NBER WORKING PAPER SERIES

ESCALATION OF SCRUTINY: THE GAINS FROM DYNAMIC ENFORCEMENT OF ENVIRONMENTAL REGULATIONS

Wesley Blundell
Gautam Gowrisankaran
Ashley Langer

Working Paper 24810
http://www.nber.org/papers/w24810

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2018

We thank Christian Belzil, Paul Ellickson, Wayne Gray, Alex Hollingsworth, Karam Kang, Anatolii Kokoza, Derek Lemoine, Ivan Rudik, Nick Ryan, Jay Shimshack, Kenneth Train, Kyle Wilson, Jonathan Williams, Tiemen Woutersen, Chenyu Yang, and seminar participants at numerous institutions for helpful comments. Terri Dykes and Jason Swift at EPA were incredibly helpful with helping us understand the details of the data and Clean Air Act enforcement generally. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Escalation of Scrutiny: The Gains from Dynamic Enforcement of Environmental Regulations
Wesley Blundell, Gautam Gowrisankaran, and Ashley Langer
NBER Working Paper No. 24810
July 2018
JEL No. C57,Q53,Q58

ABSTRACT

The U.S. Environmental Protection Agency uses a dynamic approach to environmental enforcement for air pollution, with repeat offenders subject to high fines and designation as high priority violators (HPV). We estimate the benefits of dynamic monitoring and enforcement by developing and estimating a dynamic model of a plant and regulator, where plants decide when to invest in pollution abatement technologies. We use a fixed grid approach to estimate random coefficient specifications. Investment, fines, and HPV designation are very costly to most plants. Eliminating dynamic enforcement would have large adverse impacts on the number of high priority violators and pollutants emitted.

Wesley Blundell
Department of Economics
College of Business and Economics
California State University, East Bay
25800 Carlos Bee Boulevard
Hayward, CA 94542
wesley.blundell@csueastbay.edu

Gautam Gowrisankaran
Department of Economics
University of Arizona
P.O. Box 210108
Tucson, AZ 85721-0108
and HEC Montreal
and also NBER
gowrisankaran@eller.arizona.edu

Ashley Langer
Department of Economics
University of Arizona
McClelland Hall 401
1130 E. Helen St.
Tucson, AZ 85721-0108
alanger@email.arizona.edu

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 when 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 abatement and regulatory compliance. These complications have led environmental regulators to create systems whereby repeat offenders are punished much more severely than one-time violators. The U.S. Environmental Protection Agency (EPA) designates persistent offenders of the Clean Air Act and Amendments (CAAA) as *high priority violators*, and targets them with elevated scrutiny and penalties. Figure 1 shows mean CAAA inspection rates, fines, and violation rates, separately for plants in compliance, regular (not-high-priority) violators, and high priority violators. In each case, the level of scrutiny increases dramatically across these statuses.[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 pollution abatement costs without excessively penalizing low-cost plants.

This paper addresses two related research questions on environmental monitoring and

[1]The increasing pattern for fines in Figure 1 could be due to dynamic enforcement or to those plants violating environmental norms more frequently or severely. Our analysis allows for both of these explanations.

Figure 1: EPA Clean Air Act and Amendment Enforcement by Regulatory Status

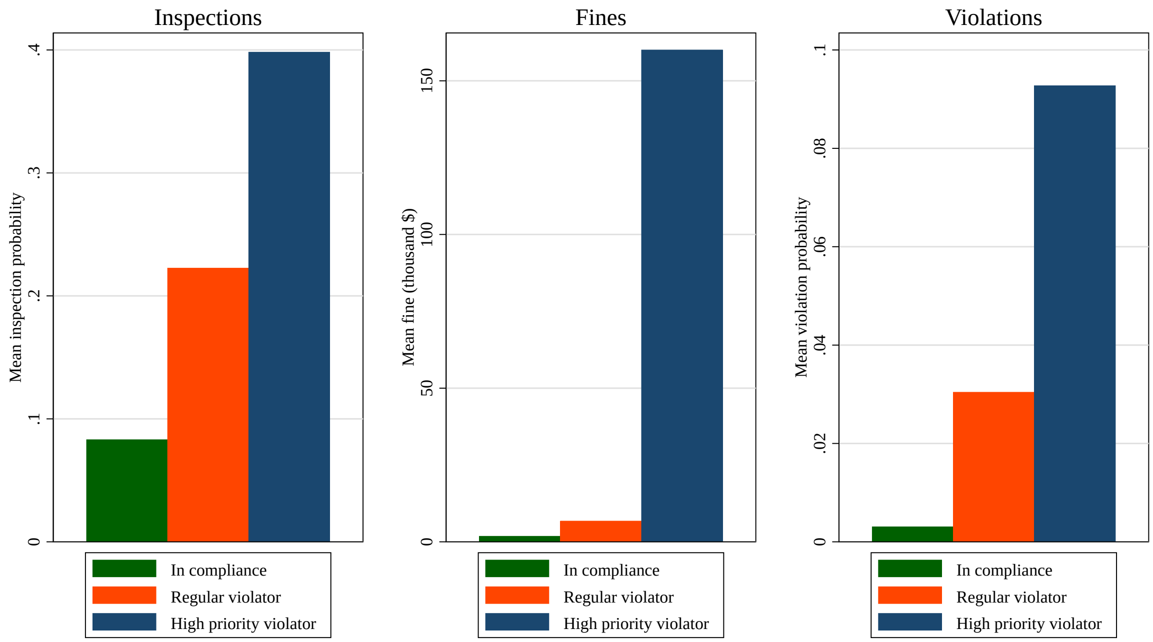

Note: authors' calculations based on estimation sample; data from 2007-13.

enforcement of air pollution laws. First, we estimate the cost to industrial facilities of complying with the current dynamic approach used by the EPA in enforcing the U.S. Clean Air Act and Amendments. Second, we simulate the value of alternative enforcement regimes in affecting plants' pollution and compliance with the CAAA.

We believe that our modeling and estimation framework may be generally useful in evaluating dynamic enforcement of regulations. Dynamic enforcement of environmental regulations is used and studied across many settings, including for the CAAA (Evans, 2016; Blundell, 2017) and the Clean Water Act (Earnhart, 2004; Shimshack and Ward, 2005) in the U.S., 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 regulations, e.g., in tax compliance and auditing (Landsberger and Meilijson, 1982) and worker health and safety regulation (Ko et al., 2010).

Despite the importance of dynamic enforcement in the real world, this literature has generally focused on documenting the responses of plants to dynamic enforcement, rather than on estimating the costs of compliance with dynamic regulations or the value of alternative enforcement regimes. One reason that it has not addressed the latter questions is that it is problematic to do so without estimating a model that formally accounts for the dynamic, optimizing behavior of plants, such as the theoretical model proposed by Harrington (1988). In particular, one would need to account for the impact of investment in decreasing future inspections, penalties, and probabilities of being classified as a high priority violator. A recent survey article (Shimshack, 2014) notes that a static model would tend to underpredict compliance rates by not accounting for future penalties from non-compliance. 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 the parameters of a discrete-time dynamic game of a plant faced with a regulator in the spirit of Harrington (1988). In our model, the regulator makes decisions regarding inspections and fines. The plant, in turn, decides whether and when to invest in pollution abatement technologies.

Each period, the regulator first decides whether or not to inspect the plant based on the plant's regulatory state and an idiosyncratic term. The regulatory state 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 violations and investments. The regulator and plant then both receive a signal about the plant's environmental compliance. The distribution from which the signal is drawn is affected by whether the plant was inspected. The signal indicates the presence of a reportable violation and whether the plant is in compliance, a regular violator, or a high priority violator. The regulator then chooses how much to fine the plant, if at all, based on the signal and the plant's regulatory state. At this stage, the plant obtains disutility from regulatory enforcement actions and designation as a high priority violator.

Following the regulator's action, if the plant is a regular or high priority violator, it makes a binary choice of whether to invest in pollution abatement. 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. Thus, a plant that is not in compliance will invest if its cost of investment is less than its expected future benefit in the form of lower future environmental scrutiny.

Our estimation makes use of extensive data that include information on virtually all industrial facilities in high polluting sectors covered by the CAAA. Our data report inspections, fines, violations, compliance status, and investment decisions for a seven-year-long panel with over 2.5 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 models. In both models, we seek to recover plants' utility functions. Although we specify a dynamic game between the plant and regulator, we do not recover the regulator's utility function but instead model regulatory actions and transitions as conditional choice probabilities (CCPs) and estimate them with simple probit and tobit regressions that account for their limited dependent nature.[2]

The models differ in their assumptions regarding the heterogeneity in the disutility parameters across plants and in their estimation approach. First, we estimate a model where the disutility parameters are the same across plants. For these specifications, we estimate the disutility parameters using a quasi-likelihood nested fixed point estimator where the dependent variable is investment. Specifically, for each candidate structural parameter, we compute the plant's dynamically optimal investment decision at each state. We then find the value of the plant's structural parameters that maximizes the probability of the observed investment actions.[3]

[2]Because we do not recover the regulator's preferences, our counterfactuals are based on plant optimization given alternative regulatory policies and do not necessarily stem from the equilibrium of the dynamic game.

[3]This forms a quasi-likelihood (and not a likelihood) because we use the regulator's estimated CCPs in the plant's dynamic optimization process.

Second, we estimate a model where the disutility parameters are heterogeneous across plants. Our heterogeneity specifications allow 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 model are the population weights of observing each of the potential disutility parameters. By choosing a finite grid, this model essentially allows us to estimate a non-parametric distribution of the cost of investments and regulatory actions and outcomes across plants (Fox et al., 2016). The generalized method of moments (GMM) estimator that we use here is computationally very tractable because the optimization problem is quick and convex.

Identification of the costs to plants of investment and enforcement is critical for understanding how enforcement affects pollution. Heuristically, to estimate these costs, one could estimate a discrete choice specification where whether a plant invests is regressed on a constant (interpreted as the cost of investment) and the decreased regulatory scrutiny that would accrue from such investment. Importantly, the actual (endogenous) future regulatory scrutiny for each plant could not be included as a regressor, as one would need to simulate how expected future regulatory scrutiny changes with investment. Our quasi-likelihood model essentially performs this regression, also allowing plants to have the option value of investment in future periods. It is therefore identified by the extents to which (a) investment lowers regulatory scrutiny among plants with observably similar characteristics and (b) plants respond to this decrease through investment. Variation in the regulatory regimes across EPA regions and industrial sectors provides substantial identifying variation here. Our second model recovers the heterogeneity in utility parameters using the equilibrium distribution of plant states and correlations in investment over time for plants. The panel nature of our data is essential here.

Summary of results. The estimates from our quasi-likelihood model where we constrain the disutility to be the same across plants confirm the findings in Duflo et al. (2018) that plants have a substantial disutility from environmental enforcement. We find that fines and classification as a high priority violator both carry substantial costs. In particular, the mean cost to a plant from being a high priority violator is equivalent to a fine of $430,000 per quarter. This cost may occur because being a high priority violator may affect the plant's reputation and relationship with the surrounding community. We also find that investment is very costly to plants, with a new investment being equivalent to a \$26 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.

The estimates from our random coefficients model show that there is substantial heterogeneity across plants in their disutility from environmental enforcement and investment costs, with a mass of 42% of plants on the quasi-likelihood estimates. However, the remaining mass of plants mostly has a substantially lower cost of investment and higher cost of fines. For instance, the majority of plants have an investment cost equivalent to a \$7 million or smaller fine. The substantial heterogeneity of compliance costs across plants may further increase the value of dynamic enforcement.

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) non-linearity of penalties across regulatory states, (2) 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 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 both more pollution and more plants in HPV status. For instance, using our random coefficients estimates, mean criteria air pollutants would rise 58% under this policy. Similarly, we find that if we eliminated the (stigma) cost of HPV status, criteria air pollutants would rise 25%.

Relation to literature. This paper relates 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 a literature that structurally estimates firm behavior in the presence of energy and environmental regulatory policies (Timmins, 2002; Ryan, 2012; Lim and Yurukoglu, 2015; Fowlie et al., 2016; Duflo et al., 2018; Houde, 2018; Kang and Silveira, 2018). Two recent papers specifically consider environmental regulation and compliance. Kang and Silveira (2018) estimate a game between the regulator and municipal water treatment plants in California in order to understand the value of regulatory discretion. Given the specifics of their sector, they focus on static regulator incentives and compliance stemming from continuous effort, rather than from discrete investments, as in our case. Duflo et al. (2018) also seek to understand the value of regulatory discretion, by estimating a dynamic model of environmental regulatory enforcement for plants in India. The dynamics in their context stem largely from the fact that plants must invest, rather than from non-linear enforcement. Though our ultimate research question is somewhat different, our conception of investment is similar to theirs. Our identification builds on theirs in that we observe multiple regulatory regimes—based on EPA regions and industrial sectors—and that we model random coefficients.

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), Fox et al. (2016), and 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 by developing moments that are useful in identifying the presence of heterogeneity and by obtaining asymptotically efficient parameter estimates.

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

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 Act and other environmental legislation. The Act was again amended in 1977 and 1990. Our data are from after 1990 and hence we consider the Clean Air Act and Amendments (CAAA).

Many of the decisions of how to structure this enforcement were 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.[4]

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. We use EPA regions for identifying variation because they represent a unit of analysis that captures both the interpretation of federal policy and geographic preferences for enforcement.

[4]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 EPAs for brevity. State and regional EPAs are required to maintain a minimum level of enforcement, but can exceed this threshold (Shimshack, 2014).

Under the current 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 out of compliance.[5] 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 all 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.

If a plant is substantially or persistently out of compliance, it 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 than "regular" violators, as in Figure 1. HPV designation also highlights violations at the plant to the federal EPA, who may become involved in returning the plant to compliance. In this way, the HPV designation is used by the federal EPA to allocate its limited resources to regulating the highest-impact plants, whether those are plants that are emitting large amounts of pollution or those whose behavior, while not particularly polluting, would set bad precedents for CAAA enforcement going forward. 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.

During the time frame of our analysis, the federal EPA further used a "watch list" to focus particular attention on HPVs that did not resolve all of their violations in a timely

[5]The EPA has programs other than the CAAA to reducing pollutants, such as the NO_X cap-and-trade program. Because the CAAA requires using the best available pollution control technologies, plants cannot simply purchase cap-and-trade permits to ensure CAAA compliance.

manner. The watch list was originally intended to increase oversight by the EPA, but was also reported in the press and hence could have increased costs by leading to increased attention from local politicians and civilian environmental protection groups (Evans, 2016).

Finally, the HPV system has been changed over time. In 2014 (after our sample period), the guidelines for plants being classified as HPVs were narrowed and the watch list was eliminated. These changes highlight the fact that the EPA is still working to determine the optimal enforcement policies, which makes evaluating the effect of dynamic incentives particularly important.

2.2 Implications of Current Regulatory Structure

Before turning to our data, it is worth understanding why dynamic regulatory enforcement may reduce pollution and add value. In a world where the plant and regulator can both perfectly observe non-compliance, the plant can instantly correct the non-compliance 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, the plant 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 the EPA, severely punishing plants when violations are uncovered may not be useful. Instead, the EPA issues only minor penalties for an initial, non-critical violation and then increases penalties and oversight as additional violations accumulate. This approach allows for the possibility that the EPA's inspection process may be noisy, and that apparent violations may not actually be present or may not be severe. Additionally, it 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. Dynamic enforcement may help achieve compliance without imposing large penalties because of the threat of high penalties from non-compliance (Harrington, 1988).

Finally, heterogeneous investment costs across plants may further add value to dynamic enforcement. If plants have heterogeneous investment costs, a dynamic enforcement approach may incentivize low-cost plants to invest in pollution abatement when they are regular violators while high-cost plants will wait until they become HPVs. This would then 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 explain our use of both databases, and provide additional details in On-Line Appendix A1.

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.[6] 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. Our analyses control for sector with the first two-digits of the NAICS code and for EPA region (obtained from the address). Our analysis keeps seven sectors with high pollution levels that we believe to have plants of broadly comparable sizes of plants: the three manufacturing sectors, mining and extraction, transportation, educational services (which includes school buses), and utilities.

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 Q4:2006 through the Q4:2014.[7] We use this dataset to create our base list of inspections, violations, fines, and investments. 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. Each record in this dataset details an action, such as an inspection, a notice of violation, a fine, or the review of an investment in pollution abatement. The unit of observation is the AFS ID, which indicates a polluting source. Each record lists a calendar date and provides information on the related EPA program,[8] 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 as having invested if we observe either a code indicating the resolution of an environmental issue or the issuance of a *Prevention of Significant Deterioration* (PSD) permit.[9]

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

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

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

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

[9]We provide additional details of exactly how we use the *Actions* dataset in Online Appendix A.1.

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.[10] 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 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. We also use this dataset to determine whether a plant has shut down, dropping plants from the sample once they have exited.

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 abatement 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 use as our quarterly fine amount the maximum amount across the two datasets. The plant identifier in this dataset is the FRS number (which indicates a plant) and not the AFS ID (which indicates a polluting source). The EPA provides a crosswalk from AFS IDs to FRS numbers. 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

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

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 dataset 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 regulatory compliance and hence do not count them as investments in our estimation.

Table 1 provides summary statistics on our main analysis data. They contain 2,555,952 plant-quarter observations covering 115,862 unique plants, of which 67.5 percent are present in every quarter of our sample period. As is well-documented in the literature (e.g., Evans, 2016), compliance is high: over our entire time frame, 95.7 percent of observations indicate compliance. Compliance is also high when considering individual plants: 88.5 percent of plants are never out of compliance, while 7.5 percent of plants have at least one quarter in which they have a violation but are never in HPV status, and only 4.0 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 (8.3%) than are plants in regular violator status (22.3%) and plants in HPV status (39.8%). Sim-

Table 1: Summary Statistics on Estimation Sample

Status:	Compliance	Regular violator	High priority violator
Regulator actions:			
Inspection (%)	8.31	22.26	39.82
Fine amount (thousands of $)	1.81	6.75	160.04
	(65.36)	(104.95)	(668.90)
$\mathbb{1}\{\text{Fine} > 0\}$ (%)	0.18	2.89	13.73
Regulatory outcomes:			
Violation (%)	0.31	3.04	9.27
Entrance into HPV status (%)	0.13	1.51	0.00
Plant actions:			
Investment (%)	0.00	4.62	17.39
Investment (from resolution code) (%)	0.00	4.35	16.22
Investment (from PSD permit) (%)	0.00	0.31	0.42
Investment (from HPV exit) (%)	0.00	0.00	0.83
Dropped investment in compliance (%)	0.35	0.00	0.00
Plant-quarter observations	2,445,020	73,493	37,439

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

ilarly, 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 4.6% of quarters when a plant is a violator and in 17.4% 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 Prevention of Significant Deterioration permits and from exiting high priority violation status. Finally, we observe codes that are indicative of investment in 0.35% 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 pertain to our main analysis data sample period. Because of the infrequency of the data and the fact that the NEI does not link perfectly with the plants in our base analysis data, we do not use these data in our structural estimation. However, we do use them to calculate the mean levels of eight different pollutants by regulatory state, which we then use to evaluate the likely pollution levels generated under counterfactual enforcement policies.[11]

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 an exact match on plant name, street address, zip code, 6-digit NAICS industrial sector, city, and state. Third, we merge plants based on an exact match on subsets of the above variables (plant name, street address, and city first, then plant name, sector, and city, then plant name and zip code). Finally, we perform a fuzzy match, with a 90% threshold, based on plant name and street address, conditional on perfectly matching on city, state, and the first four digits of the NAICS code.[12] Our match rate of 44% is similar to other papers that use the NEI data, and our match rate for manufacturing (71%) is similar to the the 77.4% match rate that Shapiro and Walker (2015) achieve between the NEI and the Census of Manufacturing.

Table 2: Summary Statistics on Mean Criteria Air Pollution Levels

Industrial sector	Observations in analysis data	Mean level in compliance	Mean level as regular violator	Mean level as HPV
Mining & extraction	758,792	138.0	383.8	1,117.0
Manufacturing: wood, petroleum, pharma	679,137	289.3	782.5	2,483.2
Manufacturing: metal	568,682	101.7	176.4	1,240.0
Transportation	166,202	190.7	202.5	207.5
Manufacturing: food, textiles	147,433	117.3	393.1	338.7
Educational services	147,161	67.3	169.8	186.3
Utilities	120,536	1,885.0	5,242.3	12,546.5

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.

Table 2 provides summary statistics on the reported criteria air pollution levels for our

[11] When we report counterfactual results, we also use estimates of the cost of NO_X, $PM_{2.5}$, SO_2, and VOCs for land emissions in the European Union as calculated using the median value of a life year and excluding health sensitivity effects in Holland et al. (2005).

[12] We complete the fuzzy match using the "reclink" command in Stata.

analysis data, by industrial sector.[13] There is substantial variation in the pollution levels across sectors. The most (least) polluting sector in our data in compliance is utilities (educational services). Across sectors, plants in violator status emit more pollution that plants in compliance. For most sectors, this effect is particularly pronounced for plants in HPV status.

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 38 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.[14] 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.

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

[14] 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.

Table 3: Investment and Resolution of Violations

Dependent variable: return to compliance		
Current investment	-0.114^{***}	(0.002)
One quarter lag of investment	0.383^{***}	(0.006)
Two quarters lag of investment	0.085^{***}	(0.007)
Three quarters lag of investment	-0.012^{**}	(0.005)
Four quarters lag of investment	-0.050^{***}	(0.005)
Number of observations	110,932	

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 25% chance that it will finish the period in HPV status, but there is now a 49% 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.

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.

Figure 3 displays the relationship between depreciated accumulated violations (using a

Figure 2: Effect of Investment on Regulatory State

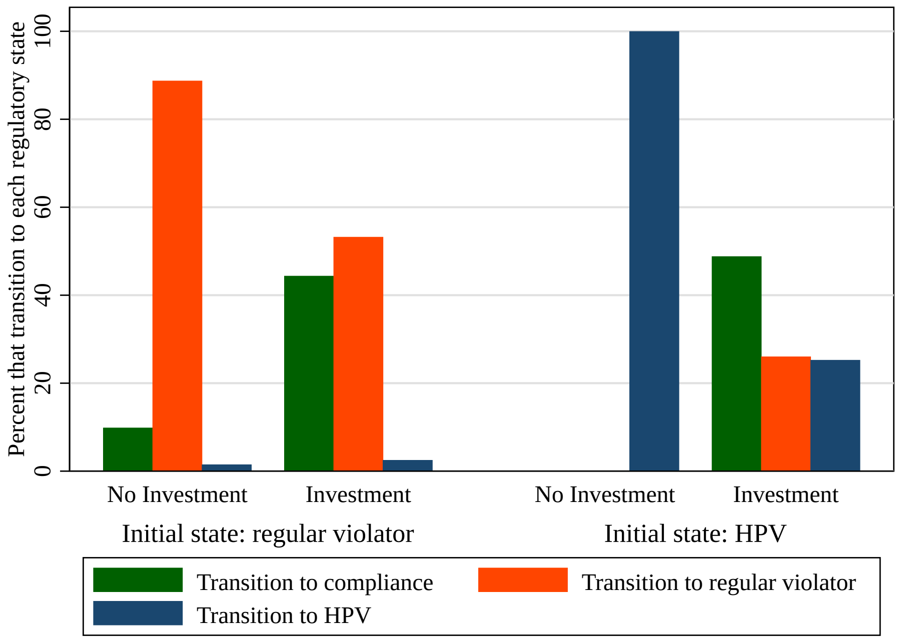

Note: authors' calculations based on estimation sample. Initial state and investment are from previous quarter.

10% quarterly depreciation rate) 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 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 depreciation level is best. Table 4 regresses inspections, fines, and violations on depreciated accumulated violations using three different depreciation rates, along with lagged HPV status. We find that the measure with the 10% depreciation rate is much more predictive of inspections and fines than the measure with a 20% depreciation rate or with a 0% depreciation rate. These results motivate our choice of including depreciated accumulated violations with a 10% quarterly depreciation rate as a state variable.

Figure 3: Depreciated Accumulated Violations and Monitoring and Enforcement

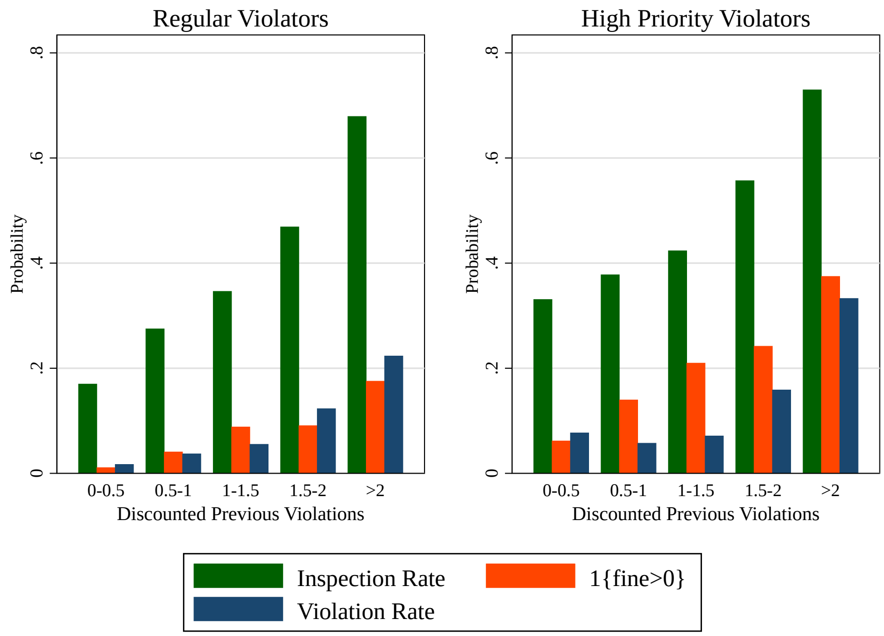

Note: authors' calculations based on estimation sample.

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^{***}	-0.008
	(0.006)	(0.004)	(0.011)
Accumulated violations with 10% depreciation	0.119^{***}	0.126^{***}	0.022
	(0.025)	(0.015)	(0.062)
Accumulated violations with 20% depreciation	-0.020	-0.059^{***}	-0.048
	(0.022)	(0.013)	(0.069)
Lagged HPV status	0.125^{***}	0.033^{***}	0.143^{***}
	(0.006)	(0.002)	(0.037)
Number of observations	110,932	110,932	110,932

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.

Heterogeneity in Regulatory Environment and Costs

Identification of our model will be aided by having heterogeneity in the regulatory environment. Figure A1 in On-Line Appendix A3 provides a scatter plot by EPA region of the cost

Figure 4: Further Investments by a Plant Following Initial Investment

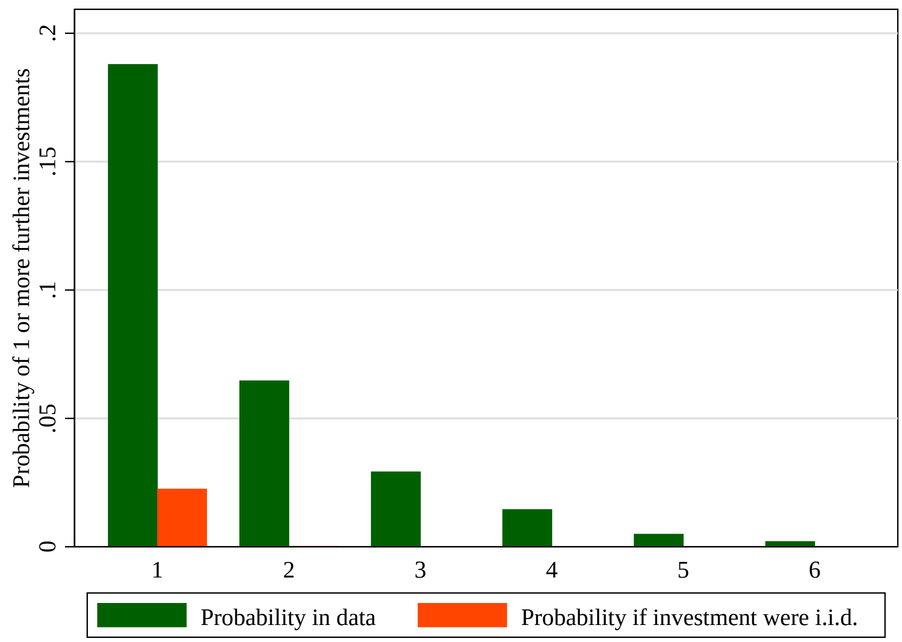

Note: authors' calculations based on estimation sample. The *i.i.d.* model is a hypothetical using the sample mean investment probability.

of being in HPV status relative to violator status in terms of inspections (the ratio between the inspection probability in HPV status to regular violator status) against the same measure in terms of fines, while Figure A2 provides a scatterplot of the analogous ratios by industrial sector. We find substantial variation and little correlation in both of these measures, with correlations of -0.23 (p=0.52) and 0.16 (p=0.74) respectively. This shows that there is variation across regulatory environments that can help identify the structural parameters and also motivate our inclusion of EPA region and industrial sector as state variables.

Finally, we evaluate the extent of heterogeneity in the correlation of investments across plants, as this heterogeneity may reflect heterogeneous investment costs. Figure 4 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 serial correlation in investment than we would expect to occur randomly. For instance, about 30% of investments are followed by at least one additional investment within the next six quarters, relative to the approximately 2.3% we would observe if investment were *i.i.d.* This suggests that a random coefficients model may be important.

4 Model

4.1 Overview

We develop a framework to estimate plant utility in the presence of dynamic regulatory monitoring and enforcement. We model each plant as playing a dynamic game with the regulator. Each period t corresponds to a quarter and the future is discounted with factor β.

Each period, the regulator first chooses whether or not to inspect. It 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. The signal also indicates the presence of a violation and whether the plant is in compliance, a violator, or a HPV. Following this, the plant, if not in compliance, makes a binary decision of whether or not to invest in pollution abatement. A plant that invests incurs a cost from its investment, but increases the chance that it returns to compliance in future periods.

We capture plant exit through the discount factor, with a lower discount factor corresponding to more exit. However, we do not endogenize the exit decision because we believe such exit is quantitatively less important than investment.[15]

We assume that the data for each plant reflect a Markov Perfect Equilibrium of the dynamic game between it and the regulator. Given the many unobserved constraints faced by the regulator, we do not estimate the regulator's utility function. Rather, we specify the regulator's policy function as a conditional choice probability (CCP) (Aguirregabiria and Mira, 2007), and then use the regulator's CCP to estimate the plants' utility functions.[16]

[15] Duflo et al. (2018) find no difference in exit rates for plants randomized into additional regulatory scrutiny in India; we believe that plants in our sample are even less likely to be at the margin for exit.

[16] Our estimation is also consistent with a plant playing against a "regulatory machine," as in Duflo et al. (2018).

4.2 Regulator and Plant Decisions

We define the *regulatory state* to be the payoff-relevant state variables over which CCPs may depend. The regulatory state in our model has five components: (1) depreciated accumulated violations with a 10% quarterly depreciation 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 discussed in Section 3.2, we believe that these summary measures capture the most important features that underlie regulatory actions.

Denote the plant's regulatory state at the start of period t to be Ω_t; Ω_t is known to the regulator and plant at the start of the period. The regulator first receives an *i.i.d.* private information shock to the value of an inspection and then decides whether or not to inspect the plant. Let $\mathcal{I}(\Omega)$ denote the CCP of an inspection in state Ω and $Ins(\Omega) \sim Binom(\mathcal{I}(\Omega))$ denote the realization. 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, in part based on the inspection. This signal has three effects. First, the state and signal together reveal whether there is a violation, with the function $Vio(\Omega, e)$. Second, also based on Ω and e, the regulator decides how much to fine the plant, if at all. Let $Fine(\Omega, e)$ denote the regulator's fine policy. Third, the state and signal also reveal the state that the plant will face at the point when it takes its action, which we denote Ω'. Let $T(\Omega, e) = \Omega'$ denote this transition. T indicates the transition into designation as being in compliance, a violator, or a HPV, and also indicates the result of $Vio(\Omega, e)$ as this will update the depreciated sum of past accumulated violations. Both Vio and T reflect environmental norms and are not choice variables of the regulator. The state Ω' is the state that is recorded in our data.

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 the regulatory state Ω_t, inspection decision Ins_t, and regulator CCPs $\mathcal{I}$.*

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 are allowed to affect the compliance signal. Assumption 1 also implies that e does not affect the future state directly, but only through its effect on violations, Vio, and state transitions, T.

Assumption 1 allows for the distribution of e_t to depend on the inspection CCP. Thus, the findings of an inspection are allowed to be a function of how often the plant is inspected given its regulatory state. However, we do not let e_t depend on the fine policy. This is useful to obtain counterfactual outcomes because it implies that the state-contingent probabilities of violations and transitions will remain the same even with counterfactual fine policies. Note that this in no way precludes the possibility that plants will invest differently with different fine structures, and that the different investments will yield different steady state distributions of regulatory states under plant optimizing behavior.

We now turn to the plant's utility function from regulatory actions. The plant obtains (dis)utility from four regulatory actions: inspections, violations, fines, and HPV status. Let $HPV(T)$ be an indicator for T designating HPV status. The flow utility for the plant from regulatory actions is:

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

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

Once the pollution signal is revealed and regulatory actions are complete, the plant moves. Recall that the state at at this point is Ω'. In any period in which Ω' indicates that the plant is a regular 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 can improve the plant's regulatory status in periods $t+1$ and $t+2$. Specifically, it enters into the state variable for the future regulatory enforcement and transition CCPs.

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 our random coefficient specifications, θ will vary across plants. However, as is common in the regulatory literature (Kang and Silveira, 2018), we do not allow the regulator to condition its monitoring and enforcement policies on each plant's individual values of these parameters. We generally expect these parameters to be negative.

4.3 Plant 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.

Let $V(\Omega)$ denote the value function at the beginning of the period, $V'(\Omega')$ denote the value function at the point right after the regulator has moved but before the plant receives its draws of ε, and $Com(T)$ be an indicator for T designating compliance.[17] We start by expositing $V'(\Omega')$:

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

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.

[17] For ease of notation, we are conditioning on the plant's parameter vector θ.

The transition state, though still not stochastic, is now potentially different than the current state, because the depreciated sum of past accumulated violations is updated.

We now exposit $V(\Omega)$, the value function at the beginning of the period:

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

where $dP(e|Ins(\Omega))$ 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 regulator CCPs.

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:

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

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

5 Empirical Implementation

Our data include a panel of plants i observed over (quarterly) time periods t and are at the level of the plant-quarter. For each plant-quarter, 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}. We first estimate the regulator's CCPs and use this to estimate plants' structural parameter vector, which is the (dis)utility to plants from regulatory actions and investment in pollution abatement. We fix $\beta = 0.95^{1/4}$ per quarter.

Starting with the regulator CCPs, we estimate probit regressions of the probability of an inspection at any state. We estimate similar probit regressions for violations based on the state and whether an inspection occurred, and tobit regressions for fines that condition on state, inspection, and violation. Finally, we estimate probits for $T(\Omega, e)$, the transition probabilities from Ω to Ω', that condition on state, inspection, violation, and fines.

To control for the state, our regressions are performed separately for plants in compliance, regular violators, and HPVs (based on their lagged reported status), and include as regressors two lags of investment, depreciated accumulated violations, and indicators for the plant's region and industrial sector. Tables A1 and A2 in On-Line Appendix A3 present estimates for these regulatory CCPs.

We estimate two models, one with homogeneous coefficients and the other with random coefficients across plants. We explain each in turn, with further details in On-Line Appendix A2, and then discuss identification.

5.1 Estimation of Model with Homogeneous Coefficients

We estimate our model with homogeneous coefficients using a quasi-likelihood nested fixed point 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 quasi-likelihood of a parameter vector θ is:

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

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. This estimator is similar to Rust (1987), except that we the regulator's CCPs are inputs, so that our estimator that is consistent with a dynamic game.[18] We obtain inference for our parameters and counterfactuals by bootstrapping our entire estimation process including the regulator's CCPs, with resampling at the plant level.

[18]We could also estimate the plant's utility function with a CCP estimator (Aguirregabiria and Mira, 2007), which is quicker to compute, but we did not, since the computational time for the nested fixed point quasi-likelihood estimator is not excessive.

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 $(\theta_1, \ldots, \theta_J)$ and that each $\theta_j,\ j = 1, \ldots, J$ occurs with probability η_j. Each plant receives a single, independent draw of θ from the multinomial distribution of potential values. The structural parameters to be estimated are therefore $\eta \equiv (\eta_1, \ldots, \eta_J)$ and no longer $(\theta_1, \ldots, \theta_J)$. We impose no restriction on the structural parameters other than what is necessary based on the fact that they are population probabilities:

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

Econometrically, the values of $(\theta_1, \ldots, \theta_J)$ are taken as given.[19] 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, 2016) 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.

Our GMM estimator has the form $\eta^* = \arg\min_\eta \|G(\eta)\| = G'(\eta)WG(\eta)$, where $G(\eta)$ is a $K \times 1$ vector of moments, G' is the transpose of G, and W is a weighting matrix. 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

[19] Fox et al. (2016) provide asymptotic results where J increases with the sample size.

values, we then estimate η by minimizing $\|G_k(\eta)\|$ subject only to the constraints in (6). This estimator is convex (Fox et al., 2016). We perform a two-step process to approximate the asymptotically efficient weighting matrix W.

Because we do not see plants from their inception onwards, we need to make some assumption about the likelihood of seeing each plant in any state. First, define a division of 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.

Our specific moments use Assumption 2 in order to be computable. Our first set of moments indicates the probabilities of being at a particular time-varying state in equilibrium, conditional on EPA region and industrial sector. Our second set of moments indicates the probabilities of being at a particular time-varying state in equilibrium times the investment probability at this state. These moments all follow closely from Nevo et al. (2016). Our third set of moments explicitly uses our panel data: it multiplies each of the the second set of moments by the sum of investments in the following six periods, as in Figure 4.

As in Nevo et al. (2016), we obtain inference for our parameters and counterfactuals by bootstrapping, with resampling at the plant level. On-Line Appendix A2 provides detail on our parameter grid, moments, and weighting matrix computation.

5.3 Identification

As discussed above, we estimate regulatory actions and the structural parameters underlying the plant's utility function. The regulatory actions and outcomes are CCPs, 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. 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 by EPA region and industrial sector and also a variety of interactions, in order to accurately capture plants' beliefs.

To understand identification of the plant utility parameters θ for our model with homogeneous coefficients, 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 regulator's CCPs (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. Identification of these two parameters therefore requires variation across states in the gross value of investment. Our data will provide this variation both for different states within an EPA region and industrial sector, and across regions and sectors.

Conditional on having identified the ratio of the two parameters, we must also identify 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 will be small, because the variance of ε is large.

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 in the expected discounted future values of each of the four regulatory levels between investing and not investing. 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 A1 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 since we must identify the distribution of values of θ rather than just the mean values of these parameters. We identify this distribution in part from the equilibrium distribution of Ω'^1 and of investment interacted with Ω'^1. For instance, if the equilibrium distribution of states is quite diffuse, this would be consistent with more variance in investment costs, all else equal. We also directly use the panel nature of our data with our third set of moments, which captures correlations over time in investments for a plant. As in Figure 4, the greater the correlation here, the more that we would expect coefficients to be heterogeneous.

6 Results

6.1 Model Estimates

As discussed in Section 5, we estimate two models to recover plants' utility functions: a model with the same parameters across plants (estimated with quasi-likelihood) and a random coefficients model (estimated with GMM). Table 5 provides results for both model.

Table 5: Estimates of Plants' Structural Parameters

	Quasi-likelihood estimates	GMM random coefficient estimates (1)	(2)	(3)	(4)	(5)
Investment utility (θ^X)	-2.948^{***}	-2.948	-1.402	-2.185	-0.564	0.545
	(0.040)					
Inspection utility (θ^I)	-0.018	-0.018	0.476	0.423	-0.739	0.511
	(0.049)					
Violation utility (θ^V)	-0.303	-0.303	-0.098	1.031	2.125	-1.784
	(0.238)					
Fine utility (millions \$, θ^F)	-0.114^{***}	-0.114	-0.208	-0.315	-0.190	-0.025
	(0.028)					
HPV status utility (θ^H)	-0.049^{**}	-0.049	-0.218	-0.144	0.032	-0.107
	(0.020)					
Weight on parameter vector	1	0.419	0.332	0.180	0.035	0.029

Note: standard errors for quasi-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 $\eta_j > 10^{-3}$.

We start with the quasi-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, with statistically significant effects for investments, fines, and HPV status. This is directly in line with the results in Duflo et al. (2018) where the authors find that both regulation and investment in pollution abatement are costly to plants. Investments, in particular, are very costly to plants: the current utility cost of an investment is equivalent to about \$26 million (2.948/0.114) 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.

HPV status is also very costly to plants, with a quarter in HPV status being equivalent to a \$430,000 fine. Note that this is in addition to the impact of HPV status on fines, inspections, and violations, which are also costly to plants. 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. Finally, though not statistically significant, a violation brings a high disutility to the plant at the mean parameter estimate, with a cost equivalent of a $2.6 million fine.

We next turn to the GMM 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 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 bootstrapped standard errors for our counterfactuals below.

Our GMM specification finds that 99% of plants can be represented by one of five parameter values.[20] We find further that there is substantial heterogeneity in the value of θ across plants. The quasi-likelihood estimates themselves—which are given in column (1)—are the estimated parameter values for 42% of plants. The next three columns, representing 55% of plants, all have smaller investment costs and larger fine coefficients than in column (1). This results in much smaller investment costs relative to fines. For instance, the investment cost in column (2) is about a half of the cost in column (1) but is equivalent to a fine cost of only $6.77 million while the investment cost in column (3) is about three quarters of the cost in column (1) but is equivalent to a fine cost of only $6.9 million. Thus, these estimates imply that investment is very costly for the substantial minority of plants given in column (1), but less costly for most of the remaining ones.

While the coefficients on investment in columns (2)-(5) are smaller on average than in the

[20]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.

quasi-likelihood results, the coefficients on HPV status are on average larger. In particular, while the quasi-likelihood coefficient on HPV status is -0.114, columns (2), (3) and (5) all have coefficients that are negative and much larger in magnitude. However, given the larger coefficients on fines for these types, the quarterly equivalent cost of HPV status are more mixed, e.g., they are \$1.05 million in column (2) but only \$457,000 in column (3). Plants with these high costs from HPV status are also more likely to obtain positive utility from inspections or violations, which is consistent with the fact that inspections may be needed to clear (costly) HPV status and violations that do not result in HPV status may also help indicate problems that can be resolved.[21]

Finally, to evaluate fit of our models, Figure A3 in On-Line Appendix A3 considers the predictions for both models as to the total number of additional investments in the six quarters after each investment, also repeating the evidence in the data from Figure 4 on this point. The random coefficients model is much closer to fitting the actual data than is the quasi-likelihood model, which underpredicts the total of repeated investments.

6.2 Counterfactuals

Using the coefficient estimates from Table 5, we now model how EPA enforcement activities, plant investments, overall compliance, and air pollution would change if the EPA changed policies or if plants' preferences were different. From Assumption 1, the state-contingent environmental compliance signal is a function only of the inspection decision and the regulator's inspection policy. This means that with the same inspection policy, the environmental signal e at any state will be the same, further implying the same findings of violations and transitions between compliance, regular violator and HPV statuses. Accordingly, all our counterfactuals change either the state-contingent fine policy or the plant utility functions, but do not change inspection policies. Note also that while we estimate parameters from the Markov Perfect Equilibrium of the dynamic game between the regulator and plant, our counterfactuals do not solve for an equilibrium of a different dynamic game but instead, for

[21] They are also in keeping with Duflo et al. (2018), who find that inspections can be beneficial to plants.

Table 6: Counterfactual Results: Changing Fine Structure

	(1) Baseline	(2) Same fines for all regular and high priority violators	(3) Fines for HPVs halved relative to baseline	(4) Fines for HPVs doubled relative to baseline
Quasi-likelihood estimates				
Compliance (%)	94.45 (0.14)	93.62 (0.24)	93.99 (0.15)	94.86 (0.13)
Regular violator (%)	4.10 (0.12)	4.02 (0.12)	4.09 (0.12)	4.10 (0.12)
HPV (%)	1.45 (0.04)	2.36 (0.22)	1.92 (0.09)	1.03 (0.04)
Investment rate (%)	0.44 (0.01)	0.44 (0.01)	0.44 (0.01)	0.45 (0.01)
Inspection rate (%)	9.25 (0.05)	9.59 (0.13)	9.40 (0.06)	9.12 (0.05)
Mean fines (1000$)	17.87 (3.06)	17.87 (3.06)	14.46 (2.61)	21.00 (3.66)
Mean violations (%)	0.53 (0.01)	0.67 (0.10)	0.59 (0.03)	0.48 (0.01)
Mean plant utility	−0.005 (0.004)	−0.004 (0.004)	−0.004 (0.004)	−0.006 (0.004)
Mean CAP (tons)	297.7 (1.8)	314.6 (4.9)	307.5 (2.8)	288.7 (1.4)
GMM random coefficient estimates				
Compliance (%)	95.07 (0.07)	87.14 (3.24)	94.16 (0.20)	95.40 (0.10)
Regular violator (%)	3.59 (0.07)	3.31 (0.11)	3.58 (0.07)	3.58 (0.07)
HPV (%)	1.34 (0.04)	9.55 (3.32)	2.26 (0.20)	1.01 (0.07)
Investment rate (%)	0.54 (0.02)	0.53 (0.03)	0.54 (0.02)	0.54 (0.02)
Inspection rate (%)	9.19 (0.05)	12.76 (1.49)	9.55 (0.11)	9.09 (0.05)
Mean fines (1000$)	16.23 (2.08)	16.23 (2.08)	16.74 (2.52)	23.50 (6.09)
Mean violations (%)	0.51 (0.01)	2.03 (0.70)	0.66 (0.04)	0.47 (0.01)
Mean plant utility	0.021 (0.007)	0.039 (0.011)	0.022 (0.007)	0.022 (0.007)
Mean NO_X (tons)	86.1 (0.4)	129.1 (16.8)	92.9 (1.8)	84.0 (0.3)
Mean SO_X (tons)	100.3 (0.8)	164.3 (24.9)	110.4 (2.5)	96.8 (0.7)
Mean lead (pounds)	13.4 (0.1)	25.6 (5.3)	15.3 (0.5)	12.9 (0.1)
Mean PM (tons)	31.6 (0.1)	41.4 (3.9)	32.9 (0.3)	31.2 (0.1)
Mean VOC (tons)	21.1 (0.0)	29.1 (3.2)	22.2 (0.3)	20.8 (0.0)
Mean HAP (tons)	1.5 (0.0)	2.2 (0.3)	1.6 (0.0)	1.5 (0.0)
Mean CAP (tons)	294.4 (1.6)	463.8 (67.2)	321.4 (7.0)	285.5 (1.3)
Mean pollution cost (1000$)	2,578.7 (11.5)	3,751.4 (454.9)	2,755.5 (44.5)	2,519.9 (10.2)

Note: all statistics report the weighted average of the long-run equilibrium mean, weighting with the number of plants by industrial sector 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 sector and region. Bootstrapped standard errors are in parentheses. Footnote 22 details pollutants and social costs.

dynamic optimization by the plant when faced with a different regulatory policy.

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 levels of different pollutants and social costs of pollution, using the NEI data to impute the mean pollution level by industrial sector, region, and status as being in compliance, a regular violator or a HPV.[22] We report values for both our quasi-likelihood and random coefficient models.[23]

Column (1) of Table 6 reports the baseline, which is calculated at the estimated parameters. We find similar results here across both specifications, with 1.45% of plants in HPV status in the long run under our homogeneous coefficients estimates and 1.34% 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 0.54% of periods instead of 0.44% of periods with the quasi-likelihood estimates.

Column (2) of Table 6 reports the long-run outcomes that would occur if the non-linearity of fines were completely removed. Here, we force the regulatory policy to have zero fines when Ω indicates compliance and the same per-quarter fines for all other violator states, conditional on industrial sector 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 sector and region that would make the equilibrium fines imposed within each industrial sector and region match those in the baseline. Under both specifications, we find that this policy increases the HPV rate substantially. For both models, most of this increase derives from having fewer plants in compliance, not from having fewer plants in regular violator status.

Focusing on the GMM estimates, the impact of eliminating non-linear fines is particularly important. We find that the percent of plants in HPV status would increase from 1.3% to 9.6%. These results emphasize that dynamic enforcement is particularly important in the presence of heterogeneous plant costs. We also note that this dramatic increase in the HPV rate is matched with a increase in regulator workload from a higher inspection rate and violation rate. However, we observe only a small decrease in the investment rate, which suggests that plants are investing at different points in time. Not shown in the table, fines would rise to $137,000 per quarter for violator plants to effectuate this policy, a level that might be too high to be politically feasible for the EPA to even implement.

[22]We report nitrogen oxides (NO_X), sulfur oxides (SO_X), lead, particulate matter (PM), volatile organic compounds (VOC), hazardous air pollutants (HAP, e.g. benzene, asbestos), and criteria air pollutants (CAP, sums CO, lead, NO_X, ozone, PM, and SO_X). Our social costs are for NO_X, $PM_{2.5}$, SO_2 and VOCs.

[23]For brevity, we omit details on pollution levels for the quasi-likelihood model.

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 all pollutants and the mean social cost of pollution are substantially higher in column (2) than in the baseline. We take this as strong evidence that non-linear fines are effective in lowering pollution, even conditioning on a mean equilibrium level of fines.

Columns (3) and (4) consider changes in the non-linearity of fines that are less dramatic than in column (2). Column (3) halves the fines at every regulatory state Ω' where the plant is in HPV status while column (4) doubles 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. As with column (2), the results are larger for the GMM random coefficient estimates than for the quasi-likelihood estimates. Note also that even though column (3) halves the state-contingent fines while in HPV status, the equilibrium level of fines is higher than the baseline, as plants spend longer in HPV status.

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) consider removing or increasing the plant's cost of being in HPV status.

Focusing on the GMM estimates in 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 increase by 47%.

Turning to columns (3) and (4), eliminating the utility cost to plants for being in HPV status would more than triple the share of plants in HPV status under the GMM random coefficient estimates. This is met with substantial increases in the inspection rate, mean number of violations per non-compliant plant, and both the pollution levels and pollution cost. If instead the utility cost to plants of being in HPV status were doubled, there would be a moderate decline in the rate of plants in HPV status and a small decrease in inspections,

Table 7: Counterfactual Results: Changing Plants' Cost Structure

	(1) Baseline	(2) No cost for HPV status, violations, or inspections	(3) No cost for being in HPV status	(4) Cost for HPV status doubled from baseline
Quasi-likelihood estimates				
Compliance (%)	94.45 (0.14)	93.59 (0.24)	94.05 (0.27)	94.69 (0.16)
Regular violator (%)	4.10 (0.12)	4.10 (0.12)	4.09 (0.12)	4.10 (0.12)
HPV (%)	1.45 (0.04)	2.31 (0.19)	1.86 (0.25)	1.20 (0.08)
Investment rate (%)	0.44 (0.01)	0.44 (0.01)	0.44 (0.01)	0.45 (0.01)
Inspection rate (%)	9.25 (0.05)	9.52 (0.08)	9.36 (0.10)	9.18 (0.05)
Mean fines (1000$)	17.87 (3.06)	27.00 (5.63)	21.75 (5.00)	15.36 (2.56)
Mean violations (%)	0.53 (0.01)	0.62 (0.03)	0.57 (0.04)	0.50 (0.01)
Mean plant utility	−0.005 (0.004)	−0.001 (0.000)	−0.004 (0.004)	−0.006 (0.004)
Mean CAP (tons)	297.7 (1.8)	321.4 (5.8)	308.0 (6.9)	291.5 (2.4)
GMM random coefficient estimates				
Compliance (%)	95.07 (0.07)	94.50 (0.39)	92.16 (1.44)	95.31 (0.18)
Regular violator (%)	3.59 (0.07)	3.53 (0.08)	3.56 (0.08)	3.59 (0.07)
HPV (%)	1.34 (0.04)	1.97 (0.39)	4.28 (1.48)	1.10 (0.16)
Investment rate (%)	0.54 (0.02)	0.55 (0.02)	0.54 (0.02)	0.54 (0.02)
Inspection rate (%)	9.19 (0.05)	9.38 (0.16)	10.25 (0.61)	9.12 (0.06)
Mean fines (1000$)	16.23 (2.08)	24.09 (15.89)	35.35 (10.30)	14.34 (3.18)
Mean violations (%)	0.51 (0.01)	0.58 (0.08)	0.78 (0.13)	0.49 (0.02)
Mean plant utility	0.021 (0.007)	0.005 (0.002)	0.027 (0.008)	0.021 (0.007)
Mean NO_X (tons)	86.1 (0.4)	89.5 (1.7)	101.4 (6.5)	84.4 (0.5)
Mean SO_X (tons)	100.3 (0.8)	105.8 (3.1)	121.6 (8.9)	97.4 (1.0)
Mean lead (pounds)	13.4 (0.1)	14.4 (0.4)	21.1 (3.2)	13.0 (0.2)
Mean PM (tons)	31.6 (0.1)	32.2 (0.4)	34.8 (1.4)	31.3 (0.1)
Mean VOC (tons)	21.1 (0.0)	21.5 (0.3)	24.1 (1.4)	20.9 (0.1)
Mean HAP (tons)	1.5 (0.0)	1.5 (0.0)	1.7 (0.1)	1.5 (0.0)
Mean CAP (tons)	294.4 (1.6)	308.7 (6.7)	366.7 (30.4)	287.0 (2.4)
Mean pollution cost (1000$)	2,578.7 (11.5)	2,670.0 (50.8)	2,971.9 (168.0)	2,531.3 (15.6)

Note: all statistics report the weighted average of the long-run equilibrium mean, weighting with the number of plants by industrial sector 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. Bootstrapped standard errors are in parentheses. Footnote 22 details pollutants and social costs.

violations, and pollution. These results both highlight the importance of a dynamic approach to enforcement and are in line with the work of Evans (2016), who finds that being placed on the EPA watch list as an HPV encouraged plants to resolve their outstanding violations.

Overall, our counterfactuals suggest not only that dynamic enforcement of the CAAA is effective in reducing non-compliance and pollution, but also that it reduces the workload of the EPA by encouraging plants to invest in pollution abatement while still in regular violator status rather than delaying investment until they reach a point where more inspections and violations are being used. Increasing the non-linearity of fines or doubling the cost to plants of being in high priority violator status does decrease non-compliance and pollution further, although the magnitudes of the changes are moderate.

7 Conclusion

This paper empirically evaluates the effectiveness of dynamic regulation in the context of Clean Air Act and Amendments enforcement. We build and estimate a dynamic game model of a plant which is faced with a regulator and which must choose when to invest in pollution abatement. We develop a random coefficients estimator that is computationally tractable and that allows for wide heterogeneity in plants' costs from regulatory scrutiny.

We find that there are substantial costs to plants of investing in pollution abatement and also of facing regulator enforcement actions, particularly fines and designation as a high priority violator. We also find that there is substantial heterogeneity across plants in their regulatory compliance costs.

Counterfactual simulations show that the non-linearity of enforcement activity and fines has led to increased compliance with the CAAA, and hence lower air pollution, while simultaneously decreasing the regulatory burden for state and federal EPAs without dramatically affecting plant investments. While pollution would be much higher under a regulatory framework with linear fines or one that imposed fewer enforcement costs on plants, further increases in the non-linearity of enforcement would not lower pollution much further.

Overall, this analysis provides the first empirical estimates of the plants' responses to the dynamic environmental regulations used around the world. A comparison of our random coefficients estimates to our estimates with homogeneous coefficients shows that heterogeneity in plant costs is particularly important in understanding dynamic regulations. Our modeling framework and results on dynamic enforcement for the CAAA may improve analysis and modeling for the evaluation of dynamic enforcement in a variety of other settings.

References

Aguirregabiria, Victor and Pedro Mira, "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 2007, *75* (1), 1–53.

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," *Econometrica*, 2018, *forthcoming.*

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.

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 Paradox," *Journal of Public Economics*, June 1999, *72* (3), 361–378.

Holland, Mike, Steve Pye, Bert Droste-Franke, and Peter Bickel, "Damages Per Tonne Emission of PM2.5, NH3, SO2, NOx, and VOCs From Each EU25 Member State (excluding Cyprus) and Surrounding Seas," Technical Report 20879, Clean Air for Europe (CAFE) Programme March 2005.

Houde, Sébastien, "How Consumers Respond To Product Certification and the Value of Energy Information," *The RAND Journal of Economics*, 2018, *49* (2), 453–477.

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, Benoît 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.

Lim, Claire SH and Ali Yurukoglu, "Dynamic Natural Monopoly Regulation: Time Inconsistency, Moral Hazard, and Political Environments," *Journal of Political Economy*, 2015.

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.

Shapiro, Joseph S. and Reed Walker, "Why is Pollution from U.S. Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade," Working Paper 20879, National Bureau of Economic Research January 2015.

Shimshack, Jay P., "The Economics of Environmental Monitoring and Enforcement," *Annual Review of Resource Economics*, 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*, March 2002, *70* (2), 603–629.

Train, Kenneth E., *Discrete Choice Methods with Simulation*, Cambridge University Press, 2009.

On-Line Appendix A1: Data Construction Details

Regulatory Actions and Outcomes

Compliance and violator statuses. During our sample period, the EPA's *Air Program Historical Compliance* dataset reported each plant's compliance status for every CAAA program. Since there is a CAAA program for each major category of air pollutant, a plant can simultaneously be in violation of multiple CAAA programs. We assume that a plant is a CAAA violator if it is a violator for any CAAA programs. For each program, we classify a plant as being a violator if compliance status is equal to "1" (in violation, no schedule), "6" (in violation, not meeting schedule), "7" (in violation, unknown with regard to schedule), "B" (in violation with regard to both emissions and procedural compliance), "D" (HPV violation), "E" (federally reportable violation), "F" (High Priority Violator on schedule), "G" (facility registry service on schedule), or "W" (in violation with regard to procedural compliance).[24]

The *Historical Compliance* dataset also reports codes indicating an unknown compliance status: "Y" (unknown with regard to both emissions and procedural compliance), "0" (unknown compliance status), "A" (unknown with regard to procedural compliance), and "U" (unknown by evaluation calculation). From our discussions with the EPA, these codes arise when a plant has not been inspected within the required time frame, but there has been no indication of a violation by the plant. Given this, we code these plants as being in compliance.[25]

We code all other plants—except those that are listed as HPVs in the *High Priority Violator History* dataset—as being in compliance. Thus, we do not use additional information on compliance in the ECHO database for some plants and pollutants, such as continuous emissions monitoring system reports.

Inspections. The *Air Facility System Actions* dataset reports multiple types of inspections, which we collapse into a single "inspection" variable. These include on- and off-site

[24] This list indicates both plants that are regular violators and HPVs. As discussed in Section 3.1, we determined HPV status from the *High Priority Violator History* dataset.

[25] Evans (2016) also considers plants in unknown compliance status to be in compliance.

full compliance evaluations conducted by either the federal or state EPA, partial compliance evaluations, and stack tests. We also consider an inspection to have occurred if the EPA issues a Section 114 letter for gathering information from the plant.[26]

In some cases we observe multiple inspections in the same quarter; e.g., if stack tests are conducted for multiple pollutants. Since our inspection variable is dichotomous, we consider these tests together to be equivalent to a single inspection.

Violations. The *Actions* dataset also reports violations. We define a violation to be the issuance of a "Notice of Violation" (NOV). An NOV is defined as "a notice sent by the State/EPA ... for a violation of the Clean Air Act." There are three codes that indicate an NOV in our data: "6A" (EPA NOV issued), "7A" (notice of noncompliance), and "7C" (state NOV issued).[27]

Plant Exits

The *Historical Compliance* dataset also allows us to understand when plants shut down. Plants may have a compliance status of "9" (in compliance: shut down). If we observe a plant in this status, we assume that it has exited. We remove it from our sample for the period with this status and all subsequent quarters.

Investment

Our data do not directly report investments or investment costs (unlike in the Duflo et al., 2018, study of pollution in India, for instance). Instead, we infer investments from the behavior of EPA regulators. We determine that an investment occurred if we observe any of the following three types of events: (1) the resolution of a major violation, (2) the issuance of a Prevention of Significant Deterioration (PSD) permit, and (3) the exit from HPV status. We now provide detail on each of these categories.

First, as shown in Table 1, the overwhelming majority of our investments come from

[26]In general, Section 114 requests are made by the EPA either to obtain plant operational information for developing future regulation or to determine plant compliance.

[27]See `https://echo.epa.gov/files/echodownloads/AFS_Data_Download.pdf`.

codes that indicate the resolution of a major violation. There are three codes in the *Actions* database that we consider evidence of this type of investment: (1) "VR" or "violation resolved," (2) "OT" or "other addressing action," and (3) "C7" or "closeout memo issued." According to the November, 2008 *Air Facility Systems National Action Types–Definitions* EPA document,[28] "a violation is resolved when it is addressed and a closeout memo has been issued, all penalties have been collected and the source is confirmed to be in physical compliance." Similarly, "other addressing action" is an addressing action for HPV cases with criminal or civil action referrals. Finally, "a closeout memo is issued when a violation is resolved with all penalties collected and the source is confirmed to be in physical compliance."

Second, a PSD permit is required for new pollution sources or for major modifications of existing sources.[29] While it is possible that major modifications of existing sources may occur for reasons other than a plant attempting to return to CAAA compliance, we believe that changes to a plant that were substantial enough to warrant a new PSD permit issuance likely imply a major investment in pollution abatement.

Finally, we also infer that an investment has occurred if a plant exits HPV status, even if we do not observe one of these codes. We make this choice because we believe that a major investment would have been necessary in order to resolve the substantial violations that would have originally merited the determination of HPV status and that the AFS data may be missing these investments.

Our definition of investment will only capture larger investments. In particular, we do not observe changes in paperwork or workplace practices that may be necessary to resolve smaller violations of the CAAA. Thus, we need to interpret our investment cost parameters as representing the cost of major investments. For this reason, our results may understate the impact of environmental monitoring and compliance at inducing plants to undertake pollution abatement investment.

[28] Downloaded September 2014.

[29] See `https://www.epa.gov/nsr/prevention-significant-deterioration-basic-information`.

On-Line Appendix A2: Computational Details

Regulatory CCPs

As described in Section 5, we first estimate the regulatory CCPs, for inspections, violations, fines, and the transition between states, $T(\Omega, e)$. We estimate inspections and violations as probits, and fines as tobits. Each regression is run separately for firms in compliance, regular violators, and high priority violators. Results of these regressions are presented in Table A1

Our estimation of the transitions between the statuses of compliance, regular violator, and HPV is complicated by the fact that there are three outcomes, rather than two. We thus evaluate the CCP in two steps. Specifically, for plants in compliance (based on their lagged reported status), we first estimate a probit specification of the likelihood of staying in compliance. We then estimate a probit specification for the subset of these plants that exit compliance on whether they transition into regular violator or HPV status conditional on exiting compliance. Both specifications include regressors for whether the plant invested in the last period or two periods ago, faced an inspection, whether it received a violation, and the amount of any fine assessed, in addition to region and industrial sector fixed effects. We perform similar two-step estimation processes for the transitions for plants in regular violator and HPV statuses (also based on their lagged reported statuses). We present results of these regressions in Table A2.

Computing the Bellman Equation

The plant's decision as to whether or not to invest at any state is based on dynamic optimization. As such, we solve for the Bellman equation for candidate parameter values, based on equations (2) and (3), to estimate both models. Specifically, for our quasi-likelihood estimator, we perform a non-linear search for θ and hence we solve for the Bellman equation for each of the candidate values of θ that are considered in the course of the non-linear search. For our GMM estimator, we solve for the Bellman equation for each of the 10,001 values in our fixed parameter grid.

The states in Ω and Ω' are discrete, except for depreciated accumulated violations. Our Bellman equation discretizes this latter variable, using 20 grid points that are evenly spaced from 0 to 9.5. The transition from Ω' to Ω, given in (2), will result in a new level of depreciated accumulated violations that does not necessarily correspond to a grid point. As such, we use linear interpolation to calculate (2).

The transition from Ω to Ω', given in (3), is stochastic, as it depends on the regulatory CCP. We perform this calculation by simulating from the estimated regulator CCP. Specifically, we first calculate the inspection probability for each state from the predicted values of our estimates. We then calculate the violation probability for each state and inspection decision. Following this, we calculate the distribution of fines for each state, inspection decision, and violation decision, using 20 evenly spaced points from the estimated residual distribution—which we denote F—ranging from $F^{-1}(0.025)$ to $F^{-1}(0.975)$. Finally, we calculate the transition probabilities between the three statuses of compliance, regular violator, and HPV, for each state, inspection decision, violation decision, and discretized fine decision.

Altogether, this gives 240 ($2\times2\times20\times3$) possible regulatory outcomes using our discretized method. We calculate the probability and mean fines for each one. The Bellman equation then integrates over these possibilities. We compute our Bellman equation until a fixed point, defined as a sup norm tolerance of 10^{-7} between subsequent iterations.

Choice of Fixed Grid Values for GMM Estimation

Our fixed grid estimator requires the ex ante specification of potential parameter grid values. We follow Fox et al. (2016) and first estimate the quasi-likelihood model and then center our fixed grid on these estimates. This requires specifying a range for the parameter grid around the quasi-likelihood estimates. We used a range of 15 (from 7.5 below the quasi-likelihood model to 7.5 above) for investment and 5 for the other parameters. We chose these ranges after experimenting to make sure that they were large enough that we did not have parameters with positive weights near the boundary.

We choose our actual grid values again by following Fox et al. (2016) and using co-prime

Halton sequences for each parameter, using the first five prime numbers as each plant has five parameters. We scale the Halton sequences over the range between the minimum and maximum values. Co-prime Halton sequences better cover the set of parameters than would taking the interaction of the same grid points for each component (Train, 2009).

We dropped the first 20 elements of each Halton sequence as recommended in the literature (Train, 2009). We use the next 10,000 elements of the Halton sequences plus the quasi-likelihood estimates themselves as our fixed grid; hence $J = 10,001$. We also experimented with $J = 8,001$ (using the first 8,000 elements of the Halton sequence) and found similar results.

Inputs to Moments

As noted in Section 5.2, we have three sets of moments. From (7), each set of moments is defined by some m_k^d and $m_k(\theta_j)$. We now detail the three sets of moments.

Our first set of moments is the steady state probability of being at any variable state conditional on the fixed state of EPA region and industrial sector. Specifically, for any moment $G_k(\eta) = m_k^d - \sum_{j=1}^{J} \eta_j m_k(\theta)$, where $k \in 1, \dots, K$ references the specific state $\omega'^1 \in \Omega'^1$ and $\omega'^2 \in \Omega'^2$, we can write:

$$m_k(\theta_j) = \Pr[\Omega'^1 = \omega'^1 | \Omega'^2 = \omega'^2, \theta_j], \tag{A1}$$

and

$$m_k^d = \sum_i \sum_t \frac{\mathbb{1}\{\Omega_{it}'^1 = \omega'^1, \Omega_{it}'^2 = \omega'^2\}}{\mathbb{1}\{\Omega_i'^2 = \omega'^2\}}. \tag{A2}$$

We compute (A1) by solving the Bellman equation given θ_j and then evaluating the steady state distribution under optimizing behavior. We use a matrix inverse formula to solve for the steady state distribution.

We note a few points about these moments. This first set of moments follows 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 11,000 states, we limit ourselves to states which have a probability of over 10^{-5} at our estimated quasi-likelihood parameters.

Our second set of moments also follows closely from Nevo et al. (2016). The m_k values for these moments are constructed from the conditional probability of being at any variable state and having an investment at that state:

$$m_k(\theta_j) = \Pr[\Omega'^1 = \omega'^1 | \Omega'^2 = \omega'^2, \theta_j] \Pr[X = 1 | \Omega', \theta_j], \tag{A3}$$

and

$$m_k^d = \sum_i \sum_t \frac{\mathbb{1}\{\Omega_{it}'^1 = \omega'^1, \Omega_{it}'^2 = \omega'^2, X_{it} = 1\}}{\mathbb{1}\{\Omega_i'^2 = \omega'^2\}}. \tag{A4}$$

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 captures the panel data aspect of investment. The m_k values for these moments are constructed from the conditional probability of being at any variable state and having an investment at that state and times the sum of investments in the next six periods:

$$\begin{gathered} m_k(\theta_j) = \Pr[\Omega'^1 = \omega'^1 | \Omega'^2 = \omega'^2, \theta_j] \times \Pr[X = 1 | \Omega', \theta_j] \times \\ \left(\textstyle\sum_{s=1}^{6} \Pr[X \ s \text{ periods ahead} = 1 | X = 1, \Omega', \theta_j] \right), \end{gathered} \tag{A5}$$

and

$$m_k^d = \sum_i \sum_t \frac{\mathbb{1}\{\Omega_{it}'^1 = \omega'^1, \Omega_{it}'^2 = \omega'^2, X_{it} = 1\} \times \left(\sum_{s=1}^{6} X_{i,t+s}\right)}{\mathbb{1}\{\Omega_i'^2 = \omega'^2\}}. \tag{A6}$$

These moments seek to match the extent of repeated investments by plants in the data—as displayed in Figure 4—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 (A5), we integrate over all potential paths conditioning on the initial state and investment decision. Each period there are ten potential paths: every interaction of (1) investment or not, (2) violation or not, and (3) regular violator and HPV statuses; plus the cases of compliance with and without violations, but without investment.[30] Over 6 periods, this then implies 10^6 = 1,000,000 possible paths for each parameter vector in our fixed grid θ_j. Thus, calculation of m_k for this set of moments is time consuming.

Overall, our estimator has 12,547 moments, composed of 4,229 of the first set and 4,159 each of the second and third set. Our computation of $m_k(\theta_j)$ results in a 12,547 $\times$ 10,001 matrix and takes approximately 3 days on an iMacPro with 8 processors, with code written in C with MPI, or 1 day on the University of Arizona high performance cluster, using 28 processors.

Weighting Matrix and Estimation of GMM Parameters η_j

We follow the standard approach in GMM estimation of weighting by an estimate of the inverse of the variance-covariance matrix to obtain asymptotically efficient estimates. We proceed in two stages. In stage 1, we estimate the model with a variance-covariance matrix that is not necessarily asymptotically efficient. Then, we use our stage 1 estimates to compute an approximation to the efficient weighting matrix. In stage 2, we reestimate our parameters using this asymptotically efficient weighting matrix. We now detail our computation of the variance-covariance matrix for both stages.

In stage 1, we calculate the variance-covariance matrix of the moments inputs m_k, at the quasi-likelihood estimates θ_Q. The diagonal elements of this matrix are calculated as:

$$Var(m_k(\theta_Q)) = \frac{E[m_k(\theta_Q)m_k(\theta_Q)] - E[m_k(\theta_Q)]^2}{N_k}, \tag{A7}$$

[30] To save computational time, we use the higher probability point for depreciated accumulated violations, rather than linear interpolation.

where N_k is the number of plant-quarter observations from the EPA region and industrial sector for moment k. This is the general formula for the variance for the mean of N_k repeated $i.i.d.$ draws from a random variable.

For the off-diagonal elements, the covariance will be zero for moments in different EPA regions or industrial sectors. We can write the covariance between moments k and l from the same EPA region and industrial sector as:

$$Cov(m_k(\theta_Q), m_l(\theta_Q)) = \frac{E[m_k(\theta_Q)m_l(\theta_{QL})] - E[m_k(\theta_Q)]E[m_l(\theta_Q)]}{N_k}. \tag{A8}$$

The first term in (A8) will be non-zero only for the three moments that pertain to the same state. In this case, the first term in the numerator of the covariance between the first and second set of moments will equal the second moment, while the first term in the numerator between the first and third set of moments or between the second and third set of moments will equal the third moment. The reason for this is that the moment from the second set will only be non-zero when the moment from the first set is non-zero, while the moment from the third set will only be non-zero when the moment from the second set is non-zero. The second term in (A8) is simply the product of the means.

In stage 1, we invert and take a Cholesky decomposition of this estimated variance-covariance matrix. We then pre-multiply $m_k(\theta_j)$ for each θ_j and m_k^d by this matrix and obtain stage 1 estimates of the weights η_j by minimizing the linear system of equations in (7) subject to the constraints in (6), via constrained least squares. We use the Matlab package lsqlin to perform this minimization process, which takes approximately 10 minutes on an iMacPro. The result is consistent, though not necessarily asymptotically efficient, estimates of η.

We then estimate the variance-covariance matrix of $G(\eta)$ using our stage 1 GMM estimates of η. From (7), the variance of $G(\eta)$ is simply the squared weighted sum of the variance conditional on the individual parameters, since the probability of each individual parameter occurring is independent across observations.

We again take a Cholesky decomposition of the inverse of this revised variance-covariance matrix, pre-multiply the matrix of moments $m_k(\theta_j)$ across all θ_j values, and re-run our estimation of the η_j weights. This provides our stage 2 estimates of η_j, which are the ones that we report.

Bootstrap Procedure for Inference

We bootstrap to obtain standard errors for both our quasi-likelihood and GMM estimates. For our GMM estimates, we provide standard errors on the counterfactual estimates only rather than also on the structural parameters.

Our bootstrap for the GMM estimator proceeds with the following repeated procedure:

1. We first draw an alternative dataset with sampling with replacement at the plant level. The new dataset has the same number of plants as the original data, though not necessarily the same number of plant-quarter observations.

2. We then use this new dataset to recalculate the regulatory CCPs.

3. Using these functions, we calculate the inputs to the moments, $m_k(\theta_j)$ and m_k^d. We limit the moments to only those states with at least a long-run probability of 1e-5. This means that the exact number of moments, m_k, varies across iterations of the bootstrapping procedure.

4. We then calculate our initial weighting matrix and estimate our first-stage GMM structural parameters η using this weighting matrix.

5. We then calculate the asymptotically efficient weighting matrix for the moments based on these first-stage estimates, and use this weighting matrix to re-estimate the structural parameters.

6. Finally, we use these estimates to calculate all of the outcomes for each potential counterfactual. We report the standard deviation of the outcomes across the bootstrap iterations as the standard error of our counterfactual outcomes.

We report results from 100 bootstrap draws, using the University of Arizona high performance cluster to perform the computations simultaneously. Our bootstrap for the quasi-likelihood process is similar: it uses the output created in steps 1 and 2 above. It then estimates the structural parameters with a non-linear search (analogous to step 4) and performs the counterfactual computation with the new structural parameters, regulator CCPs, and dataset (analogous to step 6).

For two draws, Stata did not converge when performing the probit estimation of the regulatory CCPs (presumably due to numerical imprecision) and so we skipped those draws. In one case, Matlab reported that the weighting matrix was not positive definite (again presumably due to numerical imprecision) and so we skipped that draw.

On-Line Appendix A3: Extra Figures and Tables

Figure A1: Mean Inspection Probabilities and Fines by EPA Region

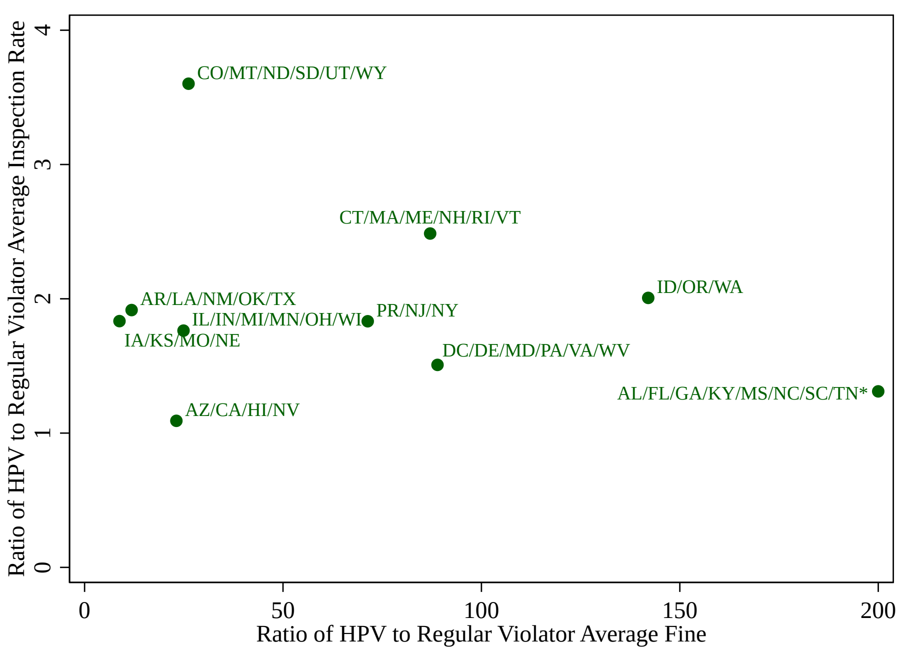

Note: authors' calculations based on estimation sample. States in each region are indicated next to value. Observations marked with a * are top-coded at a ratio of HPV to regular violator average fine of 200.

Figure A2: Mean Inspection Probabilities and Fines by Industrial Sector

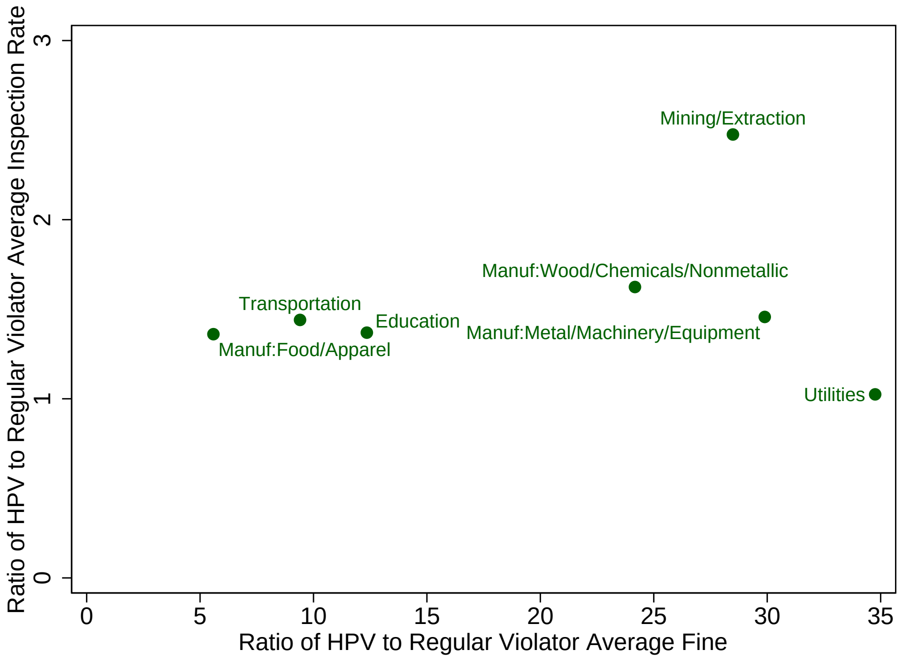

Note: authors' calculations based on estimation sample. Industrial sector measured by 2-digit NAICS code.

Figure A3: Further Investments by a Plant Following Initial Investment, in Data and Models

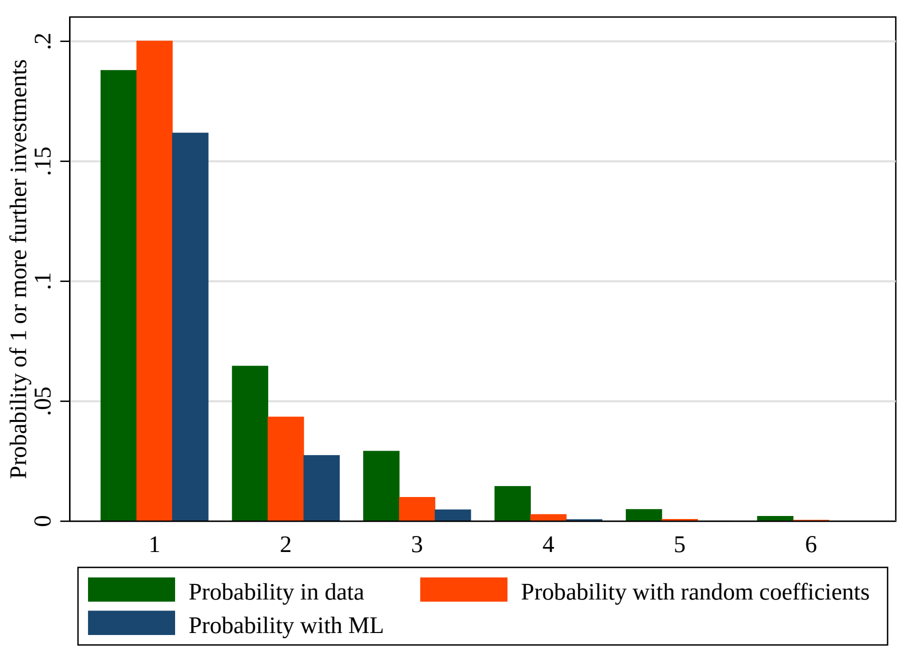

Note: authors' calculations based on estimation sample and estimated models evaluated at steady state.

Table A1: Regulatory CCPs: Inspections, Violations, and Fines

	Regular violator			High priority violator		
Dependent variable:	Inspection	Violation	Fine amount (1000$)	Inspection	Violation	Fine amount (1000$)
Depreciated Accumulated Violations	0.398***	0.332***	2.225***	0.285***	0.235***	3.938***
	(0.009)	(0.013)	(0.089)	(0.008)	(0.009)	(0.161)
Inspection		0.778***	1.278***		0.517***	2.437***
		(0.022)	(0.154)		(0.020)	(0.347)
Violation			1.555***			5.068***
			(0.276)			(0.496)
Lag of Investment	0.189***	-0.127***	-0.522*	0.025	-0.194***	-7.537***
	(0.025)	(0.046)	(0.297)	(0.018)	(0.026)	(0.491)
2nd Lag of Investment	0.345***	-0.016	-0.094	0.127***	0.171***	1.568**
	(0.024)	(0.041)	(0.271)	(0.033)	(0.040)	(0.675)
σ			6.427			18.781
			(0.101)			(0.189)
Pseudo R^2	0.093	0.152	0.072	0.071	0.095	0.036

Note: table shows results from probit regressions for "inspection" and "violation" 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 industrial sector dummies. For brevity, the table omits results for plants that are in compliance. Standard errors are clustered at the plant level.

Table A2: Regulatory CCPs: Status Transitions

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 not into comp.	Stay in HPV	Into compliance given out of HPV
Lag of Investment			1.171***	0.107	-2.527***	-0.179
			(0.024)	(0.076)	(0.122)	(0.352)
2nd Lag of Investment			0.625***	-0.282***	0.533***	-0.186*
			(0.024)	(0.053)	(0.070)	(0.104)
Depreciated Accumulated Violations			0.169***	-0.178***	0.108***	-0.340***
			(0.012)	(0.017)	(0.017)	(0.025)
Inspection	-0.700***	-0.478***	-0.011	-0.470***	0.128***	-0.057
	(0.010)	(0.029)	(0.016)	(0.034)	(0.039)	(0.042)
Violation	-4.138***	0.016	-2.017***	-1.464***	1.470***	-1.636***
	(0.024)	(0.028)	(0.103)	(0.039)	(0.066)	(0.159)
Fine	-0.006***	-4.503***	-8.243***	-0.012**	0.062**	0.017
	(0.001)	(0.866)	(0.357)	(0.005)	(0.027)	(0.050)
Pseudo R^2	0.528	0.139	0.142	0.349	0.402	0.135

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