Accumulated Violations

Table 1 showed that inspections, fines, and violations all varied substantially based on whether the plant is in compliance, a regular violator, or an HPV. We investigate here whether, even within these three broad categories, previous violations are predictive of inspections, fines, and violations. We define a summary measure called "depreciated accumulated violations" which, for plants out of compliance, is the sum of the depreciated violations, from the previous quarter back to the period the plant most recently left compliance. Our baseline results use a 10% quarterly discount factor to calculate this measure.

Figure 3 displays the relation between depreciated accumulated violations and inspections, the probability of having a positive fine, and violations. The figure splits the results into regulator and high priority violators; plants in compliance have a value of zero for depreciated accumulated violations, by construction. We find that the number of depreciated accumulated violations is a strong and positive predictor of all of these events, for both

Figure 3: Depreciated accumulated violations and monitoring and enforcement

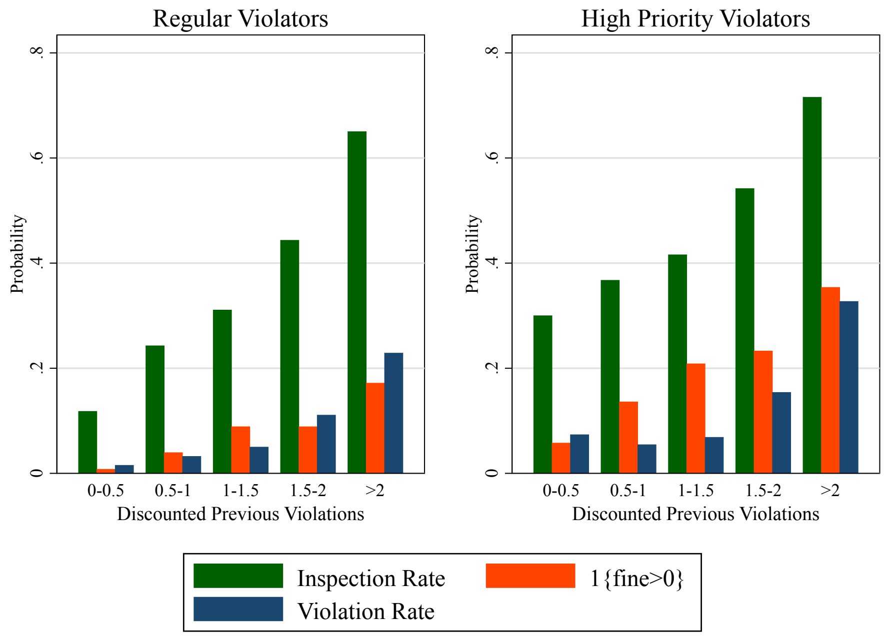

Note: authors' calculations based on estimation sample.

regular and high priority violators.

Having established that depreciated accumulated violations are a significant predictor of further regulatory actions and outcomes, we next investigate which discount factor is the most appropriate one to use. Table 4 regresses inspections, fines, and violations on depreciated accumulated violations using three different discount factors, along with lagged HPV status. We find that the measure with the 10% discount factor is much more predictive of inspections and fines than the measure with a 20% discount factor or with a 0% discount factor. These results motivate our choice of including depreciated accumulated violations with a 10% quarterly discount factor as a state variable.

Heterogeneity in Costs and Technologies

Finally, we investigate the extent of observed heterogeneity in the regulatory process. In particular, identification of our model will be aided by having heterogeneity in how inspections and fines increase with the regulatory state across EPA regions and industries. Figure 4

Table 4: Regressions of Regulatory Actions on Depreciated Accumulated Violations

Dependent variable:	Inspection	Fine amount	Violation
Accumulated violations with no depreciation	0.007	−0.014***	0.001
	(0.006)	(0.003)	(0.012)
Accumulated violations with 10% depreciation	0.119***	0.129***	−0.047
	(0.021)	(0.013)	(0.071)
Accumulated violations with 20% depreciation	−0.019	−0.067***	0.043
	(0.018)	(0.011)	(0.077)
Lagged HPV status	0.114***	0.032***	0.110***
	(0.005)	(0.002)	(0.030)
Number of observations	184,637	184,637	184,637

Note: regressions include fixed effects for 2-digit NAICS industrial sector and EPA region. Regression uses the estimation sample restricted to plants not in compliance in the previous quarter. Standard errors, which are clustered at the plant level, are in parentheses. ***, **, and * indicates statistical significance at the 1%, 5%, and 10% levels, respectively.

Figure 4: Mean inspection probability and fines by EPA region

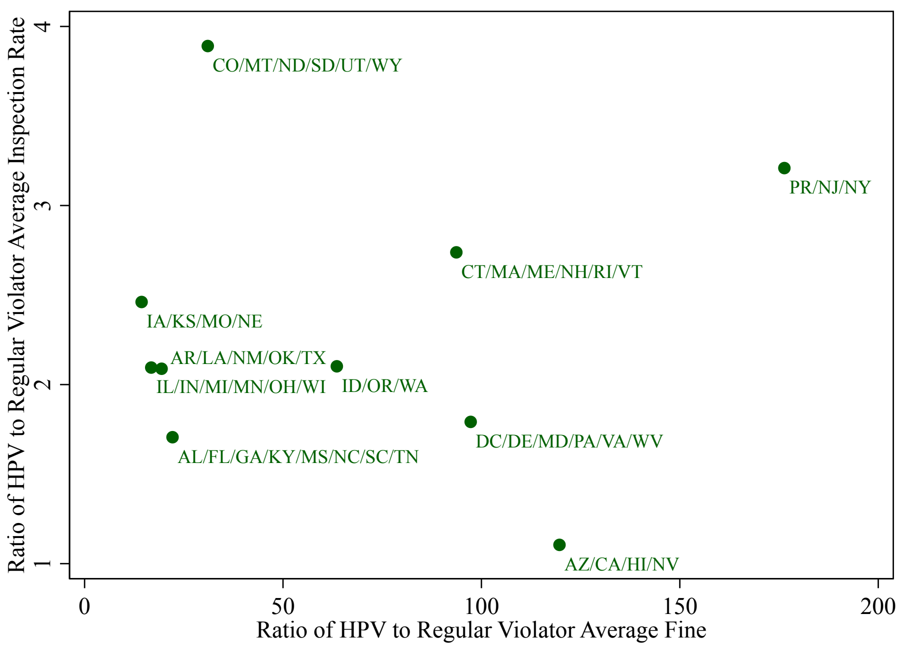

Note: authors' calculations based on estimation sample. States in each region are indicated next to value.

provides a scatter plot by EPA region of the ratio between the inspection probability in HPV status to regular violator status against the ratio of fines in HPV to regular violator status. Figure 5 provides the same information by industry. Importantly, the figures show that there is substantial variation in both of these measures. For instance, the inspection

Figure 5: Mean inspection probability and fines by EPA region

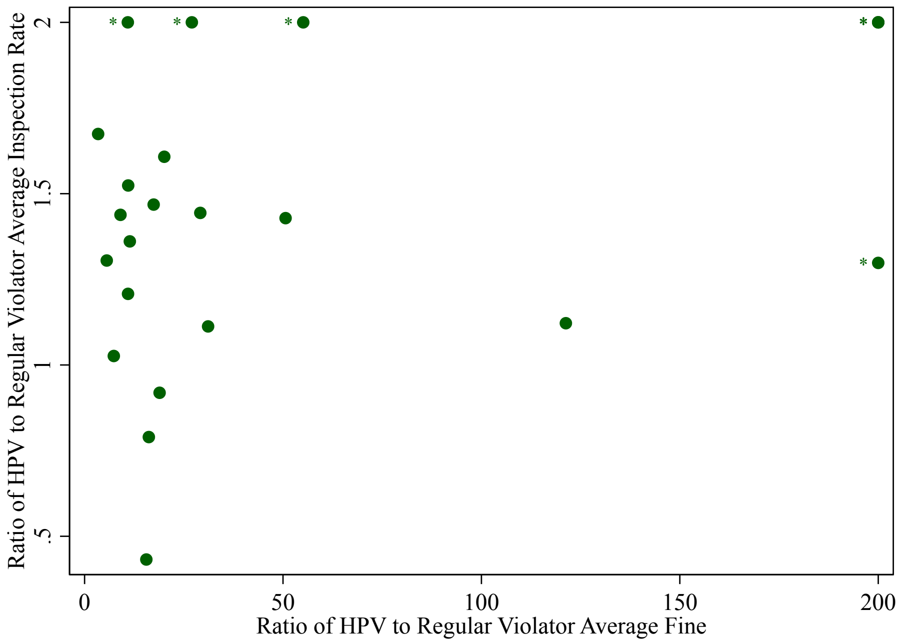

Note: authors' calculations based on estimation sample. Industries marked with a * are top-coded at a inspection ratio of 2 and/or a fine ratio of 200. The highest inspection probability ratio in an industry is over 7 and the highest fine ratio is over 1200.

probability varies across regions by a factor of 4 between the region with the lowest and highest inspection probabilities while the mean fine varies even more across regions.

In addition, these two measures are not significantly correlated with each other across regions, and the correlation across industries is noisy. This shows that there are multiple sources of variation that can help identify the structural parameters and reduces the concern that plants may sort to regions with particularly lenient environmental regulations based on plant-level unobservable characteristics. These results motivate our inclusion of EPA region and industrial sector as state variables.

Finally, we evaluate the extent to which there is heterogeneity in plant investments in our data across plants, as heterogeneity in this choice may reflect heterogeneous investment costs. Figure 6 calculates the mean total number of additional investments in the six quarters after each investment. We compare the means from the data with bars that show the rates that we would observe if investment were *i.i.d.* across our data. Our data exhibit substantially more

Figure 6: Further investments by a plant following initial investment

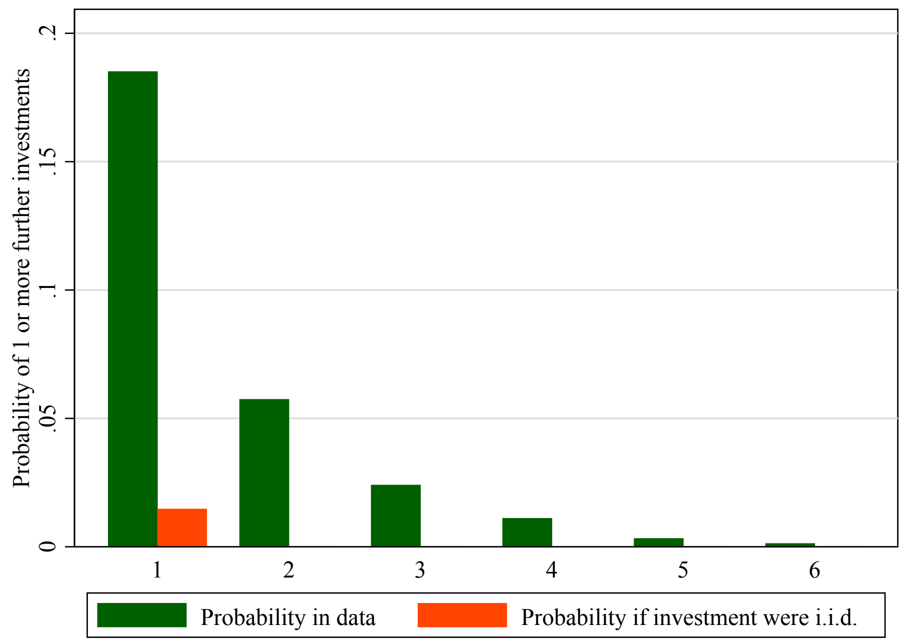

Note: authors' calculations based on estimation sample. The *i.i.d.* model is a hypothetical using the sample mean investment probability.

serial correlation in investment than we would expect to occur randomly. For instance, about 28% of investments are followed by at least one additional investment within the next six quarters, relative to the approximately 1.5% we would observe if investment were *i.i.d.* This suggests that some plants face lower investment costs (or higher regulatory costs) than others, which leads them to invest repeatedly. In the case where there is unobserved heterogeneity in costs in our data, a random coefficients model, such as the one we develop, may be important.

4 Model

4.1 Overview

We develop a framework to estimate plant behavior in the presence of dynamic regulatory monitoring and enforcement. We model each plant as a single-agent dynamic actor. Therefore in this section, we will consider the regulator faced with a single plant. Each period t

corresponds to a quarter. The plant discounts future periods with a discount factor β and chooses its strategy at each state to maximize the expected discounted value of its future utility.

The regulator is interested in having the plant comply with environmental standards. The regulatory actions are inspections, fines, violations, and the designation of the plant as being in compliance, a violator, or a HPV. In our model, the regulator first chooses whether or not to inspect. The regulator then receives a signal about the plant's environmental performance, based in part on its inspection. Using the signal, the regulator decides how much to fine the plant, if at all. It then follows the signal and environmental norms in deciding whether or not to issue a violation to the plant and/or revise its designation of whether the plant is in compliance, a violator, or a HPV.

The timing of our model assumes that these regulatory actions and the plant's observation of whether it is inspected or receives a violation, how much it is fined, and its updated compliance status, are the first thing to occur each period. Following this, each plant that is not in compliance makes a binary decision of whether or not to invest in pollution abatement technology. A plant that invests incurs a cost from its investment, but increases the chance that it returns to compliance in future periods.

We assume that the regulator commits to an ex ante stochastic policy rule regarding its regulatory decisions. Given the many unobserved constraints faced by the regulator, we do not solve for the optimal regulatory policy, but rather estimate plant utility for a given regulatory policy and also compute how equilibrium outcomes would change from different regulatory policies. We now detail the regulator's decisions and plant's actions and payoffs and then discuss dynamic optimization.

4.2 Regulator and Plant Decisions

We define the *regulatory state* to be the set of variables over which the regulator's enforcement activities and designations may depend. The regulatory state in our model has five components: (1) depreciated accumulated violations with a 10% quarterly discount rate, (2) current violator or high priority violator status, (3) two quarterly lags of investment, (4) the EPA region, and (5) the two-digit NAICS industrial sector. As we discussed in Section 3.2, we believe that these summary measures capture the most important features that underlie a plant's utility and transitions.

Denote the plant's regulatory state at the start of period t as Ω_t. The regulatory state is known to the regulator and plant at the start of the period. In principle, we could allow for the entire history of lagged investments and violations to affect regulator decisions. We limit the regulatory state to the summary measures of the depreciated sum of past accumulated violations and recent investments for tractability. The EPA region and industrial sector are very significant predictors of monitoring and enforcement that will help identify our parameters.

The regulator commits ex ante to inspection, violation, transition, and fine policy rules. The inspection policy rule specifies the probability of an inspection given the regulatory state of the plant. Denote this rule $\mathcal{I}$; $\mathcal{I}(\Omega)$ is the inspection probability at state Ω. The violation and transition rules specify the violation and transition probabilities respectively as a function of the signal it receives, in part from the inspection.[11] The fine rule specifies the size of the fine, if any, that the plant will receive, as a function of its regulatory state and signal information.

Each period starts with the potential for an inspection. Let $I(\Omega) \sim Binom(\mathcal{I}(\Omega))$ denote the binomial random variable indicating an inspection in state Ω. If the regulator decides on an inspection, the plant is then inspected.

Following the inspection phase, the regulator receives a signal e of environmental compliance; we discuss the distribution of e below. The regulator assigns a violation to the plant based on some function $V(\Omega, e)$, which was chosen ex ante and whose range is binary. Following the potential assignment of violation, the regulator decides how much to fine the plant, if at all. Here, the regulator uses an ex ante policy rule $F(\Omega, e)$.

Finally, the regulator decides on the transition from Ω to the state that the plant will face

[11] Our counterfactuals keep these three rules fixed since different rules may not be consistent with environmental standards which we do not observe.

at the point when it takes its action, Ω'. Let $T(\Omega, e)$ denote this transition. T indicates the transition into designation as being in compliance, a violator, or a HPV, and also includes the result of $V(\Omega, e)$ as this will update the depreciated sum of past accumulated violations.

Let $H(T)$ be an indicator for T designating HPV status, and $C(T)$ be an indicator for T designating compliance. The plant obtains (dis)utility from four regulatory actions: inspections, violations, fines, and HPV status. Specifically, the flow utility for a plant is:

$$U(\Omega, e) = \theta^I I(\Omega) + \theta^F F(\Omega, e) + \theta^V V(\Omega, e) + \theta^H H(T(\Omega, e)). \tag{1}$$

where $\theta^I, \theta^F, \theta^V, \theta^H$ are parameters.

We make the following assumption about the nature of the distribution of the environmental compliance signal e.

Assumption 1. *The environmental compliance signal at time t, e_t, is a function only of Ω_t, I_t, and ex ante policy rules $\mathcal{I}$, $V(\cdot,\cdot)$, $T(\cdot,\cdot)$.*

Assumption 1 rules out the possibility that an investment that is not in the regulatory state (for instance one that occurred many periods ago) could change the compliance signal. We keep two lags of investment in the regulatory state, which can therefore affect the compliance signal. Assumption 1 also implies that e does not directly affect the future state. However, we include violator status and depreciated accumulated violations in the future regulatory state, which will be a function of past violations and hence indirectly a function of e.

We allow for e_t to depend on the policy rules for inspections, violations, and transitions. Given this, the findings of an inspection are allowed to be a function of whether the plant is inspected often given its regulatory state. Importantly, we do not allow for e_t to depend on the fine policy. This is useful to obtain counterfactual computations because it implies that the state-contingent probabilities of violations and transitions will remain the same even with counterfactual fine policies. Note that this does not preclude the fact that under different fine structures, plants will invest differently, and that the different investments will yield different steady states under plant optimizing behavior. This again occurs because we allow lagged investment to enter into the state space and hence lagged investment can affect e and through that, state transitions.

Assumption 1 implies that the regulatory state Ω'_t is the only state that matters to the plant at time t. The reason for this is that, for any time t action, actions and transitions at time $t+1$ depend only on Ω'_{t+1}, I_{t+1} and e_{t+1}, which are all functions of Ω_{t+1}.

Having established that the plant's state at its time t decision point is Ω'_t, we now discuss the plant's decision process. In any period in which Ω' indicates that the plant is a violator or HPV, the plant makes a decision of whether to undertake a pollution-abating investment, X. The utility (or negative of the cost) from investment is θ^X. In addition to θ^X, the plant receives an idiosyncratic shock to its cost of not investing, ε_{0t}, and, when it is not in compliance, an idiosyncratic shock to its cost of investing, ε_{1t}. We assume that these shocks are *i.i.d.*, known to the plant prior to making its investment decision, and distributed type 1 extreme value.

Let $X = 1$ denote investment and $X = 0$ denote no investment. Then, the flow utility to the plant at this stage from action X is $X\theta^X + \varepsilon_{Xt}$. While investment is costly, the benefit to an investment at period t is that it helps clear violations in periods $t+1$ and $t+2$. Specifically, it enters into the state variable for the future regulatory enforcement and transition decisions.

Group together the structural parameters as $\theta \equiv (\theta^I, \theta^F, \theta^V, \theta^H, \theta^X)$. We assume that θ is fixed for the plant over time. In some specifications, θ will vary across plants. However, we do not allow the regulator to condition its monitoring and enforcement policies on each plant's individual values of these parameters. We would generally expect these parameters to be negative.

4.3 Dynamic Optimization

A plant that is not in compliance makes an investment decision in each period, knowing that the investment will reduce its expected future cost of regulatory enforcement. The plant's optimization therefore requires evaluating the value of being in a given state, Ω, at the start of the next period. We now exposit the value function of the plant and its probability of choosing investment at any state, conditional on its parameter vector θ.

Let $V(\Omega)$ denote the value function at the beginning of the period and let $V'(\Omega')$ denote the value function at the point right after the regulator has moved but before the plant receives its draws of ε. We start by expositing $V'(\Omega')$:

$$V'(\Omega') = C(\Omega')[\beta V(\Omega', \theta) + \gamma] \tag{2}$$
$$+(1 - C(\Omega'))[\ln(\exp(\beta V(\Omega|\Omega', X = 0)) + \exp(-\theta^X + \beta V(\Omega|\Omega', X = 1)) + \gamma].$$

The first line of (2) reflects the case of compliance. In this case, the plant transitions to the same state Ω' in the next period. Since there is no plant choice here, in expectation, the plant receives the mean value of the type 1 extreme value distribution which is γ, Euler's constant. The second line of (2) reflects the case of a plant that is a violator or high priority violator. In this case, it makes a choice of whether to invest or not. Since the value is computed ex ante to the realization of the idiosyncratic draws, we can use the familiar logit aggregation. In this case, the transition state next period depends on whether investment is made but is not stochastic. The transition state is also not the same as the current state, because the depreciated sum of past accumulated violations is updated at this point.

Having exposited $V'(\Omega')$, we now exposit $V(\Omega)$, the value function at the beginning of the period:

$$V(\Omega) = \sum_{I \in 0,1} \mathcal{I}(\Omega)^I (1 - \mathcal{I}(\Omega))^{1-I} \int [U(\Omega, e) + V'(T(\Omega, e))] \, dP(e|I, \Omega), \tag{3}$$

where $dP(e|I(\Omega))$ is the integral over environmental compliance signal e given the inspection decision and plant state. Note that the plant does not make any decision at the beginning of the period, and hence there is no maximization in (3). However, the plant must integrate over the two possible inspection decisions and regulatory actions and outcomes that can occur.

Finally, having defined value functions, we derive the probability of a plant choosing investment given a regulatory state Ω^d and its cost and utility parameters θ as:

$$\Pr(X = 1|\Omega', \theta) = \frac{\mathbb{1}\{C(\Omega') = 0\} \exp(\theta^X + \beta V(\Omega|\Omega', X = 1))}{\exp(\theta^X + \beta V(\Omega|\Omega', X = 1)) + \exp(\beta V(\Omega|\Omega', X = 0))}. \tag{4}$$

The probability in (4) will be used in deriving our estimators below. Accordingly, we have explicitly allowed it to be a function of the structural parameter vector θ.

5 Empirical Implementation

We first discuss the estimation of our two specifications, one with fixed coefficients across plants and the other with random coefficients. We then discuss identification of both specifications.

5.1 Estimation of Model with Homogeneous Coefficients

The goal of the estimation process is to recover θ, which reports the (dis)utility to plants from regulatory action and from investment in pollution mitigation technologies. Our first model assumes that θ is the same across plants.

Our data include a panel of plants i observed over time periods t and are at the level of the plant/time period. For each plant/time period, we observe the regulatory state at the point where the plant makes its investment decision, which we denote Ω'_{it} and the investment decision, X_{it}.

For this model, we employ a maximum likelihood estimator. In this model, there are no serially correlated unobservables for a plant over time, and hence, we can treat each plant / quarter as an independent observation. The log likelihood of a parameter vector θ is:

$$\log L(\theta) = \sum_i \sum_t \log([X_{it} Pr(X = 1|\Omega'_{it}, \theta) + (1 - X_{it})(1 - Pr(X = 1|\Omega'_{it}, \theta))]), \tag{5}$$

where the investment probabilities are calculated using (4). We use a nested fixed point estimator here; the investment probability equation in (4) requires solving the Bellman equation.

Our estimator here is similar to Rust (1987). We could also use conditional choice probability estimators (Hotz and Miller, 1993; Arcidiacono and Miller, 2011), which are quicker to compute, but we did not, since the computational time for the nested fixed point maximum likelihood estimator is not excessive.

Similar to Rust (1987), we estimate the evolution of the regulatory state conditional on the investment decision in a first stage. However the evolution here is more complex as it indicates the policy of the regulator, which involves as a function of its random inspection decision and its draw of the environmental compliance signal e. We estimate the regulatory outcome using reduced form specifications that condition on the plant's state. Specifically, we estimate separate probit regressions for plants in compliance, regular violators, and HPV violators that indicate the probability of an inspection at any state conditional on two lags of investment, depreciated accumulated violations, and the plant's region and industry. We estimate similar probit regressions for violations that condition on whether an inspection occurred, and tobit regressions for fines that condition on inspection and violation. We present estimated coefficients for these regulatory actions in Appendix Table 1. We also estimate similar reduced form specifications of the probability that plants transition between states, based on $T(\Omega, e)$. We present estimated coefficients for transition equations in Appendix Table 2. As discussed in Section **??**, EPA region and industrial sector bring a lot of identifying variation to our analysis. We operationalize this by allowing these variables to enter in all the regulatory regressions.

5.2 Estimation of Model with Random Coefficients

Our second model allows for the parameter vector θ to differ across plants. Specifically, in this model, we assume that θ for each plant takes on one of a fixed set of values $(\theta_1, ..., \theta_J)$. Each θ_j, $j = 1, ..., J$ occurs with probability η_j. Each plant receives an independent draw of θ from the multinomial distribution of potential values. The structural parameters to be estimated are therefore $\eta \equiv (\eta_1, ..., \eta_J)$. We impose no restriction on the structural parameters other than what is necessary based on the fact that the structural parameters are population probabilities:

$$\sum_{j=1}^{J} \eta_j = 1 \text{ and } 0 \leq \eta_j \leq 1, \ \forall j. \tag{6}$$

Thus, the structural parameters of this model are $\eta \equiv (\eta_1, ..., \eta_J)$ and no longer $(\theta_1, ..., \theta_J)$. Econometrically, the values of $(\theta_1, ..., \theta_J)$ are taken as given. We take a (large) fixed grid of these values, meant to capture the range of plausible parameter values.

We estimate the parameters here by adapting the methods of Fox et al. (2011) and Nevo et al. (2016). Specifically, we develop a computationally quick GMM estimator, allowing us to estimate many parameters. Essentially then, the model allows for a non-parametric density over the θ utility parameters.

Specifically, our GMM estimator has the form $\eta^* = \arg\min_\eta \|G(\eta)\|$, where $G(\eta)$ is a $K \times 1$ vector of moments and $\|\cdot\|$ is a simple squared norm. Each individual moment $G_k(\eta)$, $k = 1, \ldots, K$, can be written as the difference between the value of some statistic in the data, m_k^d and the weighted sum of the value of the statistic for the parametrized model, $m_k(\theta_j)$, where the weights are given by η_j:

$$G_k(\eta) = m_k^d - \sum_{j=1}^{J} \eta_j m_k(\theta_j). \tag{7}$$

We compute each m_k^d and $m_k(\theta_j)$ in an initial stage, before estimating η. This requires solving the Bellman equation and $m_k(\theta_j)$ for each of the J grid parameters. Using these values, we then estimate η by minimizing the sum of squared values of $G_k(\eta)$ subject only to the constraints in (6). This latter step is computationally very quick.

Because we do not see plants from their inception onwards, we need to make some assumption about the likelihood of seeing any plant in any state based on its underlying θ_j. First, define the following. Split the state Ω' into Ω'^1, which indicates the variable states of compliance status, lagged depreciated accumulated violations, current violation, and lagged investment, and Ω'^2, which indicates the fixed states of EPA region and industrial sector. Using this definition, we make the following assumption for our random coefficients estimation:

Assumption 2. *The observed data reflect plants that are at the steady state distribution of* Ω'^1 *conditional on a given* Ω'^2*.*

Assumption 2 would be valid if, for instance, plants enter at randomly distributed points from the steady state distribution of Ω'^1. It would also occur if they have been active a long time, in which case the distribution of Ω'^1 for any θ_j value would approach its steady state distributions.

We now discuss our specific moments, which use Assumption 2 in order to be computable. Our first set of moments is constructed from m_k values that indicate the probability of being at a particular time-varying state in equilibrium, conditional on EPA region and industrial sector. Specifically, for any $\omega'^1 \in \Omega'^1$ and $\omega'^2 \in \Omega'^2$, we can write:

$$m_k(\theta_j) = \Pr[\Omega'^1 = \omega'^1 | \Omega'^2 = \omega'^2, \theta_j], \tag{8}$$

and

$$m_k^d = \sum_i \sum_t \frac{\mathbb{1}\{\Omega_{it}'^1 = \omega'^1, \Omega_{it}'^2 = \omega'^2\}}{\mathbb{1}\{\Omega_i'^2 = \omega'^2\}}. \tag{9}$$

We compute (8) by solving the Bellman equation and then evaluating the steady state distribution under optimizing behavior.

We note a few points about these moments. These moments follow closely from Nevo et al. (2016), although we use the steady state distribution of our infinite-horizon dynamic problem, while they use the actual distribution of their finite-horizon problem. While in principle we could construct a moment from every Ω', as we have over 30,000 states, we limit ourselves to states which have a probability over 10^{-4} at our estimated maximum likelihood parameters. Finally, we normalize all moments by their standard deviation at our estimated maximum likelihood parameters, in order to increase efficiency.

Our second set of moments also follows closely from Nevo et al. (2016). The m_k values for these moments are constructed from the probability of being at a particular state times the investment probability at that state:

$$m_k(\theta_j) = \Pr[\Omega'^1 = \omega'^1 | \Omega'^2 = \omega'^2, \theta_j] \Pr(X = 1 | \Omega', \theta_j), \tag{10}$$

and

$$m_k^d = \sum_i \sum_t \frac{\mathbb{1}\{\Omega_{it}'^1 = \omega'^1, \Omega_{it}'^2 = \omega'^2, X_{it} = 1\}}{\mathbb{1}\{\Omega_i'^2 = \omega'^2\}}. \tag{11}$$

We compute these moments for every state for which we compute our first set of moments, except for states that reflect compliance, as there is no investment in these states.

Our final set of moments explicitly capture the panel data aspect of investment. The m_k values for these moments are constructed from the probability of being at a particular state times the expected value of investment at that state times the sum of investments in the next six periods:

$$m_k(\theta_j) = \Pr[\Omega'^1 = \omega'^1 | \Omega'^2 = \omega'^2, \theta_j] \Pr(X = 1 | \Omega', \theta_j) \times \tag{12}$$
$$\textstyle\sum_{s=1}^{6} E[\Pr(X \text{ } s \text{ periods ahead} = 1 | X = 1, \Omega', \theta_j)],$$

and

$$m_k^d = \sum_i \sum_t \frac{\mathbb{1}\{\Omega_{it}'^1 = \omega'^1, \Omega_{it}'^2 = \omega'^2, X_{it} = 1\} \sum_{s=1}^{6} X_{i,t+s}}{\mathbb{1}\{\Omega_i'^2 = \omega'^2\}}. \tag{13}$$

These moments seek to match the extent of repeated investments by plants in the data—as displayed in Figure 6—to the model. A more traditional correlation moment would simply multiply investment at time t with investment at time $t+1$ rather than with investment over the following six periods. We chose this formulation because we worry that investment in two subsequent quarters might partly reflect measurement error. We compute these moments for every state for which we compute our second set of moments. To calculate the investment in the 6 periods ahead in (12), we simulate over all potential paths conditioning on the initial state and investment decision.

Finally, we discuss our parameter grid. For each of the five θ components, θ^I, θ^F, θ^V, θ^H, and θ^X, we fix minimum and maximum values that capture a plausible set of potential parameters. We then use co-prime Halton sequences (using the first five prime numbers) over each component (Train, 2009), taken uniformly over the range between the minimum and maximum values. We use the first 10,000 elements of the Halton sequences as our parameter values; hence $J = 10,000$. We use Halton sequences because we expect that co-prime Halton sequences better cover the plausible parameters that would taking the interaction of the same grid points for each component.

5.3 Identification

As discussed above, we estimate regulatory actions and outcomes as well as the structural parameters underlying the plant's utility function. The regulatory actions and outcomes are estimated with simple reduced form specifications that are intended to capture the plants' expectations of regulatory actions rather than the underlying causes of regulator decision-making. Hence, identification of these functions derives from observing data across regulatory states. Since these functions represent plants' beliefs regarding the future, for them to be valid in the context of our model we need plants to not have private information about future regulatory actions and outcomes beyond the functions that we estimate. Our specifications all include fixed effects at the EPA region and industry level and also a variety of interactions, in order to accurately capture plants' beliefs.

Identification of the structural parameters θ for our model with homogeneous coefficients is somewhat different. Consider first a very simple version of this model where plants pay a cost from investment and also have a cost from fines, but do not face disutility from inspections, violations, or HPV status. This model then has two parameters.

At any violator state, a plant can calculate its expected discounted future fine if it does or does not invest. This expected future fine would condition on the regulatory actions and outcomes (which we have previously estimated) and on future actions of the plant. Let the gross value of investing be the difference in expected fines between investing and not investing. Conditional on investment costs, we would expect to see a plant invest more the higher is this gross value. The extent to which the gross value of investment results in actual investment then identifies the ratio of the cost of fines to the cost of investment.

Having discussed identification of the ratio of the two parameters, we now discuss the identification of the scale of the parameters. As is true in any logistic model, by choosing a type 1 extreme value distribution, we effectively normalize the variance of ε. The other parameters can be interpreted as their true values divided by the standard deviation of ε. Within this context, the scale of these two parameters is then identified by the variance in the investment actions within states. If, on one hand, we see a knife-edged pattern where for states below some gross value, plants never invest and for states above this value they always invest, then the scale of the estimated parameters will be large, because the variance of ε is small. If, on the other hand, we see a gradual increase in the investment probability as the gross value increases, then the scale of the estimated parameters estimated will be small, because the variance of ε is large.

In order to identify this simple model, we need variation across states in the gross value of investment, which stems from variation in the difference in expected discounted future fines from investing and not investing across states. Our data will provide this variation both for different states within an EPA region and industrial sector, and across regions and sectors.

Our actual model includes five parameters, which capture four dimensions of regulatory costs borne by the plant, plus the cost of investment. Thus to identify this model, we need independent variation in the difference, between investing and not investing, in the expected discounted future values of each of the four regulatory levels. While there is some variation in these values for different states within an EPA region and industrial sector, we believe that, in practice, variation across EPA regions and industrial sectors is very helpful in identifying these parameters. In particular, Figure 4 documents that there is substantial variation in inspections and fines across EPA regions.

Our model with random coefficients adds another dimension to our identification argument. Here, we need to identify the distribution of the values of θ, not just the mean values of these parameters. In part, we identify this distribution from the equilibrium distribution of Ω'^1 and of investment interacted with Ω'^1. For instance, if there is more spread across states in the equilibrium distribution of states, this will be consistent with more variance in

θ^X, all else equal.

However, to estimate random coefficients, it will also be useful to directly use the panel aspect of our data. We do this with our correlation moments based on (12) and (13). In particular, if we see some plants that invest often and others that rarely invest when faced with the regulatory state, then this will be evidence that there is heterogeneity in the utility parameters. The extent to which this heterogeneity varies based on regulatory actions and outcomes will then identify in which of the utility parameters the heterogeneity lies. Figure 6 provides reduced-form evidence that this heterogeneity is substantial, suggesting that these moments will provide identifying variation in practice.

6 Results

6.1 Model Estimates

Table 5: Estimates of Structural Parameters

	ML estimates	GMM random coefficient estimates (1)	(2)	(3)	(4)	(5)	(6)	
Investment utility (θ^X)	-2.993^{***}	-2.993	-1.598	-1.252	-1.213	-1.790	-1.716	-3.086
	(0.023)							
Inspection utility (θ^I)	-0.039	-0.039	-0.398	0.121	0.663	0.602	0.875	0.392
	(0.026)							
Violation utility (θ^V)	-0.129^{*}	-0.129	-0.236	-2.118	-0.121	0.002	-0.283	0.866
	(0.068)							
Fine utility (θ^F)	-0.120^{***}	-0.120	-0.293	-0.093	-0.008	-0.129	-0.472	-0.504
	(0.004)							
HPV status utility (θ^H)	-0.055^{***}	-0.055	0.165	0.050	-0.532	-0.495	-0.122	0.069
	(0.007)							
Weight on parameter vector	1	0.595	0.145	0.133	0.071	0.046	0.005	0.005

Note: standard errors for maximum likelihood estimates, which are bootstrapped with resampling at the plant level, are in parentheses. ***, **, and * indicates statistical significance at the 1%, 5%, and 10% levels, respectively. For GMM estimates, we report weights on all types j with probability $\eta_j > 10^{-3}$.

As discussed in Section 5, we estimate two different specifications of our model. The first specification has the same utility coefficients across plants and we estimate it with maximum likelihood. The second specification has random coefficients and we estimate it with GMM. Table 5 provides the estimation results for plants' utility functions θ for both specifications.

We start with the maximum likelihood results, which are on the left of Table 5. We find that investments, inspections, violations, fines, and being in HPV status are all costly for plants. Except for violations, the coefficients are all significant at the 1% level. Violations are significant at the 10% level. Note that our finding is that a current inspection or violation imposes a cost on the plant's current utility, in addition to the impact of these actions on the future regulatory state, which is also costly to plants. In terms of relative magnitudes, an inspection costs the plant the equivalent of a $800,000 fine (0.105/0.131). Similarly, a violation costs the plant the equivalent of a $900,000 fine.

We also find that investments are very costly to plants. The current utility cost of an investment is equivalent to about $25 million in fines. While this large magnitude may be in part because we are estimating a relatively small coefficient on fines, it is also likely because our investment measure captures large investments in pollution abatement.

Finally, we find that being in HPV status is also costly, with as quarterly cost to plants of about $350,000. The fact that HPV status is costly to plants suggests that the stigma from HPV status is an effective tool for ensuring compliance with environmental standards. The dynamic model is critical for understanding that HPV status itself is costly to plants, and also to understanding the magnitude of this coefficient.

We next turn to the fixed grid random coefficients estimates, which are in the remaining columns of Table 5. Recall that we allow the parameter vector θ for each plant to be one of $J = 10,000$ elements, chosen over a wide grid of potential values. For this specification, we report the probability of each of the parameters θ_j for which the population weight η_j is greater than 10^{-3}. We list the θ_j parameters in descending order of η_j. Thus, the parameters that are estimated here are the weight parameters, which are given in the last row of the table. We do not report standard errors for this specification as it would be difficult both to calculate them and to interpret them meaningfully, given that most of the estimated weights are 0. Instead, we report standard errors for our counterfactuals below, which we calculate by bootstrapping our entire estimation process at the plant level (estimation of the regulatory policy and transition, calculation of moments, and estimation of the weight parameters).

Our GMM specification finds that 99% of plants can be represented by one of three parameter values, or types, given by columns (1) - (3) of this part of the table and only 6 parameter values have a weight of more than 10^{-3}.[12] We find further that there is substantial heterogeneity in the value of θ across plants, with the mean values roughly similar to the ML estimates across the five dimensions of θ.

The largest fraction of plants—the 54% of plants in type (1)—has an investment cost that is 76% the size of the estimated investment cost in the ML estimates (2.521/3.334). However, another 25%—those in type (2)—have an investment cost that is over double the size in the ML estimates. Interestingly, the ratio of investment cost to fines is very similar for these two parameter values and also similar to the ML estimates. However, type (3), which occurs with probability 20%, has an investment cost that is only equivalent to $13 million in fines, suggesting that a substantial minority of plants find investment less costly than estimated by our ML specification.

Another difference here from the ML results is that the magnitude of the coefficients on inspections, violations, and HPV status are all much larger in absolute value than in the ML results. Moreover, while the majority of plants have negative values of these coefficients, there is a substantial minority that have positive values for them. Our takeaway is that there is substantial heterogeneity in plants' investment responses when faced with these variables and our large and heterogeneous estimated values for these coefficients capture this heterogeneity.

6.2 Counterfactuals

Using the coefficient estimates from Table 5, we now model how EPA enforcement activities, plant investments, overall compliance, and criteria air pollution would change if the EPA changed how it allocates fines or if plants did not value dimensions of regulator enforcement other than fines. From Assumption 1, the state-contingent environmental compliance signal is a function only of the inspection, violation, and state transition rules. This means that provided we keep these state-contingent regulatory policies the same, the distribution of

[12] This is consistent with Fox et al. (2016), who provide Monte Carlo evidence of the fixed grid estimator as an approximation to a model with continuous random coefficients. They also find few grid values with positive weights.

Table 6: Counterfactual Results: Changing Fine Structure

	(1) Baseline	(2) Same fines for all regulator and high priority violators	(3) Fines for HPVs halved relative to baseline	(4) Fines for HPVs doubled relative to baseline
Maximum likelihood estimates				
Compliance rate (%)	94.84 (0.29)	94.00 (0.70)	94.37 (0.35)	95.26 (0.27)
Regular violator rate (%)	3.77 (0.24)	3.70 (0.24)	3.76 (0.24)	3.78 (0.24)
HPV rate (%)	1.38 (0.12)	2.30 (0.59)	1.87 (0.20)	0.97 (0.08)
Investment rate (%)	0.41 (0.02)	0.40 (0.03)	0.40 (0.02)	0.41 (0.02)
Inspection rate (%)	8.78 (0.18)	9.14 (0.46)	8.95 (0.20)	8.65 (0.17)
Mean fines ($ thousands)	17.33 (3.02)	17.33 (3.02)	14.58 (3.34)	19.67 (3.12)
Mean violations (%)	0.50 (0.04)	0.66 (0.31)	0.57 (0.06)	0.45 (0.03)
Mean plant utility	−0.006 (0.004)	−0.006 (0.005)	−0.005 (0.004)	−0.007 (0.004)
Mean NOX (tons)	79.0 (0.7)	83.3 (1.9)	81.4 (1.0)	76.9 (0.5)
Mean SOX (tons)	109.8 (1.1)	116.0 (2.9)	113.8 (1.7)	106.4 (0.8)
Mean air pollutants (tons)	268.4 (2.8)	285.9 (7.6)	278.7 (4.3)	259.6 (2.0)
GMM random coefficient estimates				
Compliance rate (%)	95.36 (0.25)	93.90 (3.29)	94.76 (0.52)	95.74 (0.23)
Regular violator rate (%)	3.37 (0.20)	3.28 (0.21)	3.36 (0.20)	3.37 (0.20)
HPV rate (%)	1.27 (0.12)	2.81 (3.35)	1.88 (0.41)	0.88 (0.08)
Investment rate (%)	0.48 (0.03)	0.47 (0.05)	0.47 (0.03)	0.48 (0.03)
Inspection rate (%)	8.71 (0.17)	9.28 (1.62)	8.92 (0.26)	8.59 (0.17)
Mean fines ($ thousands)	15.24 (2.25)	15.24 (2.25)	13.26 (2.64)	17.88 (2.72)
Mean violations (%)	0.48 (0.03)	0.69 (0.71)	0.55 (0.06)	0.43 (0.02)
Mean plant utility	0.000 (0.007)	0.003 (0.014)	0.001 (0.007)	−0.001 (0.007)
Mean NOX (tons)	77.8 (0.6)	83.6 (11.6)	80.3 (1.3)	76.0 (0.5)
Mean SOX (tons)	108.3 (1.0)	116.1 (17.7)	112.4 (1.8)	105.4 (0.7)
Mean air pollutants (tons)	264.0 (2.5)	287.4 (45.1)	275.0 (5.6)	256.3 (1.9)

Note: all statistics report the weighted average of the long-run equilibrium mean, weighting with the number of plants by industry and region in our data. Experiment (1) presents the results of our model given the estimated coefficients and the existing regulatory actions and outcomes. Other columns change the state-contingent fines faced by plants. Experiment (2) presents results using fines that match the long-run equilibrium baseline fines for each industry and region.

outcomes at any state will be the same. Accordingly, all our counterfactuals change either the state-contingent fine policy or the plant utility functions, but do not change inspection policies or state-contingent regulatory outcomes.

Our first set of results, in Table 6, focuses on changes in the state-contingent fine policy. In all cases, we report the long-run mean values of regulatory states, regulatory actions, regulatory outcomes, investment rates, and plant utility. We also report the mean level of pollution, using the NEI data to impute the mean pollution level by industry, region, and status as being in compliance, a regular violator or a HPV. We report these values both for our ML fixed coefficients and our GMM random coefficient estimates. Column (1) of Table 6 reports the baseline, which is calculated at the estimated parameters. We find similar results here across both specifications, with 0.95% of plants in HPV status in the long run under our ML estimates and 0.8% in HPV status under our random coefficients estimates. Long-run investment rates are also similar but higher in the equilibrium of the random coefficient estimates, occurring in 0.33% of periods instead of 0.27% of periods in the ML estimates.

Column (2) Table 6 reports the long-run outcomes that would occur if the non-linearity of fines was completely removed. Here, we force the regulatory policy to have zero fines when Ω indicates compliance and the same fines for all other violator states, conditional on industry and region. In order to isolate the impact of the non-linearity of fines rather than the level of fines, we search for the value of the fines in each industry and region that would make the equilibrium fines match those in the baseline. Under both specifications, we find that this policy increases the HPV rate substantially. Most of this increase derives from having fewer plants in compliance, not from having fewer plants in regular violator status.

Under both specifications, we find that both the inspection rate and number of violations increases, with violations increasing at a faster rate than inspections. The fact that the rate of violations increases more than inspections implies that the equilibrium distribution of e has changed to have more plants out of compliance. The increase in the inspection rate also shows that this policy would be more costly to the regulator since inspections are costly to the regulator in addition to the plant (though we do not model the level of this extra cost in our analysis). Interestingly, the investment rate does not change much from the baseline here, suggesting that it is not that plants are investing much less, but rather, that they are investing at different times.

The results of eliminating non-linear fines are much more dramatic for the GMM random coefficient estimates than for the ML estimates. For instance, in the ML estimates, the percent of plants in HPV status increases by 68%, to 1.6% of all plants. Under the GMM random coefficient estimates, the HPV percent increases almost 800%, going to 7.1% of all plants. These results show that heterogeneity in plant costs—which we find to be present in the data—is a huge reason to use dynamic enforcement. Finally, we note that plant utility is higher with equal fines under the random coefficient estimates. This means that the lower investment rate is more than compensating plants for the higher inspection rate, violation rate, and HPV status rate under this policy.

Finally, given the much higher level of plants in HPV status, we also find much higher levels of air pollution. Specifically, with the GMM random coefficient estimates, the levels of NOX, SOX, and total criteria air pollution all almost double between the baseline and column (2). Thus, non-linear fines are effective in lowering pollution, even given the same mean equilibrium level of fines.

Turning to columns (3) and (4), we consider changes in the non-linearity of fines that are less dramatic than in column (2). Specifically, column (3) considers halving the fines at every regulatory state Ω' where the plant is in HPV status while column (4) considers doubling the fines at every such state. We find results that are consistent with our findings in column (2). Specifically, column (3) reports increases in the percent of plants in HPV status while column (4) shows decreases in this probability. The results are similarly larger for the GMM random coefficient estimates than for the ML estimates. Considering fines, even though column (3) halves the state-contingent fines in HPV status, we find that, in equilibrium under the random coefficient estimates, the actual amount of fines goes up dramatically relative to the baseline.

We next consider changes to the plants' cost structure, with results in Table 7. This table again reports the baseline as column (1). Column (2) reports the case where plants bear no costs from inspections, violations, or being in HPV status. Columns (3) and (4) focus on removing or increasing the plant's cost of being in HPV status.

From column (2), eliminating the cost of inspections, violations, and HPV status increases the share of plants in HPV status, the inspection rate, and the violation rate. Thus, the cost of facing enforcement activity is also critically important for reducing pollution. If plants did not face a cost of inspections, violations, and HPV status, then the rate of plants in HPV status would hugely increase.

Turning to columns (3) and (4), eliminating the utility cost to plants for being in HPV status would result in a huge increase in plants in HPV status under the GMM random

Table 7: Counterfactual Results: Changing Plants' Cost Structure

	(1) Baseline	(2) No cost for inspections, violations, or HPV status	(3) No cost for being in HPV status	(4) Cost for HPV status doubled from baseline
Maximum likelihood estimates				
Compliance rate (%)	94.84 (0.29)	94.06 (0.40)	94.39 (0.51)	95.11 (0.30)
Regular violator rate (%)	3.77 (0.24)	3.77 (0.24)	3.77 (0.24)	3.78 (0.24)
HPV rate (%)	1.38 (0.12)	2.17 (0.28)	1.84 (0.40)	1.12 (0.13)
Investment rate (%)	0.41 (0.02)	0.40 (0.02)	0.41 (0.02)	0.41 (0.02)
Inspection rate (%)	8.78 (0.18)	9.03 (0.20)	8.92 (0.23)	8.70 (0.17)
Mean fines ($ thousands)	17.33 (3.02)	25.39 (4.17)	21.78 (4.48)	14.63 (2.89)
Mean violations (%)	0.50 (0.04)	0.59 (0.04)	0.55 (0.07)	0.47 (0.03)
Mean plant utility	−0.006 (0.004)	−0.001 (0.001)	−0.005 (0.004)	−0.007 (0.004)
Mean NOX (tons)	79.0 (0.7)	82.6 (1.3)	81.0 (1.9)	77.7 (0.7)
Mean SOX (tons)	109.8 (1.1)	116.3 (2.3)	113.4 (3.6)	107.6 (1.0)
Mean air pollutants (tons)	268.4 (2.8)	284.8 (5.8)	277.8 (8.7)	262.9 (2.7)
GMM random coefficient estimates				
Compliance rate (%)	95.36 (0.25)	94.54 (0.36)	88.19 (5.27)	94.85 (0.85)
Regular violator rate (%)	3.37 (0.20)	3.38 (0.20)	3.19 (0.24)	3.37 (0.20)
HPV rate (%)	1.27 (0.12)	2.08 (0.26)	8.62 (5.41)	1.79 (0.77)
Investment rate (%)	0.48 (0.03)	0.47 (0.03)	0.48 (0.05)	0.48 (0.03)
Inspection rate (%)	8.71 (0.17)	8.98 (0.20)	11.48 (2.02)	8.84 (0.30)
Mean fines ($ thousands)	15.24 (2.25)	25.52 (4.95)	110.61 (74.80)	16.32 (4.11)
Mean violations (%)	0.48 (0.03)	0.58 (0.05)	1.36 (0.64)	0.50 (0.06)
Mean plant utility	0.000 (0.007)	0.002 (0.001)	0.016 (0.014)	0.001 (0.007)
Mean NOX (tons)	77.8 (0.6)	81.6 (1.3)	103.0 (18.3)	78.2 (1.5)
Mean SOX (tons)	108.3 (1.0)	115.1 (2.2)	144.8 (25.5)	108.9 (2.3)
Mean air pollutants (tons)	264.0 (2.5)	281.0 (5.4)	367.8 (73.8)	267.1 (7.7)

Note: all statistics report the weighted average of the long-run equilibrium mean, weighting with the number of plants by industry and region in our data. Experiment (1) presents the results of our model given the estimated coefficients and the existing regulatory actions and outcomes. Other columns change the plants' regulatory cost parameters θ.

coefficient estimates, even larger than removing the non-linearity of fines as in column (2) of Table 6. This also results in almost doubling the level of air pollutants, relative to the baseline. Thus, a large part of the effectiveness of enforcement for the EPA may be coming from the stigma of being in HPV status. If we double the costs of HPV status, we find that the results are similar to the baseline. This result is driven in part by the fact that some plants have a positive utility from being in HPV status under the random coefficient estimates. Our results here are in line with the work of Evans (2016), who finds that being placed on a federal EPA "Watchlist" as an HPV that was not returning to compliance in a timely manner was costly to plants.

7 Conclusion

This paper empirically evaluates the effectiveness of dynamic regulation in the context of Clean Air Act and Amendments enforcement by constructing a unique dataset of repeated interactions between environmental regulators and industrial emissions sources. Using these data, we first document that environmental regulators increase both enforcement activities and fines non-linearly with a plant's regulatory state. We then build and estimate a dynamic structural model of the optimal timing for a plant as investment in environmental remediation, by using variation in regulator enforcement and plant compliance characteristics. We develop a random coefficients estimator that is computationally tractable and that allows for a wide heterogeneity in the plant costs from regulatory scrutiny.

We find that there are substantial cost to plants not only of investing in pollution abatement but also of facing regulator enforcement actions: inspections, violations, fines, and being designated as a high priority violator. We also find that there is substantial heterogeneity across plants in their regulatory compliance costs.

Counterfactual simulations suggest that the non-linearity of enforcement activity and fines over regulatory states leads to increased compliance with the Clean Air Act, and hence lower air pollution, while simultaneously decreasing the regulatory burden for state and federal EPAs and compliance costs faced by plants. A comparison of our random coefficients estimates to our estimates with fixed coefficients across plants shows that heterogeneity in plant costs is particularly important in driving dynamic enforcement to meet these goals.

Overall, this analysis provides the first empirical estimates of the how plants respond to the dynamic regulations frequently used in environmental regulation around the world. Our results show that this regulatory framework drastically decreases violations of the Clean Air Act and Amendments relative to more linear fines or a regulatory structure that imposes fewer enforcement costs on plants. Our modeling framework and results on dynamic enforcement for the CAAA may improve analysis and modeling for evaluating dynamic enforcement in a variety of other settings.

References

Arcidiacono, Peter and Robert A. Miller, "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity," *Econometrica*, November 2011, *79* (6), 1823–1867.

Becker, Gary S., "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 1968, *76* (2), 169–217.

Blondiau, Thomas, Carole M. Billiet, and Sandra Rousseau, "Comparison of criminal and administrative penalties for environmental offenses," *European Journal of Law and Economics*, February 2015, *39* (1), 11–35.

Blundell, Wesley, "When Threats Become Credible: A Natural Experiment of Environmental Enforcement from Florida," *Working Paper*, 2017.

Connault, Benjamin, "Hidden Rust Models," *Working Paper*, 2016.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan, "The Value of Regulatory Discretion: Estimates from Environmental Inspections in India?," *Working Paper*, 2016.

Earnhart, Dietrich, "Panel data analysis of regulatory factors shaping environmental performance," *The Review of Economics and Statistics*, 2004, *86* (1), 391–401.

Eckert, Heather, "Inspections, warnings, and compliance: the case of petroleum storage regulation," *Journal of Environmental Economics and Management*, 2004, *47* (2), 232–259.

Evans, Mary F., "The Clean Air Act Watch List: An Enforcement and Compliance Natural Experiment," *Journal of the Association of Environmental and Resource Economists*, 2016, *3* (3), 627–665.

Fowlie, Meredith, Mar Reguant, and Stephen P. Ryan, "Market-based emissions regulation and industry dynamics," *Journal of Political Economy*, 2016, *124* (1), 249–302.

Fox, Jeremy T, Kyoo il Kim, and Chenyu Yang, "A simple nonparametric approach to estimating the distribution of random coefficients in structural models," *Journal of Econometrics*, 2016, *195* (2), 236–254.

Fox, Jeremy T., Kyoo il Kim, Stephen P. Ryan, and Patrick Bajari, "A simple estimator for the distribution of random coefficients," *Quantitative Economics*, 2011, *2* (3), 381–418.

Friesen, Lana, "Targeting enforcement to improve compliance with environmental regulations," *Journal of Environmental Economics and Management*, July 2003, *46* (1), 72–85.

Gowrisankaran, Gautam and Marc Rysman, "Dynamics of Consumer Demand for New Durable Goods," *Journal of Political Economy*, 2012, *120* (6), 1173–1219.

Gray, Wayne B. and Jay P. Shimshack, "The effectiveness of environmental monitoring and enforcement: A review of the empirical evidence," *Review of Environmental Economics and Policy*, 2011, *5* (1), 3–24.

Harford, Jon D., "Measurement error and state-dependent pollution control enforcement," *Journal of Environmental Economics and Management*, July 1991, *21* (1), 67–81.

Harrington, Winston, "Enforcement leverage when penalties are restricted," *Journal of Public Economics*, 1988, *37* (1), 29–53.

Heyes, Anthony, "Implementing Environmental Regulation: Enforcement and Compliance," *Journal of Regulatory Economics*, 2000, *17* (2), 107–29.

_ **and Neil Rickman**, "Regulatory dealing revisiting the Harrington paradox1We are grateful to Lars Hansen, Robert Cairns, Jeff Frank, two referees from this journal and seminar participants at the University of London for helpful comments. The usual disclaimer applies.1," *Journal of Public Economics*, June 1999, *72* (3), 361–378.

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

Kang, Karam and Bernardo S. Silveira, "Understanding Disparities in Punishment: Regulator Preferences and Expertise," *Working Paper*, 2018.

Ko, Kilkon, John Mendeloff, and Wayne Gray, "The Role of Inspection Sequence in Compliance with the US Occupational Safety and Health Administration's (OSHA) Standards: Interpretations and Implications," *Regulation & Governance*, March 2010, *4*, 48–70.

Landsberger, Michael and Isaac Meilijson, "Incentive generating state dependent penalty system: The case of income tax evasion," *Journal of Public Economics*, 1982, *19* (3), 333–352.

Laplante, Benot and Paul Rilstone, "Environmental inspections and emissions of the pulp and paper industry in Quebec," *Journal of Environmental Economics and management*, 1996, *31* (1), 19–36.

Lyubich, Eva, Joseph S. Shapiro, and Reed Walker, "Regulating Mismeasured Pollution: Implications of Firm Heterogeneity for Environmental Policy," Technical Report, Yale University 2018.

Nevo, Aviv, John L. Turner, and Jonathan W. Williams, "Usage-Based Pricing and Demand for Residential Broadband," *Econometrica*, 2016, *84* (2), 411–443.

Raymond, Mark, "Enforcement leverage when penalties are restricted: a reconsideration under asymmetric information," *Journal of Public Economics*, August 1999, *73* (2), 289–295.

Rust, John, "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 1987, *55* (5), 999–1033.

Ryan, Stephen P., "The costs of environmental regulation in a concentrated industry," *Econometrica*, 2012, *80* (3), 1019–1061.

Shimshack, Jay P., "The economics of environmental monitoring and enforcement," *Annu. Rev. Resour. Econ.*, 2014, *6* (1), 339–360.

_ **and Michael B. Ward**, "Regulator reputation, enforcement, and environmental compliance," *Journal of Environmental Economics and Management*, 2005, *50* (3), 519–540.

Shinkuma, Takayoshi and Shunsuke Managi, "Effectiveness of policy against illegal disposal of waste," *Environmental Economics and Policy Studies*, April 2012, *14* (2), 123–145.

Telle, Kjetil, "Monitoring and enforcement of environmental regulations: Lessons from a natural field experiment in Norway," *Journal of Public Economics*, March 2013, *99*, 24–34.

Timmins, Christopher, "Measuring the dynamic efficiency costs of regulators' preferences: Municipal water utilities in the arid west," *Econometrica*, 2002, *70* (2), 603–629.

Train, Kenneth E., *Discrete Choice methods with simulation*, Cambridge University Press, 2009.

8 Appendix

Table 1: Regulatory Actions

	Regular violator			High priority violator		
Dependent variable:	Inspection	Violation	Fine amount (thousands)	Inspection	Violation	Fine amount (thousands)
Depreciated Accumulated Violations	0.4111***	0.3251***	4.1015***	0.2955***	0.2295***	3.9487***
	(0.0082)	(0.0110)	(0.1280)	(0.0074)	(0.0080)	(0.1457)
Inspection		0.8091***	2.1097***		0.5257***	2.7643***
		(0.0182)	(0.2183)		(0.0178)	(0.3103)
Violation			2.5193***			4.9378***
			(0.3913)			(0.4488)
Lag of Investment	0.1792***	-0.0953**	-1.9349***	0.0452***	-0.2049***	-7.5272***
	(0.0220)	(0.0406)	(0.4652)	(0.0160)	(0.0237)	(0.4412)
2nd Lag of Investment	0.3790***	-0.0177	-0.9771**	0.1883***	0.1812***	1.2285**
	(0.0212)	(0.0366)	(0.4124)	(0.0300)	(0.0360)	(0.6224)
σ			10.5901			18.8278
			(0.1380)			(0.1700)
Pseudo R^2	0.1527	0.1510	0.0844	0.0791	0.0963	0.0347

Note: table shows results from logit regressions for "inspections" and "violations" and tobit regressions for "fine amount." ***, **, and * indicates statistical significance at the 1%, 5%, and 10% levels, respectively. Regressions are run separately depending on whether the plant is a regular violator or high priority violator at the end of the previous period. All regressions also include EPA region and NAICS 2-digit industry dummies. Results for plants that are in compliance are omitted. Standard errors are clustered at the plant level.

Table 2: Transition Functions

Beginning State:	Compliance		Regular violator		High priority violator	
Transition to:	Stay in compliance	Into regular vio. given out of comp.	Into compliance	Stay in regular vio. given into comp.	Stay in HPV	Into compliance given not out of HPV
Lag of Investment			1.1979***	0.0361	-2.8379***	-0.4136
			(0.0211)	(0.0671)	(0.1173)	(0.3688)
2nd Lag of Investment			0.6865***	-0.2595***	0.3131***	-0.3112***
			(0.0224)	(0.0484)	(0.0671)	(0.0910)
Depreciated Accumulated Violations			0.2416***	-0.1576***	0.1465***	-0.3315***
			(0.0092)	(0.0151)	(0.0158)	(0.0231)
Inspection	-0.7402***	-0.5135***	-0.0255*	-0.4946***	0.1375***	-0.0646*
	(0.0086)	(0.0254)	(0.0136)	(0.0301)	(0.0347)	(0.0376)
Violation	-4.1890***	0.0536**	-2.3670***	-1.4465***	1.4263***	-1.6977***
	(0.0206)	(0.0249)	(0.1366)	(0.0346)	(0.0595)	(0.1433)
Fine	-6.63e-6***	-0.0045***	-0.0062***	-4.38e-6*	6.63e-5**	2.15e-5
	(1.14e-6)	(0.0007)	(0.0010)	(2.3e-6)	(2.69e-5)	(4.94e-5)
Pseudo R^2	0.5286	0.1751	0.1318	0.3847	0.4006	0.1346

Note: table shows results from probit regressions. All regressions also include EPA region and NAICS 2-digit industry dummies. ***, **, and * indicates statistical significance at the 1%, 5%, and 10% levels, respectively. Standard errors are clustered at the plant level.