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

Wesley Blundell[†] Gautam Gowrisankaran[‡] Ashley Langer[§]

July 14, 2019

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 value 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 costly to most plants. Eliminating dynamic enforcement would have large adverse impacts on the number of high priority violators and pollutants emitted.

JEL Codes: Q53, Q58, C57
Keywords: monitoring and compliance, pollution, investment, dynamic estimation, Clean Air Act

*We thank Christian Belzil, Esther Duflo, Liran Einav, Paul Ellickson, Meredith Fowlie, Wayne Gray, Alex Hollingsworth, Karam Kang, Anatolii Kokoza, Shaun McRae, Derek Lemoine, Andy McLennan, Paul Scott, Ivan Rudik, Nick Ryan, Jay Shimshack, Kenneth Train, Kyle Wilson, Jonathan Williams, Tiemen Woutersen, Chenyu Yang, anonymous referees, and seminar participants at numerous institutions for helpful comments. We are grateful to Yujia Peng for excellent research assistance. 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.

[†]California State University, East Bay
[‡]University of Arizona, HEC Montreal, and NBER
[§]University of Arizona

1 Introduction

In the United States, the Clean Air Act and its amendments reduced damages from air pollution by $35.3 trillion from 1970 to 1990. However, these regulations impact nearly every industrial facility in the U.S., leading to combined enforcement and compliance costs to governments and plants of $831 billion over this period (Environmental Protection Agency, 1997, converted to 2007 dollars). These massive costs make it critical to understand the efficiency of regulatory monitoring and enforcement mechanisms for pollution control.

To better understand how environmental regulations are enforced, we first consider an example of a large oil refinery in Texas.[1] In 2011, after a period with only low-level violations involving emission releases, the plant was conducting work to improve efficiency when a valve that should have been left open was closed. This led to a pressure buildup in a pipeline, causing a leak and emissions of volatile organic compounds and benzene. Because these emissions came from an unauthorized source, the plant was placed in *high priority violator* (HPV) status, subjecting it to higher scrutiny and fines. In 2012, another low-level pollution release similar to the earlier ones occurred, but this time the fine imposed was doubled because the plant was in HPV status. Increased scrutiny and enhanced fines continued through a series of additional releases until the plant made two separate investments in pollution abatement—including upgrading a large number of emissions-monitoring systems—after which the plant was removed from HPV status, returning to a baseline level of scrutiny in 2013.

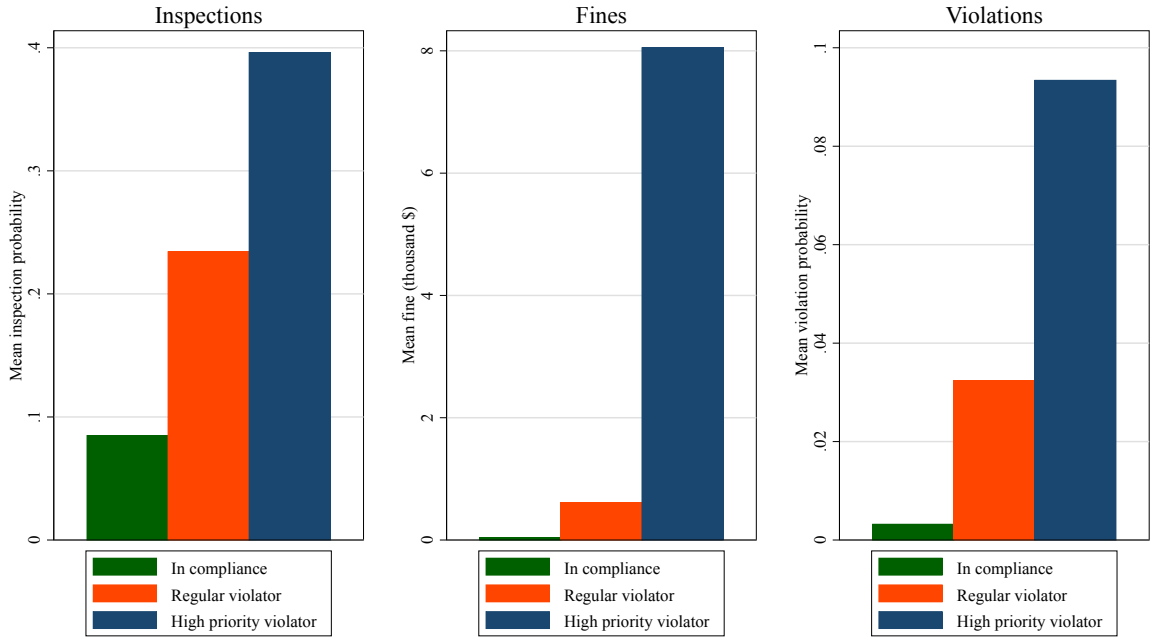
This example illustrates one way that the U.S. Environmental Protection Agency (EPA) uses *dynamic enforcement*—where regulatory actions are a function of the plant's regulatory state (Landsberger and Meilijson, 1982; Shimshack, 2014)—to enforce the Clean Air Act Amendments (CAAA): the EPA designates repeat offenders as HPVs and targets them with elevated scrutiny and penalties. Regulators may choose dynamic enforcement because it avoids over-fining plants before they have a chance to fix violations, but uses the threat of high fines as an incentive for plants to make costly investments in pollution abatement. Dynamic enforcement may add value when the imposition of fines is costly to the regulator,

[1] We obtained the information underlying this example from Texas and federal enforcement records.

when there are delays in remediation, and when the regulator cannot contract on a plant's compliance costs in its regulatory actions.

CAAA enforcement incorporates substantial state-dependent scrutiny, in part through HPV status designation. To illustrate this, Figure 1 shows mean unconditional CAAA inspection rates, fines, and violation rates, separately for plants in compliance, regular (not-high-priority) violators, and HPVs. In each case, the level of scrutiny increases dramatically across these statuses.[2]

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



Note: figure reports 2007-13 unconditional mean quarterly levels of inspections, fines, and violations by CAAA regulatory status, based on authors' calculations from the estimation sample.

This paper seeks to quantify the gains from dynamic enforcement of the CAAA. To do this, we first estimate the cost to industrial facilities of complying with the EPA's current dynamic approach. Second, we simulate the value of alternative enforcement regimes in affecting plants' emissions and compliance with the CAAA. Our modeling and estimation

[2]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.

framework are specific to the CAAA, but we believe that similar approaches may yield important general insights, since dynamic enforcement is used across many settings,[3] and it is widely recognized that dynamic enforcement can add value in theory (e.g. Landsberger and Meilijson, 1982; Harrington, 1988; Leung, 1991; Polinsky and Shavell, 1998).

Until recently, the empirical literature on enforcement has largely focused on evaluating how firm compliance relates to enforcement.[4] For instance, Magat and Viscusi (1990), examine whether inspections lower emissions at a plant, while Nadeau (1997) uses variation across plant types and states to look at the effect of enforcement on the duration of non-compliance. More recently, Shimshack and Ward (2008) show that increased enforcement can lead even compliant plants to reduce emissions, leading to "over-compliance," where plants emit well below the compliance threshold, and Stafford (2002); Keohane et al. (2009) and Blundell (2019) examine how variation in the intensity of dynamic enforcement relate to plants' compliance status.

However, this empirical literature has not been able to evaluate the value of dynamic enforcement. To do this, one needs to account for the value of regulation in lowering pollution and weigh that benefit against the compliance cost to plants and regulators. With dynamic enforcement, measuring this value requires estimating a dynamic model of the costs to plants from investment in pollution abatement relative to the costs of regulatory scrutiny. These are some of the many challenges that lead Shimshack (2014) to conclude that environmental monitoring and enforcement is "both understudied and controversial" (p. 3).

To measure the value of dynamic enforcement we develop and estimate the parameters of a discrete-time dynamic game of a plant faced with a regulator. In our model, the regulator makes decisions regarding inspections and fines. Inspections can improve the ability of the regulator to detect violations of the CAAA, such as increased air pollution. The regulator

[3]E.g., for the CAAA (Evans, 2016) 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 beyond environmental regulations, e.g., in worker health and safety regulation (Ko et al., 2010) and tax auditing in China (Maitra et al., 2007).

[4]Exceptions are two recent papers on the value of regulatory discretion, which we discuss below (Duflo et al., 2018; Kang and Silveira, 2018).

then uses its information to decide whether to transition plants to regular violator or high priority violator status and to determine whether violations have occurred, either of which can subject plants to higher inspection rates and higher fines. The regulator bears a cost from conducting inspections and imposing fines. To avoid making assumptions about the EPA's objective function, which we would have difficulty identifying, we do not estimate the regulator's utility function, but rather model the regulator's decisions using conditional choice probabilities.

Each plant, in turn, decides whether and when to invest in pollution abatement technologies. We allow for plants to bear costs from both regulatory actions (e.g. shutting down a production line to allow for an inspection) and investment in pollution abatement. Therefore, a plant that is in regular or high priority violator status will consider incurring the investment cost in order to reduce its present discounted value of future regulatory costs. Recovering these costs is the key step in being able to understand how plants will respond to counterfactual regulatory policies, such as those that do not condition enforcement activities on plant state.[5]

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 heterogeneity in plants' costs of investment and enforcement.

We estimate two main econometric models. First, we implement a quasi-likelihood fixed point estimator where mean investment and regulatory costs are identical across plants.[6] The specification and estimation of this model follow from Rust (1987)'s classic work on bus engine replacement. Second, we estimate a model where plants have heterogeneous costs. Specifically, we estimate a non-parametric random coefficients model similar to Fox et al.

[5]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.

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

(2016) where the structural parameters are the population weights of observing each of a grid of potential cost parameters. We use a generalized method of moments (GMM) estimator that is computationally very tractable, with a quick and convex optimization problem.

To intuitively understand how our model is identified, begin by imagining that we observed plants' expectations of future regulatory actions conditional on whether the plant invested in the current period or not, and that we had exogenous variation in these expectations. The cost to the plant of investment and regulatory actions could then be estimated from a static discrete choice model of plant investment decisions where the independent variables are a constant (interpreted as the cost of investment) and the change in each of the regulatory actions that the plant expects conditional on investment. Our quasi-likelihood model essentially performs this regression, also allowing plants to have the option value of investment in future periods, and using the variation in the benefit from investment across locations and industries. The quasi-likelihood model is thus identified by plants' investment response to variation in the return to investment in affecting future regulatory actions. Our second model recovers the heterogeneity in utility parameters by matching plants' steady-state equilibrium regulatory states and investment rates to our data. If observably similar plants have a high probability of being in very different states, this suggests underlying differences that the heterogeneous cost model will pick up. The panel nature of our data is useful here because it also allows us to match the serial correlation in plant investment to data.

Summary of results. The estimates from the homogeneous cost model confirm the findings in Duflo et al. (2018) that plants find environmental enforcement expensive. 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 \$10,900 per quarter. This cost may occur from high priority violator status affecting the plant's reputation and relationship with the surrounding community. We also find that investment and inspections are costly to plants, with a new investment being equivalent to a \$480,000 fine (holding constant other regulatory actions) and an inspection being equivalent to a \$8,200 fine. Because a fine may impose cost to plants beyond the amount assessed by

the EPA, these figures should be interpreted as lower bounds on the actual dollar cost to plants of investment and enforcement actions.

The estimates from our random coefficients model show that there are substantial differences in plants' regulatory and investment costs, but that nearly all plants find investment, fines, and HPV status substantially costly. In particular, about 96% of plants have investment costs of between $218,000 and $450,000 in equivalent fine costs, which are similar to, but lower than, the homogeneous cost model. This heterogeneity of compliance costs across plants may further increase the value of dynamic enforcement given that the regulator cannot contract enforcement decisions on plant costs.

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) changing the rate at which fines escalate as plants move from regular to high priority violator status and (2) how the current system of dynamic enforcement compares to a Pigouvian fine structure where plants pay a fine equal to the damages from their pollution. We examine these changes would affect plant utility and investment decisions, equilibrium plant compliance rates, regulatory actions, and overall pollution damages.

We find that if we were to completely eliminate any escalation of fines by making fines the same for all plants in (regular or high-priority) violator status, while keeping total equilibrium fines the same as in the baseline, we would have more than 22 times as many plants in HPV status, leading to 267% higher pollution damages and requiring regulators to conduct 220% more inspections. Similarly, eliminating fine escalation while holding pollution damages constant leads to over 600% higher fines. In contrast, increasing the rate that fines escalate in HPV status would have only a small effect on long-run regulatory states and pollution damages. Finally, Pigouvian fines would lead to lower pollution, but at the cost of the regulator needing to impose substantially higher fines.

Relation to literature. This paper relates to three distinct literatures. First, as noted above, we connect the theoretical and empirical literatures on the role of dynamic enforcement of environmental regulations by explicitly estimating the cost of investment in pollution

abatement. 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; Muehlenbachs, 2015; Fowlie et al., 2016; Duflo et al., 2018; Houde, 2018; Kang and Silveira, 2018). Two recent papers estimate the value of regulatory discretion and compliance. Duflo et al. (2018) estimate a dynamic model of environmental regulatory enforcement for plants in India given regulator discretion over which plants to inspect. The dynamics in their context stem largely from the fact that plants must make discrete investment decisions, rather than from dynamic 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. Kang and Silveira (2018) also seek to understand the value of regulatory discretion, by estimating a game between the regulator and municipal water treatment plants in California. 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.

Third, we apply techniques from a non-parametric 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 similar computational techniques.

The remainder of the paper is organized as follows. Section 2 documents the regulatory context and then presents our general theoretical model, that is designed to capture the key features of this environment. Section 3 details our data and provides reduced-form evidence that motivates the choices in our estimable model. Section 4 describes our empirical framework. Section 5 presents our results and counterfactuals. Section 6 concludes.

2 Dynamic Enforcement in Practice and Theory

2.1 Dynamic Enforcement Under the Clean Air Act Amendment

Congress passed the Clean Air Act in 1963 in an effort to improve air quality. While the original Act mostly provided funds for research into monitoring and limiting air pollution, a series of amendments starting in 1965 codified air pollution standards and federal enforcement of these standards. Following the National Environmental Policy Act of 1969 and the 1970 Clean Air Act Amendment, President Nixon, with the approval of Congress, created the Environmental Protection Agency (EPA) to enforce air pollution standards and other environmental legislation. The Act was last amended in 1990 to expand the scope of regulated air pollutants and increase federal enforcement authority. While the current state of air pollution regulation is the result of both the Clean Air Act and a long series of amendments, the Clean Air Act Amendments (CAAA) combine to form the current structure of air pollution regulation enforcement. We will refer to the CAAA in what follows.

The CAAA give the EPA the authority to regulate criteria air pollutants—ozone (O_3), particulate matter (PM), carbon monoxide (CO), nitrogen oxides (NO_X), sulfur dioxide (SO_2), and lead—as well as various hazardous air pollutants. The CAAA mostly mandate command-and-control regulations, the most relevant of which to our context is the requirement that plants use the best available control technology (BACT) in their production processes.[7] Regulations such as BACT generally require that plants' pollution be below thresholds that could be achieved with best practices. To ensure that plants comply with these regulations, the EPA has developed an enforcement regime that includes a system of permitting, inspections, violations, fines, and other requirements (e.g. self-reporting paperwork). This enforcement structure aims to reduce pollution by ensuring that plants are complying with the CAAA emissions and technology standards and by moving plants that

[7] While most of the CAAA are focused on command-and-control regulations, there have been some market-based approaches to regulation such as the NO_X cap-and-trade program. However, these regulations incentivize reductions in NO_X emissions beyond what is achieved via the CAAA's requirement that plants use the best available pollution control technologies. Importantly, plants cannot simply purchase cap-and-trade permits to ensure CAAA compliance.

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 CAAA and the EPA, much of the actual enforcement activity is carried out by regional- and state-level environmental protection agencies.[8] In particular, the EPA divides the country into 10 geographic 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, and assist states with major cases. Further EPA guidance explicitly states that "regions and states can take varied approaches to improving state enforcement programs" (Environmental Protection Agency, 2013, p.5). 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.

Under the enforcement system used during our sample period, all plants that are in compliance could expect to be inspected regularly. The frequency of these inspections depended not only on baseline differences across states and regions in enforcement budget and priorities, but also on the size of the plant and whether the plant was located in a National Ambient Air Quality Standards (NAAQS) non-attainment region. Counties in non-attainment status were required to have plans to return to attainment, which could lead to increased levels of scrutiny for plants in these counties.

In addition to conducting inspections and identifying violations, the EPA can issue fines to plants. Fines are calculated using two main components, the gravity of the violation and the economic benefit that the plant received from the violation (Environmental Protection Agency, 1991). The gravity component of each violation is primarily determined from the actual or potential harm of the violation, which includes (a) the level of the violation, (b) the toxicity of the pollutant, (c) the sensitivity of the environment into which the pollutant is released, and (d) the length of time of the violation. Additionally, gravity is adjusted based on a number of other factors including whether there were reporting issues (e.g. permitting

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

and self-reporting violations), the plant's history of noncompliance, and the plant's ability to pay. Our modeling of regulator fines takes these features into account through the plant's history of violations and recent investments and a series of fixed effects that seek to capture a plant's economic benefit of noncompliance, and gravity, based on the plant's industry and location. Finally, because of bankruptcy laws, political pressure, and explicit caps, the EPA is limited in the penalties it can assess. In particular, 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.

In the course of an inspection, or via a plant self-report, regulators may uncover a violation of the CAAA, and the plant will enter "violator" status. Being a violator subjects the plant to additional inspections, which could possibly uncover additional violations, and potential fines. Plants can accumulate multiple violations within violator status and will only return to compliance once those violations have been resolved. The cost to the plant of being a violator therefore comes not only from the investment cost required to resolve outstanding violations, but also from an increased level of regulatory oversight.

The EPA can designate plants with particularly egregious or repeated violations as "High Priority Violators" (HPV). The HPV designation is explicitly "designed to direct scrutiny to those violations that are most important" (Environmental Protection Agency, 1999, p.1-1) and, during our time period, is reserved for plants that meet one of ten "general" HPV criteria or five "matrix" criteria. While some violations unambiguously merit HPV designation (e.g., "Failure to obtain a Prevention of Significant Deterioration or New Source Review permit"), others either leave room for regulator discretion (e.g., "Substantial testing, monitoring, recordkeeping, or reporting violation") or are explicitly dynamic (e.g. "Violation by a chronic or recalcitrant violator"). Once a plant enters HPV status, it triggers a period of intense oversight by the EPA that includes more frequent inspections (which can lead to uncovering additional violations), higher fines, and explicit deadlines for both EPA and plant actions to resolve any outstanding violations. Plants in HPV status face higher regulatory burdens, as shown in Figure 1. As with the Texas example, plants can only exit HPV sta-

tus after resolving *all* 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 EPA further used a "watch list" to focus particular attention on HPVs that did not resolve all of their violations in a timely manner. While the watch list was originally intended to target oversight by the EPA, reporters obtained the list of included plants via Freedom of Information Act requests and publicized the list. The public disclosure of the watchlist appears to have increased plants' costs by leading to increased attention from local politicians and civilian environmental protection groups (Evans, 2016). This is in keeping with evidence from Johnson (2016) which finds that publicizing non-attainment status (in that case for OSHA regulations) can be costly to plants.

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 General Theoretical Framework

Our theoretical model of EPA enforcement and plant investment seeks to capture the framework described above in a tractable setting. Our model builds on a literature on rational compliance and optimal punishment, with a seminal conceptual model from Becker (1968). We adopt the view (from Becker) that compliance with environmental regulations is a rational decision, where a plant chooses its compliance decisions in order to maximize its surplus.

Landsberger and Meilijson (1982) expand the Becker framework to consider dynamic enforcement in a two-period model of tax compliance. They focus on policies that vary an individual's audit rate (similar to our inspection rate) based on her previous detected violations. Harrington (1988) analyzes dynamic enforcement with a similar framework, where the

regulator underpenalizes one-time violations in order to create incentives to avoid repeated violations. Mookherjee and Png (1994) generalize this idea of differential enforcement activities in a static model by formalizing the concept of *marginal deterrence*, where the regulator underpenalizes small violations in order to create strong marginal incentives to avoid large violations. These policies are both examples of what we call *escalation mechanisms*, where marginal deterrence is increasing in the extent of the violation or history of violations.

Most of the theoretical papers on escalation mechanisms show that increasing marginal deterrence can increase surplus given an implicit or explicit cost of penalties for the regulator (Landsberger and Meilijson, 1982; Harrington, 1988; Mookherjee and Png, 1994; Polinsky and Shavell, 1998; Friesen, 2003). As we noted in Section 2.1, the EPA faces such costs in enforcing the CAAA. In addition, some models specify heterogeneous plants and an inability of the regulator to contract on types as a reason for escalation mechanisms (Landsberger and Meilijson, 1982; Mookherjee and Png, 1994; Raymond, 1999; Kang and Silveira, 2018). In this case, escalation mechanisms can add value by creating a separating equilibrium across types. For instance, with heterogeneous investment costs, an escalation mechanism may incentivize low-cost plants to invest in pollution abatement when they are regular violators and fines are low while high-cost plants will wait until they become HPVs and fines are higher.

Our model of dynamic CAAA enforcement builds on these insights. We model each plant as playing a dynamic game with the regulator, with the equilibrium of the game being Markov Perfect.[9] The regulator would like plants to comply with environmental regulations, but also bears a cost from issuing fines. Violations of the CAAA arise stochastically and plants detect them concurrently with the regulator. Plants make optimizing decisions about whether to invest in remediation of CAAA violations. These investments take time and are not always successful in fixing the violation. We allow for an escalation mechanism with dynamic enforcement, as is present in the data. We also allow for heterogeneous plants and an inability to contract on plant type. The underlying reasons for dynamic enforcement in our setting are a regulator cost of issuing penalties; heterogeneous plants; the delay and

[9] As we discuss in Section 4.1, our estimation is also consistent with a plant playing against a "regulatory machine," as in Duflo et al. (2018).

stochasticity in remediation from investment; and imperfect information from inspections.

More specifically, each period t corresponds to a quarter and the future is discounted with factor β.[10] We define the *regulatory state* Ω_t to be the payoff-relevant state variables over which plant and regulatory actions may depend; Ω_t is known to the regulator and plant at the start of the period.

Each 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 the inspection probability be given by $\mathcal{I}(\Omega)$ and let $Ins(\Omega_t)$ denote the actual inspection decision. The regulator and plant receive a (unidimensional) signal e_t about the plant's environmental performance, based in part on the inspection.

The state Ω and signal e have three effects. First, Ω and e reveal whether there is a new violation, with the function $Vio(\Omega, e)$. Second, the regulator's fine policy, $Fine(\Omega, e)$, is a function of both. Third, they also determine the regulatory state that the plant will face at the point when it takes its action, which we denote $\tilde{\Omega}$. Let $T(\Omega, e) = \tilde{\Omega}$ denote this transition function and let $HPV((\tilde{\Omega}))$ denote HPV status designation under $\tilde{\Omega}$. In our framework, we allow the regulator to vary $Fine(\cdot, \cdot)$ but not $Vio(\cdot, \cdot)$, $T(\cdot, \cdot)$, or $HPV(\cdot)$, as we assume these latter three reflect environmental norms and are dictated by e. Following this, the plant, if not in compliance under $\tilde{\Omega}$, makes a binary decision of whether or not to invest in pollution abatement. Let $X \in \{0, 1\}$ denote the investment decision.

A plant chooses its investment decision in order to minimize its expected discounted sum of the costs from inspections, fines, violations, designation as a high priority violator, and investment.[11] A plant that invests incurs a cost from its investment, but increases the chance that it returns to compliance in future periods. The regulator chooses its inspection and fine policies to minimize the expected weighted sum of damages from pollution, plant investment costs, and enforcement costs.[12]

[10] While we capture *exogenous* plant exit through the discount factor, with a lower discount factor corresponding to more exit, we do not endogenize exit. 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 less likely to be at the margin for exit.

[11] Since we do not incorporate endogenous exit in our model, we do not model the profit from operations.

[12] Our estimation allows for escalation mechanisms as present in the data but does not directly estimate

2.3 Illustrative Simple Case of our Model

In order to further illustrate the value of dynamic enforcement, we develop a simple, special case of our general model, that is similar to Polinsky and Shavell (1998). For this special case, we assume a two-period model with $\beta = 1$. In both periods, at most one new violation occurs with probability p. Inspections occur with probability $\mathcal{I}(\Omega) = 1$, are costless to the regulator and plant, and perfectly reveal the presence of a violation. Thus, the signal e_t indicates the number of outstanding violations and $e_t \in \{0, 1, 2\}$. Violations are also costless. HPV status is also costless and hence not modeled.

A period 1 investment clears a period 1 violation with probability q; violations are never cleared without investment. The pollution cost to the regulator is $c_E e_t$ at period t, for some marginal pollution damage parameter c_E. The regulatory state records the history of investments and violations. Thus, for example, at period 2, the regulatory state after the inspection is $\tilde{\Omega}_2 = (X_1, e_1, e_2)$. Finally, the per-period objective function to the plant is $-\theta^X X_t - Fine(\Omega_t, e_t)$, where θ^X is the cost of investment. The regulator minimizes the sum over the two periods of $c_E e_t$, $\theta^X X$, and its cost of fines.

We focus on the case with a period 1 violation—so $e_1 = 1$, as this is the only case where the regulator might want to incentivize period 1 investment. The simplest policy that a regulator could choose would be a linear fine policy $c_F e_t$. When θ^X is known and contractable and the cost of investment is sufficiently low relative to other costs, the regulator incentivizes period 1 investment by choosing the lowest c_F that would compel the plant to invest.

With a linear fine policy, the regulator has to issue fines for the period 1 violations even though this has no effect on investment. Thus, this fine lowers the regulator objective function. An alternative is for the regulator to choose a static escalation mechanism: it could fine only when $e_t = 2$, which would remove the cost of fining when $e_t = 1$ but the plant has not had a chance to invest, and would still incentivize investment in period 1. For this reason, the regulator can incentivize investment for the same values of θ^X as the linear fine policy with lower expected fines, thereby adding surplus. This policy is not explicitly

the regulator's objective function.

dynamic (since it does not depend on the regulatory state Ω_t, but only on the current signal e_t) but increases marginal deterrence in period 2 since it will result in no fines in period 1. Because expected fines are lower, the regulator will further choose to incentivize investment for more values of θ^X, thereby adding further surplus in some cases.

A dynamic escalation mechanism would increase surplus relative to the static escalation mechanism. In this case, the regulator could fine when $e_{t-1} > 0$, $X_{t-1} = 0$, and when it wants to incentivize investment. Choosing this policy for the same set of θ^X as above will mimic the same investment incentives but with no fines paid in equilibrium (since plants whose investment does not succeed in returning the plant to compliance are not fined), and hence no fine costs. Thus, the regulator will choose to incentivize investment for even more values of θ^X.

If instead, the regulator faces a distribution of θ^X values and cannot contract on θ^X, dynamic enforcement also adds value by better selecting the set of plants which it incentivizes the invest. For simple investment cost type distributions, the same relationship hold: the regulator will incentivize investment for more values of θ^X with dynamic enforcement than with a static escalation mechanism which would in turn incentivize investment for more values of θ^X than with linear fines.

Overall, our illustrative simple case shows that escalation mechanisms add value by increasing the marginal deterrence for period 2 violations relative to period 1 violations. Dynamic escalation mechanisms add more value by increasing the set of actions over which the regulator can condition. Our estimation recovers the regulator's actual enforcement mechanisms and plants' costs from enforcement and investment. Our counterfactuals investigate the extent to which dynamic enforcement leads to lower fines or pollution relative to static enforcement.

3 Data and Reduced Form Evidence

3.1 Description of Data

Our main analyses primarily use four publicly available databases: two CAAA enforcement databases, a pollution damage database, and a database of pollution levels. The first enforcement source is the Environmental Compliance History Online (ECHO) database, which records information on investment and regulatory compliance at the plant level over time. The second enforcement database records the National Ambient Air Quality Standard (NAAQS) attainment status of each county for each pollutant in each year. We combine these enforcement databases with county-level marginal pollution damages from the AP3 model and plant-level emissions information from the National Emissions Inventory (NEI). We now explain our use of all four databases, with 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 four ECHO components: (1) the *Facility Registry Service* dataset, (2) the *Air Facility System Actions* dataset, (3) the *Air Program Historical Compliance* dataset, and (4) the *High Priority Violator History* dataset. 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, for EPA region, and for the gravity of violations based on sector and county. We keep seven sectors with high pollution damages that we believe to have plants of broadly comparable costs of investment and enforcement: 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 regu-

lators, from Q4:2006 through the Q4:2014.[13] 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,[14] and the penalty amount when the action is a fine.[15] 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.[16] This measure of investment is imperfect in that it only captures large (likely capital) investments rather than smaller investments in improving plant processes that may also reduce pollution. To our knowledge, there is no comprehensive national database that contains these types of smaller process investments.

Third, the *Air Program Historical Compliance* dataset records the historical compliance status for each plant and EPA program at the AFS ID and quarter level. These data derive from a combination of self-reports by plants and regulator inputs. We follow the literature (Laplante and Rilstone, 1996; Shimshack and Ward, 2005) in treating the self-reported data as accurate.[17] We use this dataset to determine whether a plant is in compliance or a violator

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

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

[15]It is possible for plants to contest fines in court. However, Helland (2001) finds that fewer than 4% of fines are successfully contested by plants, a number that is in keeping with our own analysis of the Integrated Compliance and Information System's (ICIS) Federal Enforcement and Case Data.

[16]We also collected data from the Texas Commission on Environmental Quality (TCEQ) on all changes in pollution abatement devices at major air polluters during our time frame. These data confirm that our measure of investment matches well with observed changes in abatement technology. Details of the TCEQ data and its relationship to the ECHO data as well as additional details of exactly how we use the *Actions* database are in Appendix A1.

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

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.

Our main analysis data merge together the above four 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 sample 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 sample 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: Summary Statistics on Estimation Sample

Status:	Compliance	Regular violator	High priority violator
Regulator actions:			
Inspection (%)	8.49	23.43	39.61
Fine amount (thousands of $)	0.04	0.61	8.05
	(0.78)	(1.41)	(11.35)
$\mathbb{1}\{\text{Fine}> 0\}$ (%)	0.15	2.75	11.92
Regulatory outcomes:			
Violation (%)	0.32	3.24	9.34
Entrance into HPV status (%)	0.14	1.64	0.00
Plant actions:			
Investment (%)	0.00	4.91	17.50
Investment (from resolution code) (%)	0.00	4.62	16.35
Investment (from PSD permit) (%)	0.00	0.34	0.43
Investment (from HPV exit) (%)	0.00	0.00	0.80
Dropped investment in compliance (%)	0.37	0.00	0.00
Plant / quarter observations	2,252,570	66,992	36,346

Note: authors' calculations based on estimation sample. Standard deviations are reported in parentheses. Regulatory actions and outcomes are based on lagged status. Plant actions are based on current status.

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 (8.5%) than are plants in regular violator status (23.4%) and plants in HPV status (39.6%). Similarly, fines are much higher for violators and even higher for HPVs. Violating plants are more likely to incur further violations. Violating plants are also much more likely to enter

HPV status than are plants in compliance.

We find that investment occurs in 4.9% of quarters when a plant is a violator and in 17.5% 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 0.37% of plant / quarters in compliance, but do not count these as investment, as noted above.

NAAQS Attainment Status Database

In order to understand the level of environmental enforcement a specific plant may face, we make use of data on whether a given county is entirely or partly in non-attainment of a NAAQS standard during our sample period. We consider each pollutant that was covered by the NAAQS during this period. In particular, we use information on non-attainment for 8-hour ozone (1997 and 2008 standards), carbon monoxide (1971 standard), lead (1978 and 2008 standards), PM-10 (1987 standard), and PM-2.5 (1997 and 2006 standards) in each year from the EPA's "Green Book."[18]

The AP3 database

In order to measure the harm of different types of pollution, we use the AP3 database (Muller et al., 2011) for elevated (e.g. smokestack-level rather than ground-level) emissions to get the marginal damages of each of the criteria air pollutants in each county in 2011.[19] The AP3 data comes from an integrated assessment model that explicitly considers the impact of pollution emitted in different locations, and thereby takes into account differences in local populations and underlying pollution levels. We supplement these data with a national estimate of the

[18] We do not include information on the 1979 1-hour ozone standard because it was revoked at the start of our period on June 15, 2005; the 1971 nitrogen dioxide standard because all areas were in attainment as of September 22, 1998; or the 2010 sulfur dioxide standard because the original areas were not designated until October 4, 2013.

[19] While the criteria air pollutants are technically ozone (O_3), particulate matter (PM), carbon monoxide (CO), nitrogen oxides (NO_X), and sulfur dioxide (SO_2), we will also include the damages from volatile organic compounds (VOC) since it is a precursor to ozone and ammonia (NH_3) because it is a precursor to PM2.5.

marginal damages of lead from Zahran et al. (2017).[20]

This database of pollution damages by county serves two purposes in our analysis. First, we use it to construct a measure of the expected gravity of emissions from a plant of a given industry in a given county. Second, we use it to quantify the costs of different regulatory regimes in our counterfactual analysis.

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 sample period.[21] 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 use these data in only two ways. First, in combination with the marginal emission damage data discussed below, we use the NEI data to understand each plant's potential for harm and the extent of any potential deviation from the CAAA standards. Second, we use the NEI data 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. Finally, following the EPA's analysis of the costs and benefits of the CAAA (Environmental Protection Agency, 1997), we focus our use of the NEI data on criteria air pollutants instead of hazardous air pollutants, which are more difficult to measure.

Table 2 provides summary statistics on the reported criteria air pollution damages for our analysis data, by industrial sector. The damages reported in Table 2 are reported in thousands of dollars in order to aggregate across pollutants with very different damages per ton emitted. There is substantial variation in the pollution damages 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.

[20] Zahran et al. (2017) measures the effect of leaded aviation fuel on the level of lead in children's blood and associates this with changes in long-run earnings. This is likely a lower-bound on the marginal damages of lead.

[21] On-Line Appendix A1 provides details of how we merge the NEI database to the ECHO database.

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 (NAICS 21)	687,400	$501	$3,829	$4,789
Utilities (NAICS 22)	112,554	$14,892	$58,630	$77,941
Manufacturing: food, textiles (NAICS 31)	139,826	$642	$2,831	$2,510
Manufacturing: wood, petroleum (NAICS 32)	617,572	$895	$2,800	$5,894
Manufacturing: metal (NAICS 33)	539,000	$319	$1,967	$2,652
Transportation (NAICS 48)	157,326	$416	$1,008	$2,881
Educational services (NAICS 61)	132,209	$785	$1,730	$1,943

Note: table reports summary statistics on total criteria air pollution damages in thousands of dollars per plant / quarter observation in our analysis data, matched to the NEI data based on being in compliance, a regular violator, or a HPV.

Recall that in our dynamic model, 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.[22] 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

[22]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.

periods. Finally, the lack of a current effect of investments motivates our timing assumption that investment is made at the end of each period, after the regulator's actions and regulatory outcomes.

Table 3: Investment and Resolution of Violations

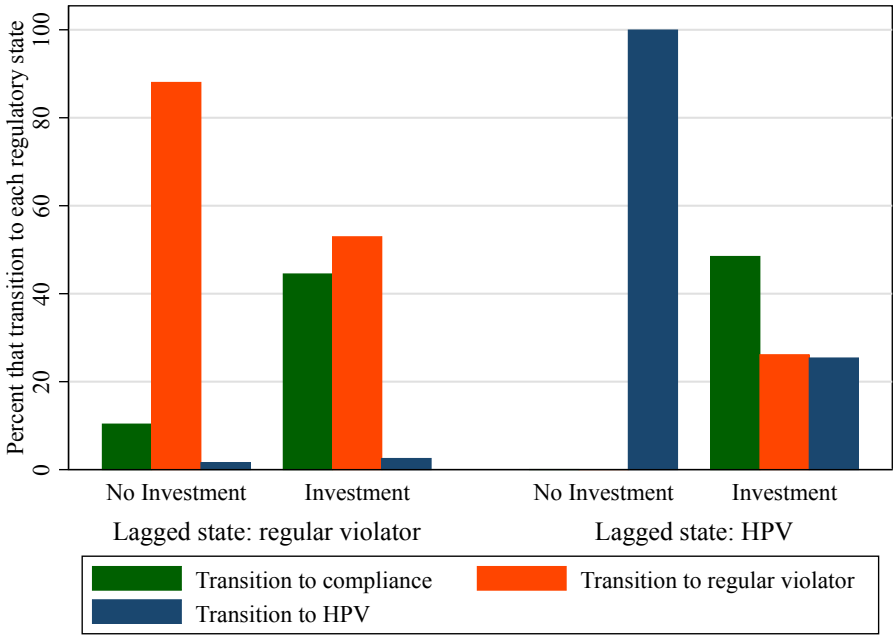
Dependent variable: return to compliance		
Current investment	-0.115^{***}	(0.002)
One quarter lag of investment	0.380^{***}	(0.006)
Two quarters lag of investment	0.083^{***}	(0.007)
Three quarters lag of investment	-0.012^{**}	(0.005)
Four quarters lag of investment	-0.051^{***}	(0.005)
Number of observations	103,338	

Note: regression includes EPA region, NAICS 2-digit industrial sector, and gravity dummies. 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.

Finally, as we discussed in Section 3.1, we assume that any investments that we observe in our data that occur while a plant is in compliance are economic investments (e.g. designed to increase productivity) rather than prophylactic efforts to improve environmental compliance. We do this because regressions of whether a plant transitions out of compliance given recent investment, region, industry, and gravity dummies show that investments in compliance

Figure 2: Effect of Investment on Regulatory State



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

actually *increase* the likelihood that a plant's transitions to both regular and high priority violator in the following two quarters.[23] This result is consistent with the evidence presented in Keohane et al. (2009) that shows that the EPA was more likely to bring lawsuits against plants with recent large investments, a result that they attribute to increased regulatory scrutiny after major investments.

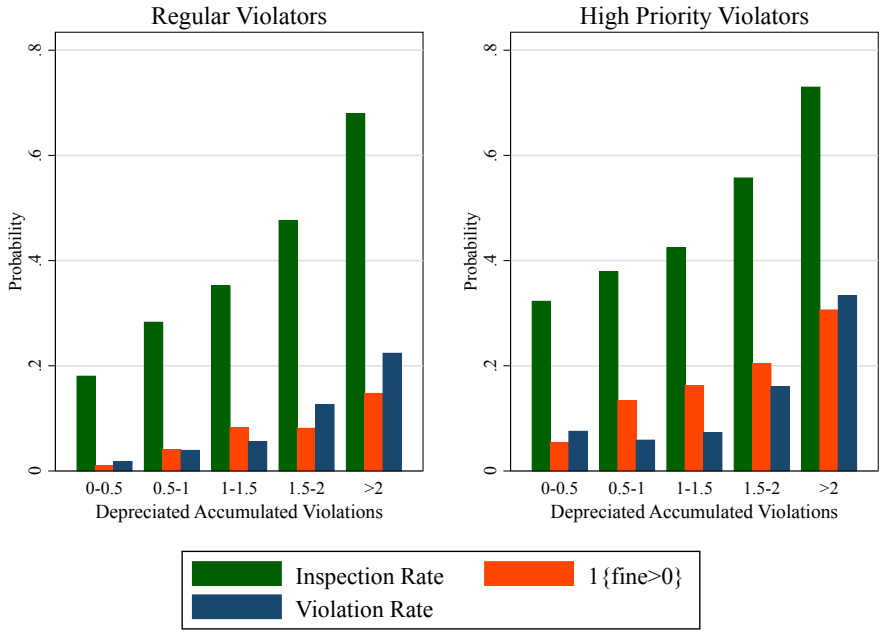
Depreciated Accumulated Violations

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,

[23]Table A1 of On-Line Appendix **??** presents regression coefficients.

from the previous quarter back to the period the plant most recently left compliance.

Figure 3: Depreciated Accumulated Violations and Monitoring and Enforcement



Note: authors' calculations based on estimation sample.

Figure 3 displays the relationship between depreciated accumulated violations (using a 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.[24]

Gravity State

As we discussed in Section 2.1, one of the key components of the EPA's determinants of fines is the gravity of the associated violation. The gravity of a violation is primarily determined by its actual or potential harm, which varies with the pollutants emitted and plant location.

[24]Table A2 in On-Line Appendix A3 justifies a 10% quarterly depreciation rate of accumulated violations relative to other possible depreciation rates.

Gravity is not directly recorded in the ECHO database.[25]

We construct a version of plant-specific expected gravity that aims to capture plants' expectations of the actual and potential harm of a violation as well as the regulatory scrutiny brought about by a plant being in a NAAQS non-attainment area. We focus on the idea that the distribution of pollution across plants in an industry forms the basis of expectations about pollution quantities, both in terms of the mean amount of pollution and the extreme level of pollution if it were an outlier in its industry.

For a given plant in a given county, we therefore take every plant in the same industry nationally, and use the NEI pollution database and the AP3 damages database to calculate the damages from criteria air pollutants if each of those plants were located in this county. From this distribution, we take the mean of this distribution as the plant's expected actual damages of a violation and the 90th percentile of this distribution to be the expected potential damages of a violation. We then combine this information with the NAAQS non-attainment database to sort plants into five gravity bins: below and above the national median for actual and potential damages, further splitting those above the median in both categories into attainment and non-attainment status during our sample period.[26] Table A3 in On-Line Appendix A3 provides summary statistics on gravity in our sample.

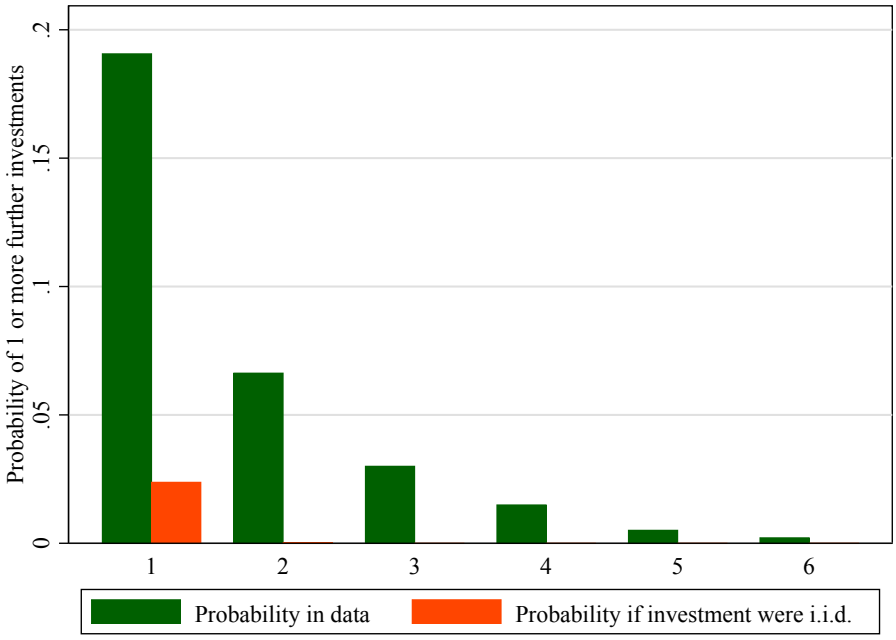
Heterogeneity in Regulatory Environment and Costs

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 of being in HPV status relative to violator status in terms of inspections (the ratio between

[25]While the data do include the pollutant implicated in the violation, this variable is only reported for 14.6% of violations, because this data field does not fall under the federal Minimum Data Requirements of what must be reported to the EPA for every plant. Further, when the "pollutant" is reported, it is often a generic entry such as "facility-wide permit violations" (conditional on an entry, 34.8% of pollutants list this code).

[26]Our measure of non-attainment is whether a county is in non-attainment for any pollutant during any year of our sample period. In our data, 87% of counties are either fully in attainment or out of attainment for at least one pollutant in every year of our sample period and only 1.6% are out of attainment for a minority of years. Hence, treating all counties that are not always in attainment equivalently is a reasonable approximation that allows us to avoid making assumptions about plants' expectations regarding transitions in and out of non-attainment status.

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.

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.[27]

[27]We calculated this figure by taking the investment rate in the data, which is 0.00404, and assuming that each period, each plant invested with this probability. We then calculated the expected number of investments under this scenario.

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 Empirical Framework

4.1 Estimable Model

Our data are at the level of the plant / quarter and include a panel of plants i observed over (quarterly) time periods t. For each plant / quarter, we observe the regulatory state at the point where the plant makes its investment decision—which is $\tilde{\Omega}_{it}$—and the investment decision, X_{it}.

We specialize our model developed in Section 2.2 to these data. Given identification challenges, 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 plants' utility functions. Following the evidence in Section **??**, we let the regulatory state Ω have six 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, (5) the two-digit NAICS industrial sector, and (6) the gravity of the violation, as measured by county non-attainment status and potential environmental damages for plants based on county and industrial sector. Table A3 in On-Line Appendix A3 provides summary statistics on gravity for our estimation sample.

Recall that Ω and e together determine the distribution of enforcement actions and regulatory state transitions. We make the following assumption about e.

Assumption 1. *The environmental compliance signal at period 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 current period violations, Vio, and state transitions, T.[28]

In our model, the regulator chooses stage-contingent inspection policies and state- and signal-contingent fine amounts. Note that Assumption 1 allows for the distribution of e_t to depend on the state-contingent inspection policies. This limits our counterfactuals to ones where inspection policies are the same as in the data. Since we do not observe e, it would be difficult to model how changes in inspection probabilities would affect the resulting distribution of e and through that, violations and transitions. We do evaluate counterfactuals with different fine policies (as well as ones with different values of plants' structural parameters) because the fine policy is a function of e but does not affect transitions conditional on the regulatory state (which includes lagged investments).

We now exposit the plant's utility function from regulatory actions. We let the flow utility for the plant from regulatory actions be:

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

where $\theta^I, \theta^F, \theta^V$, and θ^H are parameters. Note that (1) implies that plants can have a cost from not only fines, but also inspections, additional violations, and being an HPV (consistent with the evidence in Section 2.1), though not from regular violator status.

Recall that once the pollution signal is revealed and regulatory actions are complete, the state at this point is $\tilde{\Omega}$, and the plant can invest if it is not in compliance. The cost of investment is $\theta^X + \varepsilon_{Xt}$. Both ε_{0t} and ε_{1t} are idiosyncratic cost shocks. 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. Plants that are in compliance receive a single shock ε_{0t} and do not make any active decision.

[28] We do not specify the distribution of e_t because we never directly recover it. Instead, we estimate regulator CCPs—which are a function of e_t—and impose a functional form directly on the CCPs.

Group together the structural parameters as $\theta \equiv (\theta^I, \theta^F, \theta^V, \theta^H, \theta^X)$. We generally expect these parameters to be negative, except for θ^X, which we expect to be positive. We assume that θ is fixed for the plant over time. In our random coefficient specifications, θ will vary across plants. In this case, we assume that θ is not contractable, i.e., the regulator cannot choose different enforcement contracts for different plants based on θ.

On-Line Appendix A2 provides the Bellman equations for the plant's dynamic optimization problem.

4.2 Estimation of Regulator CCPs

We estimate the regulator's CCPs and then use them to estimate plants' utility functions. Each CCP is estimated separately for plants in compliance, regular violators, and HPVs (based on their lagged reported status). We start by estimating probit regressions of the probability of an inspection at any state, controlling for two lags of investment; region, industry, and gravity dummies; and depreciated accumulated violations (for plants not in compliance). We estimate similar probit regressions for violations based on the state and whether an inspection occurred. To predict fines, we estimate tobit regressions that further add whether a violation occurred to the regressors in the violation CCP.

Finally, to understand plants' transitions between states, we estimate multinomial logits for $T(\Omega, e)$—which are the transition probabilities from Ω to $\tilde{\Omega}$—that condition on fines and all the regressors in the fines CCP. We examined robustness to further controls such as interactions between region and industry. The results were similar but our bootstrapped datasets lacked variation in some cells because of small numbers of observations, making it difficult to estimate CCPs for them and hence to obtain bootstrapped standard errors to our counterfactuals.

Tables A4, A5, A6, and A7 in On-Line Appendix A3 present marginal effect estimates for these regulatory CCPs. In general, the marginal effects work in the way one would expect: inspections increase the likelihood of violations, violations increase fines and increase the likelihood of transitioning to a more severe regulatory state, and investments increase the

likelihood of a return to compliance.[29]

4.3 Empirical Implementation with Homogeneous Coefficients

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.

We fix $\beta = 0.95^{1/4}$ per quarter. This incorporates both time-discounting at the quarterly rate of 0.0098 and an exogenous probability of exit, which is 0.0031 per quarter in our data.

We estimate θ in 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-log-likelihood of a parameter vector θ is:

$$\log L(\theta) = \sum_i \sum_t \log \left(\left[X_{it} Pr(X = 1|\tilde{\Omega}_{it}, \theta) + (1 - X_{it})(1 - Pr(X = 1|\tilde{\Omega}_{it}, \theta)) \right] \right), \quad (2)$$

where the $Pr(X = 1)$ values are obtained from investment probabilities at the fixed point of the Bellman equation.

We use a nested fixed point estimator similar to Rust (1987). One difference is that in Rust (1987), the state transitions conditional on actions are exogenous, while here, they derive from the regulator's CCPs, making our estimator consistent with a dynamic game.[30] 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.

[29] One important point to note is that investments increase the likelihood of an inspection and inspections increase the likelihood that a high priority violator transitions out of HPV status. This is because major investments often need to be inspected by the regulator within a fixed timeframe, and often a regulator will need to inspect a plant to confirm that the investment has, indeed, returned the plant to compliance.

[30] 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.

4.4 Empirical Implementation 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:

$$\sum_{j=1}^{J} \eta_j = 1 \text{ and } 0 \leq \eta_j \leq 1,\ \forall j. \tag{3}$$

Econometrically, the values of $(\theta_1, \ldots, \theta_J)$ are taken as given.[31] We take a (large) fixed grid of these values, meant to capture the range of plausible parameter values.

We estimate the parameters here by adapting the methods of Fox et al. (2011) and Nevo et al. (2016). Specifically, this framework leads to a computationally quick GMM estimator, allowing us to estimate many parameters, approximating a non-parametric density over the θ utility parameters (Fox et al., 2016).

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:

$$G_k(\eta) = m_k^d - \sum_{j=1}^{J} \eta_j m_k(\theta_j). \tag{4}$$

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

[31] 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 (3). 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 an assumption about the likelihood of seeing each plant in any of its possible states. First, define a division of the state $\tilde{\Omega}$ into $\tilde{\Omega}^1$, which indicates the fixed states of EPA region, industrial sector, and gravity, and $\tilde{\Omega}^2$, which indicates the variable states of compliance status, lagged depreciated accumulated violations, current violation, and lagged investment. 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* $\tilde{\Omega}^2$ *conditional on a given* $\tilde{\Omega}^1$.

Assumption 2 would be valid if, for instance, plants enter at randomly distributed points from the steady state distribution of $\tilde{\Omega}^2$ given $\tilde{\Omega}^1$. It would also occur if they have been active a long time, in which case the distribution of $\tilde{\Omega}^2$ for any θ_j value would approach its steady state distributions. It also rules out a situation where all plants are still adapting to a new regulatory regime.

We compute our specific moments using Assumption 2. Our first set of moments indicates the probabilities of being at a particular time-varying state in equilibrium, conditional on $\tilde{\Omega}^1$. 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.

4.5 Identification

As we 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. If this assumption did not hold, and plants could predict future regulatory actions beyond our CCPs, then this would lead to serially correlated unobserved state variables, rendering a dynamic estimation much more complicated. Our specifications all include fixed effects by EPA region, industrial sector, and gravity 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 costs 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. These expected future fines would depend on the plant's expected future states and the fines at those states, with and without investment. We measure these expectations using the estimated regulator's CCPs and future actions of the plant, highlighting the importance of accurately estimating the CCPs in our context. The plant will then compare its difference in expected discounted fines to the cost of investment, and invest if it receives a benefit from doing so. Conditional on a given fine cost parameter, the investment cost parameter is then identified by the rate at which the probability of investment increases with the increase in the difference in expected discounted fines. Our data will provide variation in the difference in expected discounted fines both for different states within an EPA region, industrial sector, and gravity, and across these fixed states.

Note that our investment variable captures large investments rather than small process investments, since the latter are not available in our data. Our model implicitly captures these process investments through their impact on expected future fines but it does not endogenize them. In other words, it does not allow them to vary in counterfactual policy environments. It is possible that plants invest more in these processes when they are faced with higher marginal fines. In this case, we would understate the importance of dynamic enforcement since we would not capture this response.

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 observables explain much of the variation in investment rates. 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 observables explain little of the variation in investment rates.

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 industrial sector, region, and gravity, we believe that, in practice, variation across these fixed states is very helpful in identifying these parameters. In particular, Figure A1 documents that there is substantial variation in how inspections and fines increase with HPV status across EPA regions.

Our model with random coefficients requires an additional identification argument since we must identify the distribution of values of θ rather than just the mean values of these parameters. If we see some plants repeatedly investing in remediation while other plants

in the same state invest very infrequently, this would suggest that plants have quite different investment costs. More generally, persistence in decisions over time beyond what can be explained by the Markovian structure of the dynamic model with a single θ will identify heterogeneity of types. Thus, with more heterogeneity in plant costs, there will be a higher occurrence of extreme states, e.g., many plants in HPV status with high depreciated accumulated violations concurrent with many plants in compliance.

Our model chooses the level of heterogeneity to match the steady state distribution of states in the data (our first set of moments) and investment rates in those states (our second set of moments) to those predicted by the model. In addition, we directly match the level of investment serial correlation in the data with our third set of moments, which reflects correlations over time in investments for a given plant. As in Figure 4, the greater the correlation here, the more cost heterogeneity we would expect.

5 Results

5.1 Model Estimates

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 4 provides results for both models. The table reports utility parameters as well as the probability that a plant has each of those utility parameters. For the quasi-likelihood model, since there is one set of coefficients, this probability is 1. We report bootstrapped standard errors for the quasi-likelihood model. For the random coefficient estimates, however, we allow the parameter vectors θ to be chosen from a wide grid of potential values. We report the estimated probability, η_j, of observing each of the parameters, θ_j, in the last row of Table 4. We only report those θ_j parameters for which the probability η_j is greater than 1%, and we list the θ_j parameters in descending order of η_j. 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.

Table 4: Estimates of Plants' Structural Parameters

	Quasi-likelihood estimates	GMM random coefficient estimates (1)	(2)	(3)	(4)	(5)	(6)
Negative of investment cost ($-\theta^X$)	-2.872^{***}	-2.334	-1.326	-2.498	-2.541	-1.988	0.153
	(0.037)						
Inspection utility (θ^I)	-0.049	-0.194	0.444	-0.096	0.897	0.001	-2.483
	(0.051)						
Violation utility (θ^V)	-0.077	0.143	0.127	0.650	-0.101	-2.169	-2.006
	(0.195)						
Fine utility (millions \$, θ^F)	-5.980^{***}	-5.181	-6.073	-6.766	-8.460	-7.494	-7.524
	(0.874)						
HPV status utility (θ^H)	-0.065^{***}	-0.029	-0.234	-0.078	-0.411	0.070	-2.437
	(0.014)						
Weight on parameter vector	1	0.439	0.178	0.167	0.128	0.047	0.018

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 all parameter vectors with weight $\eta_j > 0.01$.

We start with the quasi-likelihood results, which are on the left of Table 4. 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.[32] 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.

We next discuss the relative magnitude of our coefficients. The cost of an investment is equivalent to about \$480,000 (2.872/5.980 multiplied by \$1 million) in fines, while one quarter in HPV status is equivalent to a \$10,900 fine. This direct HPV cost is in addition to the impact of HPV status on fines, inspections, and violations, which are also costly to plants. Finally, though not statistically significant, the point estimates suggest that each inspection is equivalent to a fine of \$8,200, and each violation is equivalent to a \$12,900 fine.

While it is straightforward to discuss the relative magnitude of our coefficients, understanding their absolute magnitude is complicated by the fact that fines may be costly to a plant beyond just the amount assessed by the EPA. Resolving fines likely involves additional

[32]Note that investment is represented as a cost, but regulator actions are represented as utility *benefits*, so negative coefficients on regulator actions imply that these actions are costly to plants.

legal work for the plant and may involve reputational costs to plants beyond the EPA (as Evans (2016) suggests HPV status does). In this case, the absolute magnitude of the other coefficients would be higher than their cost relative to a dollar of fines.

One way to evaluate the potential scale of our coefficient estimates is to compare our estimate of investment cost to estimates from the literature on the cost to plants of pollution abatement capital expenditures. Becker (2005) uses the U.S. Census Bureau's Pollution Abatement Costs and Expenditures (PACE) survey to get estimates of average air pollution abatement capital expenditures per plant given non-zero outlays. Using 2007 dollars, he finds that these expenditures average $1.1 million and argues that these are an understatement of the true cost because regulatory compliance may necessitate production process changes that are costly and because the PACE survey does not include the cost of permits or sacrificed output. Dividing $1.1 million by our $480,000 estimate of the investment costs relative to fines suggests that the true cost to a plant from the imposition of a dollar of fines may be more than $2. Or, in other words, more than half of the cost of fines may be from non-monetary costs.

We next turn to the GMM random coefficients estimates, which are in the remaining columns of Table 4. Our GMM specification estimates that six values of θ account for nearly 98% of plants.[33] Nearly half (44%) of the weight is on a set of coefficients that are similar (but not identical) to the quasi-likelihood coefficients. In particular, for plants in this group, investments are equivalent to a $450,000 fine, HPV status is equivalent to a $5,600 fine per quarter, and each inspection is equivalent to a $37,400 fine. Violations are estimated to actually increase utility slightly, which means that for these plants, violations do not themselves lower utility, although we do find that they are correlated with transitions to HPV status.

Interestingly, the second most common set of coefficients, with 18% weight, has much lower investment costs (equivalent to a $218,300 fine) and higher HPV costs (equivalent to a

[33]Our small number of values 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 and find few grid values with positive weights. It is also consistent with the constrained linear optimization that solves for the weights—where the η_j values must be between 0 and 1 and sum to 1—being similar to a LASSO estimator.

$38,500 fine per quarter). These plants find inspections beneficial.[34] In fact, across the five most common coefficient estimates, which represent nearly 96% of plants, the plants with the highest HPV costs and lowest investment costs are the ones that find inspections beneficial.

Column (6) shows that 1.8% of plants have a small but negative mean cost (or benefit) of investment (equivalent to a $20,300 fine per investment). Note that these plants have extremely high costs of inspections (equivalent to a $330,000 fine), violations (a $266,600 fine), and HPV status (a $323,900 fine per quarter), and may be very adverse to environmental enforcement activities relative to investment.

For the first five columns, the GMM investment costs relative to fine costs range from $218,000$ to $450,000$. This is much smaller than the range in other regulatory enforcement coefficients relative to their means. For instance, HPV costs relative to fine costs range from $-9,300$ to $48,600$ per quarter. Thus, the GMM coefficients suggest that there is more heterogeneity in plants' HPV, inspection, and violation cost than there is in plants' investment cost.

Finally, to evaluate model fit, 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.[35]

5.2 Counterfactuals

Using the coefficient estimates from Table 4, 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 regula-

[34] This is in keeping with Duflo et al. (2018), who find that inspections can be beneficial to plants.

[35] We have also estimated our two models for a single industrial sector, mining, allowing the "industry" dummies to now represent 4-digit NAICS subsectors within mining (e.g. coal mining vs metal ore mining). We chose this sector because the sub-industries are fairly homogeneous relative to some other industries in both their production activities and the EPA programs for which we observe violations in the data. Results for the single sector are quite similar to those presented here.

tor'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, our counterfactuals change the state-contingent fine policy and/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 dynamic optimization by the plant when faced with a different regulatory policy.

We conduct two sets of counterfactual experiments. Our first set evaluates the value of dynamic enforcement. Here, we first look at how regulatory states, pollution, and investment would change if the regulator fined all plants in regular and high priority violator status identically, but total equilibrium fines accessed within each region, industry, and gravity state remained the same as the baseline.[36] We then compare this to a counterfactual where the fines for plants in HPV status are doubled, thereby *increasing* the rate at which fines escalate with regulatory status. Finally, we conduct a related counterfactual experiment where we remove dynamics, but solve for the level of fines where total pollution damages are the same as in the baseline. This gives us a sense of how high fines would need to be to make a static fine policy with no escalation mechanism achieve the same cost of pollution as the current, dynamic policy.[37]

Our second set aims to understand how escalation mechanisms relate to the efficient policy of charging each plant for the damages of its pollution, much like a Pigouvian tax (Pigou, 1947). Because we believe that some of the cost to plants of fines is non-monetary, we conduct this experiment in three ways: (1) where the fine cost to plants is entirely monetary, so the efficient fine is the full damages, (2) where the fine cost to plants is twice the imposed fine (as might be inferred from Becker (2005)), so the efficient fine is half of the damages, and (3) where the fine cost to plants is 10 times the imposed fine, so the efficient fine is one-tenth of the damages.

[36]For this counterfactual, we assume that the regulator never fines plants when they are in compliance, and we set the cost of HPV status to zero to fully remove dynamic enforcement.

[37]For this counterfactual, we make the same assumptions as for the second counterfactual.

Table 5 presents the result of our first set of counterfactual experiments, those focusing on the value of dynamic enforcement. 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 pollution damages. The top panel of Table 5 reports the counterfactual results for our quasi-likelihood model while the bottom panel reports results for the random coefficient model.

Table 5: Counterfactual Results: Changing the Escalation Rate of Fines

	(1) Baseline	(2) Same fines for all regular and high priority violators; total fines constant	(3) Fines for HPVs doubled relative to baseline	(4) Same fines for all regular and high priority violators; total pollution constant
Quasi-likelihood estimates				
Compliance (%)	94.66 (0.07)	91.45 (0.12)	95.06 (0.07)	94.81 (0.08)
Regular violator (%)	3.91 (0.05)	3.78 (0.07)	3.91 (0.04)	3.49 (0.01)
HPV (%)	1.43 (0.02)	4.77 (0.05)	1.03 (0.02)	1.70 (0.09)
Investment rate (%)	0.44 (0.00)	0.43 (0.01)	0.45 (0.02)	0.51 (0.01)
Inspection rate (%)	9.43 (0.10)	10.60 (0.06)	9.31 (0.12)	9.52 (0.12)
Mean fines (1000$)	0.32 (0.00)	0.32 (0.01)	0.38 (0.02)	1.51 (0.06)
Mean violations (%)	0.54 (0.00)	1.08 (0.01)	0.50 (0.02)	0.60 (0.00)
Mean plant utility	−0.007 (0.000)	−0.003 (0.000)	−0.008 (0.000)	−0.013 (0.000)
Mean pollution cost (1000$)	1,541.9 (171.3)	1,866.7 (222.7)	1,499.6 (127.9)	1,541.9 (217.5)
GMM random coefficient estimates				
Compliance (%)	92.85 (0.05)	64.71 (0.33)	93.23 (0.02)	92.22 (0.20)
Regular violator (%)	3.42 (0.05)	2.48 (0.05)	3.42 (0.02)	2.70 (0.01)
HPV (%)	1.33 (0.00)	30.40 (0.39)	0.95 (0.00)	2.67 (0.19)
Investment rate (%)	0.52 (0.00)	0.45 (0.03)	0.53 (0.00)	0.63 (0.01)
Inspection rate (%)	9.16 (0.08)	20.19 (0.25)	9.05 (0.03)	9.63 (0.07)
Mean fines (1000$)	0.30 (0.00)	0.30 (0.02)	0.35 (0.00)	1.90 (0.01)
Mean violations (%)	0.52 (0.00)	4.93 (0.22)	0.47 (0.00)	0.71 (0.01)
Mean plant utility	0.006 (0.000)	0.075 (0.000)	0.005 (0.000)	0.001 (0.000)
Mean pollution cost (1000$)	1,487.2 (236.5)	3,976.5 (17.8)	1,447.4 (44.0)	1,487.2 (346.3)

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. Pollution costs are in dollars per quarter. Bootstrapped standard errors are in parentheses.

Column (1) of Table 5 reports the baseline, which is calculated at the estimated parameters. We find similar results here across both specifications, with 1.4% of plants in HPV

status in the long run under our homogeneous coefficients estimates and 1.3% 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 0.44% of periods instead of 0.52% of periods with the quasi-likelihood estimates.

Column (2) of Table 5 reports the long-run outcomes when the escalation mechanism is removed but equilibrium region-industry-gravity fines remain constant. Both the quasi-likelihood and random coefficient estimates show that this leads to large increases in the share of plants in HPV status and a large increase in pollution. In particular, continuing to focus on the quasi-likelihood estimates, we find that the percent of plants in HPV status would increase from 1.3% to 4.8%, while the random coefficient estimates suggest that 30.4% of plants would be in HPV status. All of this change in the share of plants in HPV status comes from a reduction in the share of plants in compliance, as the share of plants in regular violator status actually falls. 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 an increase in regulator workload from a higher inspection rate (20.2% of periods rather than 9.2% for the random coefficient model) and violation rate (4.9% of periods rather than 0.5%). However, we observe only a small decrease in the investment rate (from 5.2% to 4.5% of periods for the random coefficients model), which suggests that plants are investing at different points in time.

Finally, given the much higher level of plants in HPV status, we also find much higher levels of air pollution. Specifically, with the random coefficient estimates, damages from criteria air pollutants rise from \$1.5 million per plant / quarter to \$4.0 million per plant quarter. We take this as strong evidence that dynamic fines are effective in lowering pollution, even conditioning on a mean equilibrium level of fines.

Column (3) of Table 5 *increases* the rate at which fines escalate with regulatory state by doubling the fines for plants in HPV status. As might be expected, this change decreases the share of plants that are HPV status to 1% from 1.3%, while simultaneously decreasing the inspection and violation rates slightly and increasing the investment rate slightly. With this fine policy, average pollution damages drop from \$1.49 million to \$1.45 million per plant /

quarter. We take this as evidence that while there is some benefit to increasing the rate at which fines escalate with regulatory status, this benefit is limited.

Finally, column (4) of Table 5 removes the escalation of fines with regulatory state, but holds pollution damages constant by allowing fines to vary. In this case, there is a slightly higher share of plants in HPV status (2.7% vs 1.3%) with a related slight increase in the inspection and violation rates and a slight decrease in the investment rate. What is striking, however, is how mean fines increase by 634%. To the extent that regulators bear costs from imposing fines, this result shows just how expensive it would be for the regulator to have a fine policy that does not escalate across regulatory states.

Table 6 presents the results of our second set of counterfactual experiments, where we explicitly look at how dynamic enforcement adds value by increasing marginal deterrence for plants in regular violator status (at the expense of under-penalizing these plants while they are in regular violator status). Importantly, actual fines are approximately 13 times higher in high priority violator status than in regular violator status (as shown in Figure 1 and Table 1), while damages are only 2.1 times higher (the mean from Table 2). This means that the existing fine policy escalates much more steeply with regulatory state to induce greater marginal deterrence among one-time violators than would a Pigouvian fine policy.

Column (1) of Table 6 repeats our baseline results, while column (2) shows how long-run outcomes would change if the regulator charged plants in compliance no fines, but charged plants in regular and high priority violator status their full marginal damages. The first thing to note is that these fines are extremely large: 180 times (18,000%) higher than in the baseline at $53,910 per plant / quarter. Even with this massive increase in fines, the share of plants in HPV status actually increases from 1.3% to 1.7%. Importantly, the share of plants in regular violator status drops substantially, from 3.4% to 1.6%. This is exactly what we would expect from both the theoretical literature on escalation mechanisms (Mookherjee and Png, 1994) and our theoretical model: the dynamic enforcement policy will lead to more one-time violations in order to achieve a reduction in repeated violations. This effect is also reflected in the fact that the Pigouvian fines lead to a 13.5% reduction in pollution damages, so the dynamic enforcement approach leads to inefficiently high pollution if it were costly to

impose fines. The benefit, of course, is that a regulator faced with a cost of imposing fines needs to impose substantially lower fines to achieve this level of pollution.

Table 6: Counterfactual Results: Pigouvian Fines

	(1) Baseline	(2) Fines equal to damages for violators	(3) Fines equal to 1/2 damages for violators	(4) Fines equal to 1/10 damages for violators
Quasi-likelihood estimates				
Compliance (%)	94.66 (0.05)	97.56 (0.13)	96.99 (0.25)	95.62 (0.23)
Regular violator (%)	3.91 (0.02)	1.77 (0.10)	2.20 (0.03)	2.99 (0.05)
HPV (%)	1.43 (0.03)	0.67 (0.23)	0.81 (0.22)	1.40 (0.18)
Investment rate (%)	0.44 (0.01)	0.83 (0.02)	0.78 (0.00)	0.63 (0.01)
Inspection rate (%)	9.43 (0.09)	9.00 (0.00)	9.06 (0.04)	9.28 (0.12)
Mean fines (1000$)	0.32 (0.01)	55.30 (1.07)	28.17 (4.41)	6.33 (0.42)
Mean violations (%)	0.54 (0.02)	0.42 (0.03)	0.44 (0.02)	0.54 (0.01)
Mean plant utility	−0.007 (0.000)	−0.353 (0.007)	−0.187 (0.026)	−0.048 (0.002)
Mean pollution cost (1000$)	1,541.9 (132.8)	1,318.5 (287.7)	1,319.6 (560.7)	1,326.5 (231.4)
GMM random coefficient estimates				
Compliance (%)	92.85 (0.09)	94.34 (1.54)	93.74 (0.02)	91.21 (0.01)
Regular violator (%)	3.42 (0.06)	1.57 (0.02)	1.88 (0.12)	2.49 (0.09)
HPV (%)	1.33 (0.03)	1.69 (1.56)	1.97 (0.13)	3.89 (0.08)
Investment rate (%)	0.52 (0.01)	0.84 (0.01)	0.80 (0.02)	0.67 (0.03)
Inspection rate (%)	9.16 (0.17)	9.12 (0.52)	9.21 (0.13)	9.86 (0.16)
Mean fines (1000$)	0.30 (0.00)	53.91 (9.20)	27.39 (1.97)	6.21 (0.38)
Mean violations (%)	0.52 (0.01)	0.51 (0.17)	0.54 (0.03)	0.75 (0.05)
Mean plant utility	0.006 (0.000)	−0.339 (0.061)	−0.171 (0.012)	−0.029 (0.003)
Mean pollution cost (1000$)	1,487.2 (432.4)	1,286.8 (118.2)	1,287.6 (141.7)	1,294.9 (203.6)

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. Pollution costs are in dollars per quarter. Bootstrapped standard errors are in parentheses.

Columns (3) and (4) of Table 6 present the counterfactual results if plants' cost of fines is 2 times (column (3)) or 10 times (column (4)) the fine imposed by the regulator. In this case, the efficient fine to charge is either one-half or one-tenth of total damages. The basic idea from column (2) is repeated: Pigouvian fines lead to fewer plants in regular violator status than the baseline, but more plants in high priority violator status and substantially higher fines. Given the results from Becker (2005), we expect that plants' cost of fines is 2-3 times the fine imposed by the regulator, which would suggest that column (3) is the counterfactual that best approximates efficient penalties.

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 escalation 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. Further, we demonstrate empirically the theoretical insight that increasing the extent to which penalties escalate with regulatory state will lead to under-deterrence of first-time violations in order to reduce the rate of repeat violations. In our counterfactual experiments, this increases pollution relative to the efficient outcome, but drastically decreases the fine the regulator must impose.

6 Conclusion

This paper empirically evaluates the value of dynamic enforcement in the context of Clean Air Act Amendments. We build and estimate a dynamic game model of a plant which is faced with a regulator and must choose when to invest in pollution abatement. We estimate a random coefficients model 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 state-dependence in 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 fines for violators that are independent of the plants' state, further increases in the rate fines escalate with regulatory state would not lower pollution much further.

Overall, while we believe that this analysis provides substantial evidence that dynamic enforcement is valuable, our approach is limited in certain ways. Identification of our model relies on a series of assumptions, including that plants' perceptions of regulatory actions match our regulatory conditional choice probabilities. Future research could improve on our approach by modeling regulator decisions for a single state, where more detailed data may make understanding the determinants of the variation in investment costs possible, although this would come at the cost of less underlying variation in enforcement policies. Further, by modeling the regulator using conditional choice probabilities, we give up the ability to vary inspection policies and regulatory state transition functions in our counterfactuals.

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*, 169–217.

Becker, Randy A, "Air Pollution Abatement Costs Under the Clean Air Act: Evidence from the PACE Survey," *Journal of Environmental Economics and Management*, 2005, *50* (1), 144–169.

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*, 2019.

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.

Environmental Protection Agency, "Clean Air Act Stationary Source Civil Penalty Policy," October 1991.

_ , "The Benefits and Costs of the Clean Air Act, 1970 to 1990," October 1997.

_ , "National Strategy for Improving Oversight of State Enforcement Performance," December 2013.

Environmental Protection Agency, Office of Enforcement and Compliance Assurance, "The Timely and Appropriate (T&A) Enforcement Response to High Priority Violations (HPVs)," June 1999.

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*, 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.

Harrington, Winston, "Enforcement Leverage When Penalties are Restricted," *Journal of Public Economics*, 1988, *37* (1), 29–53.

Helland, Eric, "Prosecutorial Discretion at the EPA: Some Evidence on Litigation Strategy," *Journal of Regulatory Economics*, 2001, *19* (3), 271–294.

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

Johnson, Matthew S, "Regulation by Shaming: Deterrence Effects of Publicizing Violations of Workplace Safety and Health Laws," *Unpublished manuscript*, 2016.

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

Keohane, Nathaniel O., Erin T. Mansur, and Andrey Voynov, "Averting Regulatory Enforcement: Evidence from New Source Review," *Journal of Economics & Management Strategy*, March 2009, *18* (1), 75–104.

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.

Leung, Siu Fai, "How to Make the Fine Fit the Corporate Crime? An Analysis of Static and Dynamic Optimal Punishment Theories," *Journal of Public Economics*, 1991, *45* (2), 243–256.

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

Magat, Wesley A and W Kip Viscusi, "Effectiveness of the EPA's Regulatory Enforcement: The Case of Industrial Effluent Standards," *The Journal of Law and Economics*, 1990, *33* (2), 331–360.

Maitra, Pushkar, Russell Smyth, Ingrid Nielsen, Chris Nyland, and Cherrie Zhu, "Firm Compliance with Social Insurance Obligations Where There is a Weak Surveillance and Enforcement Mechanism: Empirical Evidence from Shanghai," *Pacific Economic Review*, 2007, *12* (5), 577–596.

Mookherjee, Dilip and Ivan PL Png, "Marginal Deterrence in Enforcement of Law," *Journal of Political Economy*, 1994, *102* (5), 1039–1066.

Muehlenbachs, Lucija, "A Dynamic Model of Cleanup: Estimating Sunk Costs in Oil and Gas Production," *International Economic Review*, 2015, *56* (1), 155–185.

Muller, Nicholas Z, Robert Mendelsohn, and William Nordhaus, "Environmental Accounting for Pollution in the United States Economy," *American Economic Review*, 2011, *101* (5), 1649–75.

Nadeau, Louis W, "EPA Effectiveness at Reducing the Duration of Plant-Level Noncompliance," *Journal of Environmental Economics and Management*, 1997, *34* (1), 54–78.

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

Pigou, Arthur Cecil, *A Study in Public Finance*, third ed., Macmillan, 1947.

Polinsky, A Mitchell and Steven Shavell, "On Offense History and the Theory of Deterrence," *International Review of Law and Economics*, 1998, *18* (3), 305–324.

Raymond, Mark, "Enforcement Leverage When Penalties are Restricted: A Reconsideration Under Asymmetric Information," *Journal of Public Economics*, 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.

_ **and** _ , "Enforcement and Over-Compliance," *Journal of Environmental Economics and Management*, 2008, *55* (1), 90–105.

Shinkuma, Takayoshi and Shunsuke Managi, "Effectiveness of Policy Against Illegal Disposal of Waste," *Environmental Economics and Policy Studies*, April 2012, *14* (2), 123–145.

Stafford, Sarah L., "The Effect of Punishment on Firm Compliance with Hazardous Waste Regulations," *Journal of Environmental Economics and Management*, September 2002, *44* (2), 290–308.

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.

Zahran, Sammy, Terrence Iverson, Shawn P McElmurry, and Stephan Weiler, "The Effect of Leaded Aviation Gasoline on Blood Lead in Children," *Journal of the Association of Environmental and Resource Economists*, 2017, *4* (2), 575–610.

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).[38]

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.[39]

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

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

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

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 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.[40]

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).[41]

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

[40] 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.

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

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 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,[42] "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.[43] 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.

This 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

[42]Downloaded September 2014.

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

pollution abatement investment.

To verify that our measure of investment does indeed capture investments in pollution abatement capital equipment, we collected additional data from the Texas Commission on Environmental Quality (TCEQ). The TCEQ data provide information on the installation and removal of pollution control devices for all plants covered by Texas Administrative Code, Title 30, Rule 101.10. This regulation applies to plants with the highest emissions, which is a subset of plants in Texas that are regulated by the EPA. The installation of control devices forms a direct marker of an investment, corresponding to our definition.

We matched the Texas data manually to our base data using firm/regulated entity name, city, and address. Although the set of plants that is regulated by this statute is a subset of the set that show up in our EPA data, we are able to match 1,044 out of 2,109 of the EPA plants in Texas to a plant in the TCEQ data. (Note also that not every plant covered by this regulation will have an abatement device.) In all, the TCEQ data contained 1,520 plants with a change in an emissions source or abatement device during our period, so our 1,044 matched observations represent 69% of these. (Note further that the TCEQ data cover more industries than the 7 in our study, but the TCEQ data do not report industry.) Overall, we believe that our match rate is high enough to make meaningful statements regarding the abatement device changes for larger plants in Texas.

We first investigated whether an investment in the EPA dataset correlated with the installation of an abatement device in the TCEQ data. One issue is that the timing of investment in the two datasets is somewhat different. On the one hand, the EPA data record an indirect measure of investment that only appears in the data once the EPA has confirmed that the violation has been resolved and hence we might expect the EPA measure to lag the Texas measure. On the other hand, the Texas measure of investment only occurs after TCEQ has recorded it in their system following a plant visit, which is supposed to occur within a year of the device installation. TCEQ also does not require self-reporting for abatement devices. Thus, the TCEQ measure may lag the EPA measure.

Despite these limitations, we find a strong and significant relationship between the EPA investment measure and the TCEQ abatement device installation measure. Specifically, we

found that 45% of EPA investments have a TCEQ abatement device installation within 4 quarters conditional on the plant being observed in the both datasets (and unconditionally, the figure is 29%). Similarly, a regression of EPA investment on TCEQ abatement device installation gives a coefficient of 0.031 with a t-stat of 16.9.

We also used the TCEQ abatement device measure to figure out whether additional EPA actions should be included in our measure of investment. We identified three groups of actions that could plausibly be added: (1) an indicator for whether a penalty was paid (C3); (2) an indicator for a violation being withdrawn (WD); and (3) indicators for the EPA determining that the plant was no longer deemed to be in violation due to a rule change or to the plant not being subject to the rule (2L, 2M, NM, NN). Overall, we found only 18 of these actions, compared to 1,094 EPA investments for plants in Texas. Of these 18, only 5 had a TCEQ abatement device change within 4 quarters. Thus, we decided not to add these codes to our definition of investment.

Finally, we investigated whether the installation of an abatement device in compliance in the TCEQ data predicted avoidance of violator status. Specifically, we regressed exit from compliance on recent TCEQ abatement device installation, defined as a TCEQ abatement device installation in the current quarters or within the previous four quarters. We find that, similar to EPA investment, TCEQ abatement device installation in compliance actually increases the likelihood of future violator status. Also, as with the EPA investment variable, TCEQ abatement device installation in violator status predicts a return to compliance.

NEI Database Merge

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.[44] 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.

A2 Computational Details

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, $\tilde{V}(\tilde{\Omega})$ 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.[45] We first exposit $V(\Omega)$, the value function at the beginning of the period:

$$V(\Omega) = \sum_{I \in 0,1} \mathcal{I}(\Omega)^I (1 - \mathcal{I}(\Omega))^{1-I} \int \left[U(\Omega, e) + \tilde{V}(T(\Omega, e)) \right] dP(e|I, \Omega), \tag{A1}$$

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 (A1). However, the plant must integrate over the regulator CCPs.

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

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

We now exposit $\tilde{V}(\tilde{\Omega})$:

$$\begin{aligned}
&\tilde{V}(\tilde{\Omega}) = Com(\tilde{\Omega})[\beta V(\tilde{\Omega},\theta)+\varepsilon_0] + (1-Com(\tilde{\Omega}))\times \\
&\int\int \max\{\beta V(\Omega|\tilde{\Omega},X=0)+\varepsilon_0, -\theta^X+\beta V(\Omega|\tilde{\Omega},X=1)+\varepsilon_1\} f(\varepsilon_0)f(\varepsilon_1)d(\varepsilon_0)d(\varepsilon_1) \\
= \quad & Com(\tilde{\Omega})[\beta V(\tilde{\Omega},\theta)+\gamma] + (1-Com(\tilde{\Omega}))\times \qquad \text{(A2)} \\
&[\ln(\exp(\beta V(\Omega|\tilde{\Omega},X=0)) + \exp(-\theta^X+\beta V(\Omega|\tilde{\Omega},X=1)) + \gamma],
\end{aligned}$$

where $f(\cdot)$ is the density of the type 1 extreme value distribution. The first part of (A2) reflects the case of compliance. In this case, the plant transitions to the same state $\tilde{\Omega}$ 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 part of (A2) 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. The transition state, though still not stochastic, is now potentially different than the current state, because lagged investments and depreciated accumulated violations are both updated.

Finally, having defined value functions, we derive the probability of a plant choosing investment given a regulatory state $\tilde{\Omega}$ and its cost and utility parameters θ as:

$$\Pr(X=1|\tilde{\Omega},\theta) = \frac{(1-Com(\tilde{\Omega}))\exp(\theta^X+\beta V(\Omega|\tilde{\Omega},X=1))}{\exp(\theta^X+\beta V(\Omega|\tilde{\Omega},X=1)) + \exp(\beta V(\Omega|\tilde{\Omega},X=0))}. \qquad \text{(A3)}$$

The probability in (A3) will be used in deriving our estimators below. Accordingly, we have exposited it as a function of the structural parameter vector θ.

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 A7.

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 (A2) and (A1), 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 $\tilde{\Omega}$ 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 $\tilde{\Omega}$ to Ω, given in (A2), 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 (A2).

The transition from Ω to $\tilde{\Omega}$, given in (A1), 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 4.4, we have three sets of moments. From (4), 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 $\tilde{\Omega}^1$. Specifically, for any moment $G_k(\eta) = m_k^d - \sum_{j=1}^{J} \eta_j m_k(\theta)$, where $k \in 1, \ldots, K$ references the specific state $\tilde{\omega}^2 \in \tilde{\Omega}^2$ and $\tilde{\omega}^1 \in \tilde{\Omega}^2$, we can write:

$$m_k(\theta_j) = \Pr[\tilde{\Omega}^2 = \tilde{\omega}^2 | \tilde{\Omega}^1 = \tilde{\omega}^1, \theta_j], \tag{A4}$$

and

$$m_k^d = \sum_i \sum_t \frac{\mathbb{1}\{\tilde{\Omega}_{it}^2 = \tilde{\omega}^2, \tilde{\Omega}_{it}^1 = \tilde{\omega}^1\}}{\mathbb{1}\{\tilde{\Omega}_i^1 = \tilde{\omega}^1\}}. \tag{A5}$$

We compute (A4) 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 $\tilde{\Omega}$, 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:

$$m_k(\theta_j) = \Pr[\tilde{\Omega}^2 = \tilde{\omega}^2 | \tilde{\Omega}^1 = \tilde{\omega}^1, \theta_j] \Pr[X = 1 | \tilde{\Omega}, \theta_j], \tag{A6}$$

and

$$m_k^d = \sum_i \sum_t \frac{\mathbb{1}\{\tilde{\Omega}_{it}^2 = \tilde{\omega}^2, \tilde{\Omega}_{it}^1 = \tilde{\omega}^1, X_{it} = 1\}}{\mathbb{1}\{\tilde{\Omega}_i^1 = \tilde{\omega}^1\}}. \tag{A7}$$

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:

$$m_k(\theta_j) = \Pr[\tilde{\Omega}^2 = \tilde{\omega}^2 | \tilde{\Omega}^1 = \tilde{\omega}^1, \theta_j] \times \Pr[X = 1 | \tilde{\Omega}, \theta_j] \times \left(\sum_{s=1}^{6} \Pr[X \ s \text{ periods ahead} = 1 | X = 1, \tilde{\Omega}, \theta_j] \right), \tag{A8}$$

and

$$m_k^d = \sum_i \sum_t \frac{\mathbb{1}\{\tilde{\Omega}_{it}^2 = \tilde{\omega}^2, \tilde{\Omega}_{it}^1 = \tilde{\omega}^1, X_{it} = 1\} \times \left(\sum_{s=1}^{6} X_{i,t+s} \right)}{\mathbb{1}\{\tilde{\Omega}_i^1 = \tilde{\omega}^1\}}. \tag{A9}$$

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 (A8), 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.[46] 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.

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

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:

$$Var(m_k(\theta_Q)) = \frac{E[m_k(\theta_Q)m_k(\theta_Q)] - E[m_k(\theta_Q)]^2}{N_k}, \tag{A10}$$

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:

$$Cov(m_k(\theta_Q), m_l(\theta_Q)) = \frac{E[m_k(\theta_Q)m_l(\theta_Q)] - E[m_k(\theta_Q)]E[m_l(\theta_Q)]}{N_k}. \tag{A11}$$

The first term in (A11) 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 (A11) 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 (4) subject to the constraints in (3), 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 (4), 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.

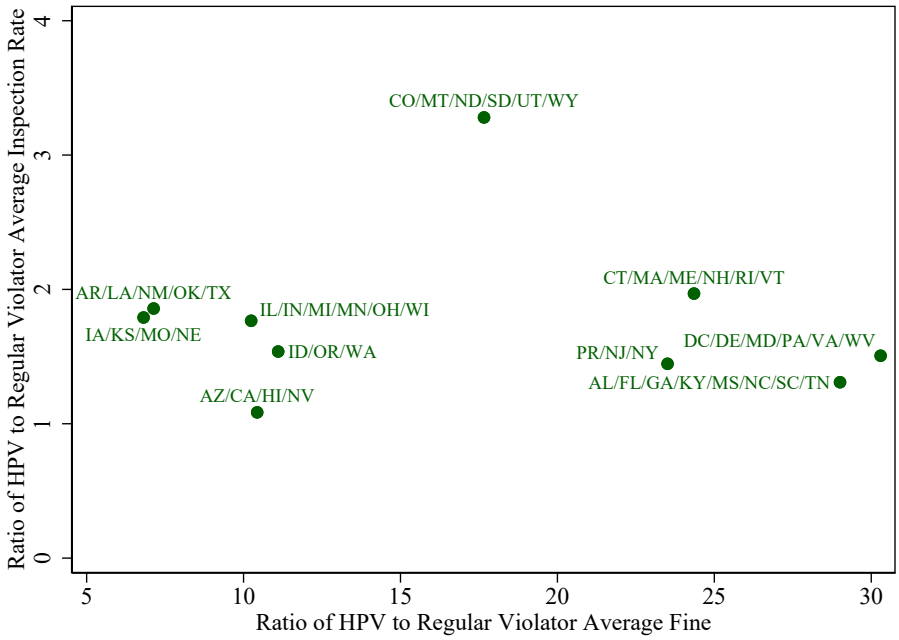
A3 Extra Figures and Tables

Table A1: State Transitions After Investment in Compliance

Outcome: transition to regular violator status		
One quarter lag of investment	1.29***	(.09)
Two quarters lag of investment	1.21***	(.17)
Outcome: transition to HPV status		
One quarter lag of investment	.48***	(.12)
Two quarters lag of investment	1.11***	(.17)

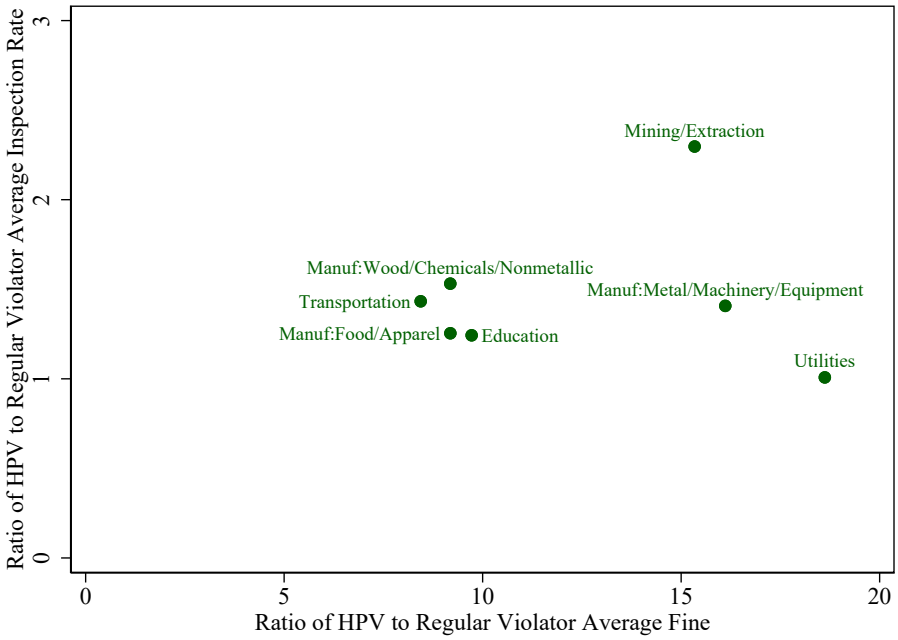
Note: table shows estimates from a multinomial logit regression. Regression includes EPA region, NAICS 2-digit industrial sector, and gravity dummies. Regression uses the estimation sample restricted to plants in compliance in the previous quarter. Standard errors, which are clustered at the plant level, are in parentheses. ***, **, and * indicates statistical significance at the 1%, 5%, and 10% levels, respectively.

Figure 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.

Table A2: Regressions of Regulatory Actions on Depreciated Accumulated Violations

Dependent variable:	Inspection	Fine amount	Violation
Accumulated violations with no depreciation	0.004	-0.014^{***}	-0.000
	(0.007)	(0.004)	(0.001)
Accumulated violations with 10% depreciation	0.132^{***}	0.128^{***}	0.008
	(0.025)	(0.016)	(0.006)
Accumulated violations with 20% depreciation	-0.031	-0.059^{***}	-0.006
	(0.022)	(0.013)	(0.004)
Lagged HPV status	0.115^{***}	0.032^{***}	0.006^{***}
	(0.006)	(0.002)	(0.001)
Number of observations	103,338	103,338	103,338

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.

Table A3: Percent of Observations With Gravity State by Regulatory State

Gravity	Actual damage	Potential damage	NAAQS attainment	In compliance	Regular violator	HPV
1	Low	Low	Either	37.19	36.29	38.98
2	Low	High	Either	2.89	2.44	2.08
3	High	Low	Either	4.07	4.16	3.64
4	High	High	Yes	28.22	29.34	26.58
5	High	High	No	27.63	27.77	28.72
Total:				100	100	100

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

Table A4: Regulatory CCPs Marginal Effects: Inspections

	In compliance	Regular violator	HPV
Regulator actions			
Plant time-varying state			
Lag investment	–	0.050	0.010
2nd lag investment	–	0.100	0.035
Deprec. accum. violations	–	0.126	0.105
Plant fixed state			
Non attainment (when high gravity)	-0.025	-0.022	0.000
High gravity (when attainment)	-0.001	-0.022	-0.017
SE EPA region (versus SW)	-0.120	-0.026	0.042
Industry utilities (versus manuf. food)	0.108	0.193	0.135
Mean	0.086	0.272	0.428
Pseudo R^2	0.093	0.091	0.081

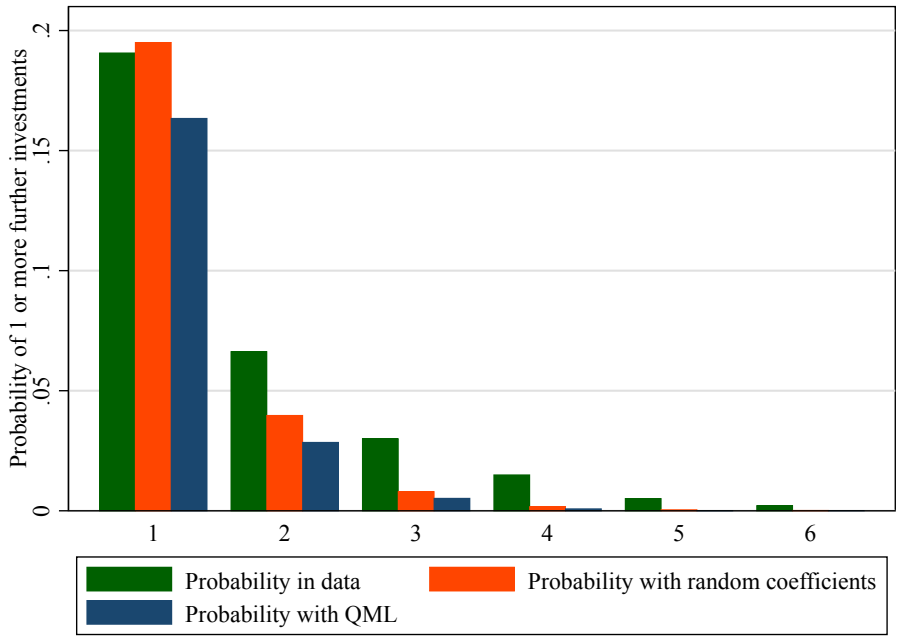
Note: table shows results from probit regressions. Regressions are run separately depending on whether the plant is in each regulatory status (compliance, a regular violator, or high priority violator) at the end of the previous period. Each entry is the marginal effect of having the independent variable move from the mean of each regulatory status to one standard deviation above the mean, with the exception of the plant fixed states, where the marginal effect is calculated as described in the table. All regressions also include EPA region, NAICS 2-digit industrial sector, and gravity dummies.

Table A5: Regulatory CCPs Marginal Effects: Violations

	In compliance	Regular violator	HPV
Regulator actions			
Inspection	0.021	0.063	0.085
Plant time-varying state			
Lag investment	–	-0.007	-0.026
2nd lag investment	–	-0.001	0.029
Deprec. accum. violations	–	0.026	0.041
Plant fixed state			
Non attainment (when high gravity)	0.001	0.001	0.010
High gravity (when attainment)	-0.000	0.006	-0.010
SE EPA region (versus SW)	-0.002	-0.010	-0.026
Industry utilities (versus manuf. food)	-0.001	-0.003	-0.013
Mean	0.000	0.102	0.156
Pseudo R^2	0.185	0.152	0.099

Note: table shows results from probit regressions. Regressions are run separately depending on whether the plant is in each regulatory status (compliance, a regular violator, or high priority violator) at the end of the previous period. Each entry is the marginal effect of having the independent variable move from the mean of each regulatory status to one standard deviation above the mean, with the exception of the plant fixed states, where the marginal effect is calculated as described in the table. All regressions also include EPA region, NAICS 2-digit industrial sector, and gravity dummies.

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 A6: Regulatory CCPs Marginal Effects: Fines

	In compliance	Regular violator	HPV
Regulator actions			
Violation	0.000	0.020	0.289
Inspection	0.000	0.024	0.186
Plant time-varying state			
Lag investment	–	0.002	-0.622
2nd lag investment	–	0.002	0.131
Deprec. accum. violations	–	0.000	0.000
Plant fixed state			
Non attainment (when high gravity)	0.000	0.005	0.196
High gravity (when attainment)	0.000	-0.001	-0.116
SE EPA region (versus SW)	0.000	-0.150	0.134
Industry utilities (versus manuf. food)	0.000	-0.005	0.079
Mean	0.035	0.637	8.268
Pseudo R^2	0.187	0.245	0.114

Note: table shows results from tobit regressions. Regressions are run separately depending on whether the plant is in each regulatory status (compliance, a regular violator, or HPV) at the end of the previous period. Each entry is the marginal effect of having the independent variable move from the mean of each regulatory status to one standard deviation above the mean, with the exception of the plant fixed states, where the marginal effect is calculated as described in the table. All regressions also include EPA region, NAICS 2-digit industrial sector, and gravity dummies.

Table A7: Regulatory CCPs Marginal Effects: Status Transitions

Beginning State:	Compliance		Regular violator		High priority violator	
Transition to:	Into regular violator	Into HPV	Into compliance	Into HPV	Into compliance	Into regular violator
Regulator actions						
Fines						
	0.000	0.000	-0.047	0.000	-0.017	0.000
Violation	0.674	0.000	-0.123	0.000	-0.126	0.000
Inspection	0.006	0.004	-0.006	0.013	-0.013	-0.002
Plant time-varying state						
Lag investment	–	–	0.313	0.000	0.442	0.000
2nd lag investment	–	–	0.137	0.000	-0.042	0.000
Deprec. accum. violations	–	–	0.032	0.000	-0.031	0.000
Plant fixed state						
Non-attainment (given highest gravity)	0.000	0.000	0.004	0.002	0.007	-0.003
Highest gravity and attainment (versus lowest)	-0.000	-0.000	-0.011	-0.000	0.004	-0.003
SE EPA region (versus SW)	0.001	-0.004	0.187	-0.149	-0.013	0.048
Utility sector (versus manuf. food)	-0.000	0.000	-0.010	0.011	0.007	0.003
Mean	0.000		1.000		2.000	
Pseudo R^2	0.5049		0.1760		0.3224	

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