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

Wesley Blundell† Gautam Gowrisankaran‡ Ashley Langer§

February 10, 2020

Abstract

The U.S. Environmental Protection Agency uses a dynamic approach to enforcing air pollution regulations, with repeat offenders subject to high fines and designation as *high priority violators* (HPV). We estimate the value of dynamic 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 raise pollution damages by 164% with constant fines or raise fines by 519% with constant pollution damages.

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

*We thank Dan Ackerberg, Christian Belzil, Esther Duflo, Chase Eck, Liran Einav, Paul Ellickson, Meredith Fowlie, Wayne Gray, Alex Hollingsworth, Karam Kang, Anatolii Kokoza, Derek Lemoine, Bentley MacLeod, Kelli Marquardt, Andy McLennan, Shaun McRae, Ivan Rudik, Nick Ryan, Paul Scott, 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 generous in helping us to understand the details of the data and Clean Air Act enforcement generally. Gowrisankaran acknowledges funding from the National Science Foundation (Grant SES-1824348).

†Washington State University
‡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, since these regulations impact nearly every industrial facility in the U.S., combined enforcement and compliance costs to governments and plants over this period were also large: \$831 billion (Environmental Protection Agency, 1997, converted to 2007 dollars). While the benefits appear to justify the costs, the sheer magnitude of these costs makes it critical to understand the efficiency of regulatory monitoring and enforcement mechanisms for pollution control.

To better understand how environmental regulations are enforced, first consider an example of a large oil refinery in Texas.[1] In 2011, after a period with only low-level violations, the plant was conducting work to improve productive 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 within the facility, 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 and monitoring, after which it 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 history of past actions (Landsberger and Meilijson, 1982; Shimshack, 2014)—to enforce the Clean Air Act Amendments (CAAA). Specifically, 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

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

pollution abatement. Dynamic enforcement may add value when the imposition of fines is costly to the regulator and also when the regulator cannot contract on a plant's compliance costs with its regulatory policies.

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

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

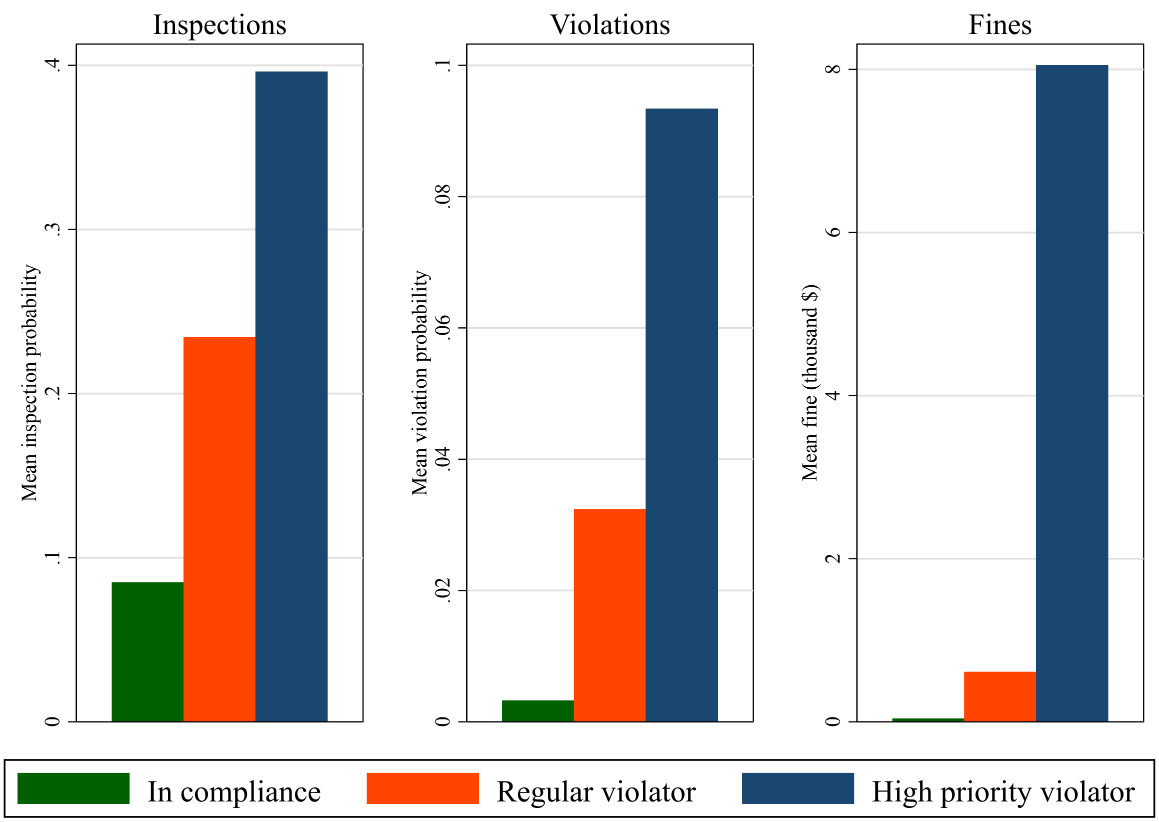

Note: figure reports 2007-13 unconditional mean quarterly levels of inspections, violations, and fines by CAAA regulatory status, based on authors' calculations from the estimation sample.

[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.

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. We then simulate the value of alternative enforcement regimes in affecting plants' emissions and compliance with the CAAA. Our modeling and estimation 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. While the theoretical value of dynamic enforcement is well established, our contribution is to provide evidence on the degree to which this value holds empirically.

In order to measure the value of dynamic enforcement, one needs to account for its benefit in lowering pollution damages and weigh that against the compliance costs to plants and regulators. 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. In our model, the plant and regulator play a discrete-time dynamic game. The regulator makes decisions regarding inspections and fines. Inspections help the regulator obtain more precise information about CAAA compliance. Using its information—including from the inspection when one is performed—the regulator determines whether violations have occurred and decides whether to transition plants to regular or high priority violator status. Both outstanding violations and elevated regulatory status 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, we do not estimate the regulator's utility function, but rather model the regulator's decisions using conditional choice probabilities (CCPs).

Plants decide whether and when to invest in pollution abatement technologies and potentially 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 investing in order to reduce its present discounted value of future regulatory costs. Recovering these costs is key to understanding how plants will respond to counterfactual regulatory policies, such as those that do not condition enforcement activities on plant state.

Our estimation makes use of extensive data with information on virtually all industrial facilities in high polluting industries covered by the CAAA. Our data report inspections, violations, fines, compliance status, and investment decisions for a seven-year-long panel with over 2.3 million plant / quarter observations. These data allow us to estimate plant costs with a non-parametric random coefficients model. We specify a fixed grid of potential cost parameters and estimate the population weights of each. We use a generalized method of moments (GMM) estimator that is computationally very tractable, with a quick and convex optimization problem. Using our estimated cost parameters, we evaluate the gains from dynamic enforcement by computing the pollution damages, assessed fines, and other outcomes when plants optimize under counterfactual regulatory policies.

Relation to literature. This paper relates to three distinct literatures. First, there is an empirical literature on the enforcement of environmental regulations that has largely focused on estimating the relationship between compliance and enforcement.[3] A number of papers also show that dynamic enforcement exists across a variety of contexts.[4] We add to this literature by estimating the value of dynamic enforcement of environmental regulations.

Second, we build on the structural environmental economics literature (e.g., Timmins, 2002; Ryan, 2012; Lim and Yurukoglu, 2015; Muehlenbachs, 2015; Fowlie et al., 2016; Duflo et al., 2018; Houde, 2018; Kang and Silveira, 2018). In particular, Duflo et al. (2018) and Kang and Silveira (2018) estimate regulator preferences in order to evaluate the value of regulator discretion. Duflo et al. (2018) consider a dynamic model of air pollution regulation in India (but do not investigate dynamic enforcement), while Kang and Silveira (2018) consider a static model of water pollution enforcement in California. Though our settings are different,

[3]For instance, Magat and Viscusi (1990), examine whether inspections lower emissions at a plant, Nadeau (1997) uses variation across plant types and states to look at the effect of enforcement on the duration of non-compliance, 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 (2020) examine how variation in the intensity of dynamic enforcement relates to plants' compliance status.

[4]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).

these papers also highlight the value of heterogeneous enforcement.

Third, we use 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.

2 Dynamic Enforcement in Practice and Theory

2.1 Dynamic Enforcement Under the Clean Air Act Amendments

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, the Environmental Protection Agency (EPA) was created 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. The Clean Air Act combines with its amendments 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 (Pb)—as well as various hazardous air pollutants. The CAAA mostly mandate command-and-control regulations, which require that plants' pollution be at or below thresholds that could be achieved with the best technologies and practices.[5] 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.,

[5]The CAAA include some market-based regulations, such as the NO_X cap-and-trade program. However, these regulations incentivize reductions in NO_X emissions beyond the command-and-control requirements. Importantly, plants cannot simply purchase cap-and-trade permits to ensure CAAA compliance.

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 encouraging plants that are out of compliance to return to compliance via plant investments in improved processes or technology.

While the CAAA and EPA dictate the structure of CAAA enforcement, much of the actual enforcement activity is carried out by regional- and state-level environmental protection agencies.[6] 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). Thus, EPA regions and states represent geographic areas across which the interpretation of federal policy and preferences for enforcement may vary.

Under the enforcement system used during our sample period, all plants—in compliance or otherwise—could expect to be inspected regularly. The frequency of these inspections depended not only on baseline differences across states and regions in enforcement budgets 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 area. Non-attainment areas were required to have plans to return to attainment, which could lead to increased levels of scrutiny for plants in these areas.

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

[6] 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).

violations, but also from an increased level of regulatory oversight.

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 duration of the violation. Additionally, gravity is adjusted based on a number of other factors including whether there were reporting issues (e.g., permitting and self-reporting violations), the plant's history of noncompliance, and the plant's ability to pay.[7] 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 and EPA. Thus, there is an advantage to the EPA of obtaining compliance without issuing numerous large penalties.

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

[7]While regulatory enforcement can be tailored to individual plants to some extent via adjustments for ability to pay, enforcement is not allowed to vary based on the EPA's perception of plants' underlying costs. In particular, Environmental Protection Agency (1991) states on p. 22:"... in order to promote equity, the system for penalty assessment must have enough flexibility to account for the unique facts of each case. Yet it still must produce consistent enough results to ensure similarly-situated violators are treated similarly."

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 status 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.

The use of HPV status has been contentious, and the EPA continues to update enforcement policies. During the time frame of our analysis, the EPA used a "watch list" to focus particular attention on HPVs that did not resolve all of their violations in a timely manner. 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), who finds that publicizing non-compliance (in that case for OSHA regulations) can be costly to plants. Further, 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 evaluating the effect of dynamic incentives is 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 (Bentham, 1789; Becker, 1968). We adopt their view that compliance, in our case 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 or enforcement for the regulator (Landsberger and Meilijson, 1982; Harrington, 1988; Leung, 1991; 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 studies consider 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. Each plant plays a dynamic game with the regulator. Our estimation is consistent with the equilibrium of the game being Markov Perfect or with pre-commitment on the part of the regulator.[8] The regulator would like plants to comply with environmental regulations, but also bears a cost from conducting inspections and issuing fines. CAAA violations arise stochastically and plants detect them concurrently with the regulator. Plants make optimizing decisions about whether to invest in remediation of violations. These investments take time and are not always successful in fixing violations. We allow for an escalation mechanism with dynamic

[8]Pre-commitment is very similar to a plant playing against a "regulatory machine" (as modeled by Duflo et al., 2018).

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 are a regulator cost of enforcement; heterogeneous plants; delay and stochasticity in remediation from investment; and imperfect information from inspections.

Each period t corresponds to a quarter and the future is discounted with factor β.[9] Let the *regulatory state* Ω_t 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 $\mathcal{I}(\Omega)$ denote the inspection probability and Ins the actual inspection decision. The regulator and plant then receive a signal e_t, which provides information on the presence and severity of CAAA compliance issues, including emissions from multiple pollutants, plant reporting concerns, and technology maintenance problems.

Specifically, the signal $e_t \equiv (e_t^1, \ldots, e_t^5)$, is the predictor of compliance issues beyond the state. It has five potentially correlated dimensions and a joint distribution that depends on Ins. First, violations depend on e^1, through the function $Vio(\Omega, e^1)$. Second, fines depend on e^2, through $Fine(\Omega, e^2)$. Third, e^3, e^4, and e^5 determine transitions to compliance, regular violator, and HPV status, through $\tilde{\Omega} \equiv T(\Omega, e^3, e^4, e^5)$. In our framework, $\mathcal{I}(\cdot)$ and $Fine(\cdot)$ are policies chosen by the regulator, whereas $Vio(\cdot)$ and $T(\cdot)$ are dictated by e and CAAA standards.

Following the regulator action, 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, violations, fines, designation as a high priority violator, and investment.[10] 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

[9] 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 than plants in India.

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

its inspection and fine policies to minimize the expected weighted sum of damages from pollution, plant investment costs, and enforcement costs.

In order to further illustrate the value of dynamic enforcement, On-Line Appendix A1 develops a simple, special case of this model, that is similar to Polinsky and Shavell (1998). Our simple case highlights how static escalation mechanisms add value by allowing the regulator to increase the marginal deterrence for multiple violations relative to individual violations. Dynamic escalation mechanisms, including the approach the EPA uses to enforce the CAAA, add more value in theory by allowing the regulator to condition on more variables. The remainder of our paper investigates the extent to which this theoretical result holds in practice, by specializing this model to our empirical context.

3 Data and Empirical Foundations

Before we turn to our empirical framework, Section 3.1 describes our data sources and Section 3.2 develops the empirical assumptions that allow us to take our theoretical model to the data.

3.1 Description of Data

Our main analyses principally use four publicly available databases. We summarize our use of the databases here, with details on data construction in On-Line Appendix A2.

Primarily, we use the Environmental Compliance History Online (ECHO) enforcement database. The ECHO database provides plant industry and county, enforcement actions, measures that we use to determine investment, and compliance, regular violator, and HPV status. We infer that a plant has invested if the ECHO data indicate either an environmental issue resolution code or the issuance of a *Prevention of Significant Deterioration* (PSD) permit.[11] Our measure of investment is imperfect in that it only captures large (likely capital) investments rather than smaller investments in improving plant processes that may

[11]We also infer investments for plants that exited HPV status and eliminate investments in compliance (see Section 3.2).

also reduce pollution. To our knowledge, there is no comprehensive national database that contains these types of smaller process investments. We also collected data from the Texas Commission on Environmental Quality (TCEQ) on all changes in pollution abatement devices at major air polluters in Texas during our time frame. These data confirm that our measure of investment matches well with observed changes in abatement technology.

We collapse the ECHO data from the pollution source (AFS ID) level to the plant (FRS number) level using a crosswalk provided by the EPA, and aggregate to the quarter level. We limit our study to the seven most polluting North American Industry Classification System (NAICS) industrial sectors, as listed in Table 2 below. This forms our analysis data, which are at the plant / quarter level and extend from Q1:2007 until Q3:2013.[12]

Table 1: Investment Rates by Regulatory Status

	Compliance	Regular violator	High priority violator
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.

Table 1 summarizes investment rates by regulatory status. Our data contain 2,355,908 plant / quarter observations. As is well-documented in the literature (e.g., Evans, 2016), compliance is high: 95.6 percent of observations indicate 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 an HPV. We derive the vast majority of these investments (94%) 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

[12]The ECHO enforcement actions data start 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. Our seven industries capture 74% of plant / quarters with inspections, 75% of plant /quarters with violations, and 78% of plant / quarters with positive fines during our sample period, among plants that report to ECHO.

status. Finally, we observe codes that are indicative of investment in 0.37% of plant / quarters in compliance, but do not count these as investments.

Not shown in the table, our data cover 107,705 unique plants, of which 66.7 percent are present in every quarter of our sample period. Compliance is also high when considering individual plants: 88.4 percent of plants are never out of compliance, while 7.4 percent of plants have at least one quarter in which they are a regular violator but are never in HPV status. Only 4.2 percent of plants have at least one quarter in which they are in HPV status.

We combine the ECHO enforcement data with three additional datasets. First, the National Emissions Inventory database measures emissions every three years. Our study focuses on emissions of criteria air pollutants (and not hazardous air pollutants) as the data quality for these pollutants is much better (Environmental Protection Agency, 1997). We merge the 2008 and 2011 NEI data from ECHO's Air Emissions Data to our base data using the FRS number and year. We use the NEI data in combination with the AP3 data described next to understand each plant's expected gravity of a violation. Further, we use the NEI data to calculate the mean levels of six pollutants by regulatory state, which are necessary for our counterfactuals.

Second, we use the AP3 database (Clay et al., 2019) for elevated (e.g., smokestack-level rather than ground-level) emissions to get the marginal damages for criteria air pollutants in each county in 2011. We supplement the AP3 data with a national estimate of the marginal damages of lead from Zahran et al. (2017).[13]

Third, the National Ambient Air Quality Standards (NAAQS) database indicates whether a given county is entirely or partly in non-attainment of NAAQS during our sample period. These data enter into our measure of the expected gravity of a violation.

Table 2 provides summary statistics on the reported criteria air pollution damages for our analysis data, by industry. There is substantial variation in the pollution damages across industries. For plants in compliance, the most (least) polluting industry in our data is utilities (educational services). Across industries, plants in regular violator status emit more harmful

[13]Zahran et al. (2017) measure the effect of leaded aviation fuel on the level of lead in children's blood and associate this with changes in long-run earnings. This is likely a lower bound on the marginal damages of lead (Hollingsworth and Rudik, 2019).

Table 2: Summary Statistics on Mean Criteria Air Pollution Levels

Industry	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.

pollutants than plants in compliance. In addition, for most industries, plants in HPV status have higher pollution damages than plants in regular violator status.

3.2 Empirical Foundations of the Estimable Model

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, industrial sector, and expected gravity of violations. In practice, we need to summarize this information for tractability. In this section we provide evidence to motivate our state space and other modeling choices, with further substantiating tables and figures in On-Line Appendix A4.

Investment

We first investigate the role of current and past investment in affecting violator status by regressing whether a plant returns to compliance (from regular or high priority violator status) on current investment, and four quarter lags of investment. We find that investment in the previous quarter is a very strong predictor of a return to compliance, increasing the probability of a return by 38 percentage points. Investment two quarters ago is a weaker, though still statistically significant and positive predictor. In contrast, current investment, and further lags of investment are all negative predictors.[14] Based on these regressions, our state space allows for two lags of investment to affect the regulatory state. We also assume that current investment does not impact a plant's likelihood of returning to compliance in the current period (but can in the subsequent two periods). Finally, the lack of a positive current effect of investment motivates our timing assumption that investment occurs at the end of each period, after the regulator's actions and regulatory outcomes.

Figure 2: Effect of Investment on Regulatory State

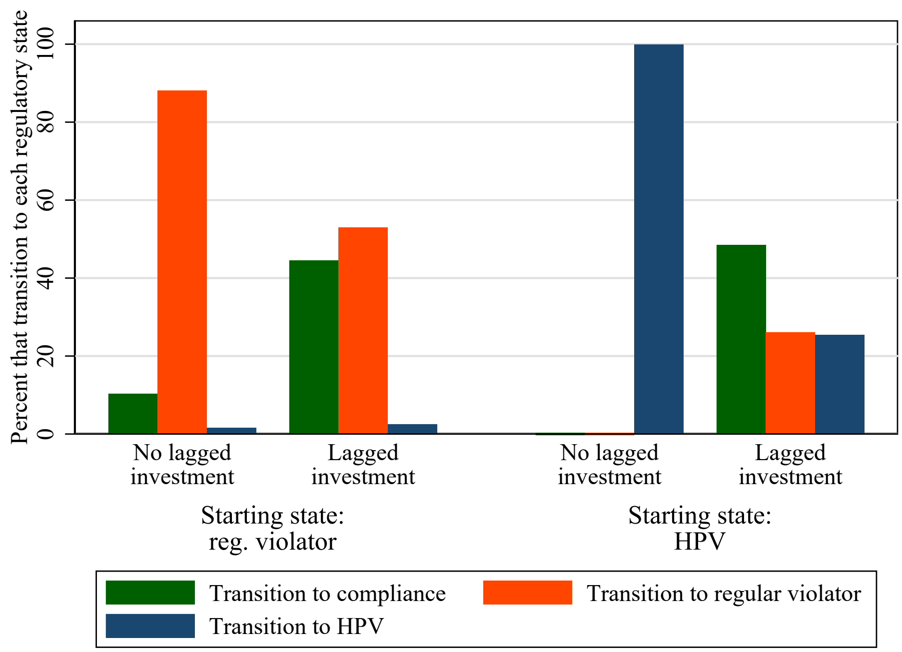

Note: authors' calculations based on estimation sample.

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.

[14] The negative coefficient on current investment may be due to plants in violation investing when additional problems arise.

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 increases the probability that a plant returns to compliance, but does not result in compliance with certainty.

Finally, we consider investments in compliance, to investigate whether these might help prevent future violations. We estimate whether a plant transitions out of compliance given recent investment, region, industry, and gravity state dummies. We find that investments in compliance *increase* the likelihood that a plant transitions to both regular and high priority violator status in the following two quarters.[15] We therefore 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.

Depreciated Accumulated Violations

Figure 1 showed that inspections, violations, and fines all varied substantially based on whether the plant is in compliance, a regular violator, or an HPV. We further investigate whether, within these categories, previous violations are predictive of inspections, violations, and fines. We define a summary measure for plants out of compliance called "depreciated accumulated violations" which is the sum of the depreciated violations from the previous quarter back to the period the plant most recently left compliance. We find that for both regular and high priority violators, depreciated accumulated violations is a strong and positive predictor of inspections, the probability of having a positive fine, and violations.[16] We therefore include depreciated accumulated violations as a state variable that can affect plants' expected regulatory burden.

[15]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 (economic) investments, a result that they attribute to increased regulatory scrutiny after major investments.

[16]We use a 10% quarterly depreciation rate for accumulated violations, as this results in better predictors for these variables than other depreciation rates.

Gravity State

As we discussed in Section 2.1, one of the key components of the EPA's determination 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. Gravity is not directly recorded in the ECHO database.[17]

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.[18] 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 as the expected potential damages of a violation. We then combine this information with the NAAQS non-attainment database to sort plants into five gravity state 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.

[17]While the data do include the pollutant implicated in the violation, this field is only reported for 14.6% of violations because it 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).

[18]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, implying that this is a reasonable and computationally tractable approximation.

Heterogeneity in Regulatory Environment and Costs

Our model allows plants to find both regulatory policies (inspections and fines) costly. We assess how these enforcement decisions vary across different regions and industries by measuring the ratio of fines in HPV status to regular violator status and correlating these ratios with the analogous ratios of inspection rates. We find a correlation of 0.06 (p=0.87) across regions and 0.09 (p=0.86) across industries. These low correlations imply that regions and industries differ in how enforcement escalates with the regulatory state, which will help us identify the costs of fines separately from the costs of inspections.

Finally, our data exhibit substantially more serial correlation in investment than we would expect to occur randomly. 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

This section specializes the model we developed in Section 2.2 to our empirical context, presents our estimation approach, and discusses identification. On-Line Appendix A3 provides additional details.

4.1 Estimable Model

We do not estimate the regulator's utility function. Rather, we specify the regulator's policy function as CCPs (Aguirregabiria and Mira, 2007), and then use the regulator's CCPs to estimate plants' utility functions. Following the evidence in Section 3.2, we let the regulatory state Ω have six components: (1) EPA region, (2) two-digit NAICS industrial sector, (3) expected gravity of potential violations, as measured by county non-attainment status and potential environmental damages for plants based on county and industry, (4) depreciated accumulated violations with a 10% quarterly depreciation rate, (5) regular violator or high priority violator status, and (6) two quarterly lags of investment.

The regulatory state needs to capture all information that affects the distribution of current and future regulatory actions in order for plants to have the same priors on expected regulatory enforcement as our model. Formally, we impose:

Assumption 1. *The environmental compliance signal at period t, e_t, is a function only of the regulatory state Ω_t, the regulator inspection CCPs $\mathcal{I}$, and the inspection decision Ins_t.*

Assumption 1 imposes that e is a function of the regulatory state and the regulator's inspection policy and decision. It rules out the possibility that an investment that is not in the regulatory state (for instance one that occurred many periods ago) could change e. We keep two lags of investment in the regulatory state, and both are allowed to affect the compliance signal.

Assumption 1 is stronger than what we require for estimation: for estimation, we could have directly assumed that plants' have priors on violations, fines, and transitions based only on Ω and the inspection decision, rather than assuming that the underlying signal is a function of Ω and other information. However, Assumption 1 is critical for our counterfactual experiments because it makes explicit how plants' priors will change under different regulatory regimes: it implies that a plant at a given regulatory state Ω_t faced with a given inspection decision and inspection policy will face the same distribution of e—and hence the same distribution of violations and transitions—even under counterfactual fine policies.

Note that e depends on the inspection policy in addition to the regulatory state and the inspection decision. Our conditioning of e on the inspection policy allows the frequency of inspections to affect the expected distribution of signals, which we believe adds to the credibility of our counterfactuals. However, it also implies a limitation of our potential counterfactuals: changing the regulator's inspection policy may change the distribution of e in ways we cannot observe. This limits us to changing the fine policy but not the inspection policy in our counterfactuals.

We let the flow utility for the plant from regulatory actions be:

$$U(\Omega, e) = \theta^I Ins(\Omega) + \theta^V Vio(\Omega, e^1) + \theta^F Fine(\Omega, e^2) + \theta^H HPV(T(\Omega, e^3, e^4, e^5)), \quad (1)$$

where $HPV(\cdot)$ denotes HPV status designation, and $\theta^I, \theta^V, \theta^F$, 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 is $\tilde{\Omega}$, and the plant can invest if it is not in compliance. Our data are at the level of the plant / quarter and include a panel of plants observed over time. For each plant / quarter, we observe the regulatory state at the point where the plant makes its investment decision—which is $\tilde{\Omega}$—and its investment decision. 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.

Group together the structural parameters as $\theta \equiv (\theta^I, \theta^V, \theta^F, \theta^H, \theta^X)$. 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 estimated model, θ will vary across plants. We assume that θ is not contractable, i.e., the regulator cannot choose different enforcement contracts for different plants based on θ.

4.2 Estimation of Regulator CCPs

We estimate plants' expectations of regulator actions, which are Ins, Vio, $Fine$, and T, with CCPs. We specify inspections as a probit of the plant's state Ω. The remaining CCPs are a function of the state Ω, whether an inspection occurred, and the signal e. Econometrically, $(e^1, \ldots, e^5)$ are the residuals in the latent predictors for these CCPs.

We allow for $(e^1, \ldots, e^5)$ to be correlated. Rather then estimating Vio, $Fine$, and T jointly, we estimate the marginal density of Vio, the conditional density of $Fines$ given whether a violation occurred (by including this variable in the regression), and the conditional density of T given the fines assessed and whether a violation occurred. To condition on the state, we estimate the CCPs separately for plants in compliance, regular violators, and HPVs and include indicators for two lags of investment; region, industry, and gravity state dummies; and depreciated accumulated violations (for plants not in compliance). We estimate Vio with a probit, $Fine$ with a tobit, and T with multinomial logits. Our CCPs include interactions of inspection and gravity state except in cases where this led to convergence problems.[19]

4.3 Empirical Implementation of Random Coefficients Model

Our model allows for the parameter vector θ to differ across plants.[20] Specifically, we assume that θ for each plant takes on one of a fixed set of values $(\theta_1, \ldots, \theta_J)$ and that each parameter vector $\theta_j,\ j = 1, \ldots, J$, occurs with probability η_j. Each plant receives a single, independent draw of θ from the multinomial distribution of potential values. The structural parameters that we estimate are therefore $\eta \equiv (\eta_1, \ldots, \eta_J)$ and not $(\theta_1, \ldots, \theta_J)$. We impose no restriction on the structural parameters other than what is necessary based on the fact that they are population probabilities:

$$\sum_{j=1}^{J} \eta_j = 1 \text{ and } 0 \leq \eta_j \leq 1,\ \forall j. \tag{2}$$

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

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

Our GMM estimator has the form $\eta^* = \arg\min_\eta \|G(\eta)\| = \arg\min_\eta G'(\eta)WG(\eta)$, 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

[19]On-Line Appendix A4 provides marginal effects for the CCPs. In general, the results match our expectations.

[20]In addition to our random coefficient model, we estimate a homogeneous coefficients model.

the parametrized model, $m_k(\theta_j)$, where the weights are $\eta_1, \ldots, \eta_J$:

$$G_k(\eta) = m_k^d - \sum_{j=1}^{J} \eta_j m_k(\theta_j). \tag{3}$$

We compute each m_k^d and $m_k(\theta_j)$ in an initial stage, before estimating η. This requires solving the relevant Bellman equation and $m_k(\theta_j)$ for each of the J grid parameters. Using these values, we then estimate η by minimizing $\|G_k(\eta)\|$ subject only to the constraints in (2). We perform a two-step process to improve the efficiency of the 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 region, industry, and gravity state—and $\tilde{\Omega}^2$—which indicates the variable states of compliance status, lagged depreciated accumulated violations, current violation, and lagged investments. 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 rules out a situation where all plants are still adapting to a new regulatory regime.

We compute three sets of specific moments using Assumption 2. Each moment in the first set indicates the equilibrium share of being at a particular time-varying state, conditional on $\tilde{\Omega}^1$. Each moment in the second set indicates the conditional equilibrium share of plants at a particular time-varying state times the share investing at this state. These moments all follow closely from Nevo et al. (2016). Our third set of moments explicitly uses our panel data: each multiplies a second set moment by the corresponding sum of investments in the following six periods. As in Nevo et al. (2016), we obtain inference for our parameters and counterfactuals by bootstrapping, with resampling at the plant level.

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.

4.4 Identification

To understand how the utility parameters θ in our model are identified, consider first a two-parameter version of the homogeneous coefficients model where plants find investment and fines costly but do not face costs from inspections, violations or HPV status and where the idiosyncratic investment cost shocks are zero. In this model, at any violator state, a plant would observe its expected change in discounted future fines conditional on investment. If investment reduced expected discounted future fines by more than the cost of investment, then the plant will invest. Therefore, if the ratio of investment costs to fine costs, $\frac{\theta^X}{-\theta^F}$, was less than the expected change in future fines, the plant would invest. Under this simple model, the parameter ratio is identified from the lowest expected change in future fines at which plants invest.

Conditional on having identified the ratio of the two parameters, we can identify the scale of the parameters by adding in the type 1 extreme value investment cost shocks. The scale is identified by the rate at which the investment probability increases with the expected change in future fines. The steeper is this rate, the larger is this scale.

Our actual model includes five parameters per plant, 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 how investment changes the expected future level of each of these four dimensions. While there is some variation in these changes for different states, $\tilde{\Omega}^2$, within a region, industry, and gravity state, the additional variation across these fixed states, $\tilde{\Omega}^1$, is very helpful in identifying these parameters.

This identification argument hinges on accurately measuring plants' expectations of future regulatory actions with and without investment. We calculate these expected regulatory actions using the estimated regulator's CCPs and future actions of the plant. For these CCPs 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, this would lead to serially correlated unobserved state variables, invalidating our Assumption 1 and requiring very different estimation methods. Our specifications all include fixed effects by region, industry, and gravity state as well as a variety of interactions in order to accurately capture plants' beliefs.

Our random coefficients model 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 some plants repeatedly invest while other plants in the same state invest very infrequently, this would suggest variation in 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. Persistence implies that more heterogeneity will lead to a higher occurrence of extreme states, e.g., many plants in HPV status and many plants in compliance.

We identify the distribution of regulatory costs even though, in our data, the substantial majority of plants never leave compliance. Our model assumes that leaving compliance is not a function of the plant's type, θ. This allows us to identify the distribution of random coefficients based only on the behavior of plants not in compliance.

Our model chooses parameters that most closely match the steady state equilibrium dispersion across states and investment rates in those states to data. We also match the serial correlation in investment in the data with our third set of moments. The greater the correlation here, the more cost heterogeneity we would expect.

Finally, 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. If plants invest more in these processes when they are faced with higher marginal enforcement, we would understate the importance of dynamic enforcement.

5 Results

5.1 Model Estimates

We provide structural parameter estimates in Table 3 for our main model and a homogeneous coefficients specification estimated via quasi-maximum likelihood.[21] 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, and we report bootstrapped standard errors. 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 3. We report the six θ_j parameters with the highest probabilities η_j, 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.

We start with the quasi-likelihood results, which are on the left of Table 3. 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.[22] This is consistent with Duflo et al. (2018), who find that both regulation and investment in pollution abatement are costly to plants.

We next turn to the GMM random coefficients estimates. This specification estimates that six values of θ account for nearly 98% of plants.[23] Nearly half (44%) of the weight is on a set of coefficients that are similar to the quasi-likelihood coefficients. Given that we are estimating utility parameters, we consider the ratios of coefficients. In particular, for plants of this type, investments are equivalent to a \$450,000 fine (2.334/5.181 multiplied by

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

[22]We report the negative of the investment cost, so a negative θ^X implies costly investment.

[23]Heiss et al. (2019) note that this estimator is similar to a LASSO and hence may generate a small number of positive parameters due to an implicit penalization of additional positive parameters.

Table 3: 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	-2.540	-1.988	0.153
	(0.041)						
Inspection utility (θ^I)	-0.049	-0.194	0.444	-0.096	0.897	0.001	-2.483
	(0.049)						
Violation utility (θ^V)	-0.077	0.143	0.128	0.650	-0.100	-2.169	-2.006
	(0.197)						
Fine utility (millions \$, θ^F)	-5.980^{***}	-5.181	-6.073	-6.766	-8.460	-7.494	-7.524
	(1.005)						
HPV status utility (θ^H)	-0.065^{***}	-0.029	-0.234	-0.078	-0.411	0.070	-2.437
	(0.015)						
Weight on parameter vector	1	0.438	0.174	0.170	0.126	0.049	0.019

Note: for the quasi-likelihood approach, we estimate the costs themselves, whereas for the GMM random coefficient approach, we estimate the weights (in the bottom row) on each potential vector of costs. For GMM estimates, we report the 6 parameter vectors with the highest weight. Standard errors for quasi-likelihood estimates, which are bootstrapped with resampling at the plant level, are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

$1 million), HPV status is equivalent to a $5,600 fine per quarter, and each inspection is equivalent to a $37,400 fine. Unlike for the quasi-likelihood estimates, violations increase utility slightly, which means that for these plants, violations do not themselves lower utility, although they do positively correlate with transitions to HPV status.

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 legal work for the plant and harm its reputation more broadly (as Evans (2016) and our estimates suggest HPV status does). This would imply that the cost to a plant of a $1 fine may be substantially larger than $1, which would in turn imply that if an investment is equivalent to $450,000 in fines, then it may actually cost the plant substantially more than $450,000 to invest.

One way to evaluate the potential absolute magnitude of our coefficient estimates is to compare our estimates of investment costs 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. In 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 \$450,000 estimate of the investment costs relative to fines suggests that the true cost to a plant from the imposition of a dollar of fines is \$2.4, with correspondingly higher monetary costs for other regulatory actions. Because Becker (2005) views his estimates as a lower bound on the cost of investment, we assume a \$1 fine is equivalent to \$3 in costs in some counterfactuals.

Interestingly, the second most common set of coefficients, with 17% weight, has much lower investment costs (equivalent to a \$218,300 fine) and higher HPV costs (equivalent to a \$38,500 fine per quarter). These plants find inspections beneficial.[24] In fact, across the five most common coefficient estimates, which represent 95.7% of plants, the plants with the highest HPV costs and lowest investment costs are the ones that find inspections beneficial.

Column (6) shows that 1.9% 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 five coefficients with the most weight, representing 95.7% of plants, the GMM investment costs relative to fine costs range from \$218,000 to \$450,000. This range 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 for these plants. Thus, the GMM coefficients suggest that there is more heterogeneity in plants' HPV, inspection, and violation cost than there is in plants' investment costs.[25]

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

[25] On-Line Appendix A4 provides evidence on model fit and sensitivity checks in which we estimate our model for a single industry, for the ten most populous states (with state fixed effects in the CCPs instead of region fixed effects), and with richer CCPs.

5.2 Counterfactuals

Using the coefficient estimates from Table 3, we now model how EPA enforcement activities, plant investments, overall compliance, and air pollution damages would change under different EPA policies. Because we do not recover regulator preferences, our counterfactuals are based on plant optimization given alternative regulatory policies and do not necessarily stem from the equilibrium of a dynamic game. As we discussed in Section 4.1, we limit our counterfactuals to ones with the same state-contingent inspection policy and only vary the state-contingent fine policy and plants' structural parameters.

We conduct two sets of counterfactual policies. Our first set evaluates the value of dynamic enforcement. Here, we first examine how outcomes would change if the regulator fined all plants in regular and high priority violator status identically for a given region, industry, and gravity state, keeping *total assessed fines* the same as the baseline for each such group. We compare this to a similar counterfactual where the regulator fined all plants in regular and high priority violator status identically for a given region, industry, and gravity state, but where it kept *total pollution damages* the same as the baseline within each group.[26] Finally, we consider a counterfactual where the fines for plants in HPV status are doubled, thereby *increasing* the escalation rate of fines.

Table 4 presents the results of these counterfactual experiments. We report the long-run mean values of regulatory states, regulatory actions, investment rates, plant utility, and pollution damages.

Column (1) of Table 4 reports the observed rates of each outcome in our data, while column (2) reports the baseline, which is calculated at the estimated parameters. In general, our model reproduces the data well: the frequency at which plants are in each regulatory state, the investment, inspection, and violation rates, and the mean pollution damages are similar. Assessed fines are slightly higher in the baseline than in the data.

Column (3) of Table 4 reports the non-dynamic case when equilibrium total fines are the same as in the baseline. We find large increases in the share of plants in HPV status and in

[26] For these counterfactuals, 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.

Table 4: Counterfactual Results: Changing the Escalation Rate of Fines

	(1) Data	(2) Baseline	(3) Same fines for all violators; fines constant	(4) Same fines for all violators; pollution damages constant	(5) Fines for HPVs doubled relative to baseline
Compliance (%)	95.62	95.11 (0.22)	66.72 (13.91)	94.49 (0.62)	95.52 (0.24)
Regular violator (%)	2.88	3.47 (0.25)	2.53 (0.57)	2.72 (0.56)	3.47 (0.26)
HPV (%)	1.50	1.42 (0.05)	30.75 (14.43)	2.79 (0.65)	1.01 (0.03)
Investment rate (%)	0.40	0.54 (0.05)	0.47 (0.06)	0.65 (0.09)	0.55 (0.05)
Inspection rate (%)	9.65	9.41 (0.06)	20.54 (5.41)	9.88 (0.23)	9.28 (0.05)
Fines (thousands $)	0.18	0.32 (0.03)	0.32 (0.03)	1.98 (1.62)	0.36 (0.03)
Violations (%)	0.55	0.54 (0.01)	5.00 (2.20)	0.74 (0.10)	0.49 (0.01)
Plant utility	—	0.006 (0.034)	0.077 (0.091)	0.001 (0.039)	0.005 (0.034)
Pollution damages (mil. $)	1.65	1.53 (0.03)	4.04 (1.19)	1.53 (0.03)	1.48 (0.02)

Note: each statistic is the long-run equilibrium mean, weighting by the number of plants by region, industry, and gravity state in our data. Plant utility reports the average flow utility across types and states including ε except for Euler's constant. Column (1) presents the value of each statistic in our data. Column (2) presents the results of our model given the estimated coefficients and the existing regulatory actions and outcomes. Other columns change the state-contingent fines and the HPV cost faced by plants. Columns (3) and (4) impose the same fines for all regular and high-priority violators for a given fixed state. Column (5) doubles the fines for plants in HPV status. All values are per plant / quarter. Bootstrapped standard errors are in parentheses.

pollution damages. In particular, we find that the share of plants in HPV status would rise from 1.4% to 30.8%. This increase comes mostly from a reduction in the share of plants in compliance and is matched by an increase in regulator workload from a higher inspection rate (from 9.4% to 20.5% of plant / quarters) and violation rate (from 0.5% to 5.0%). However, the investment rate drops only moderately (from 0.54% to 0.47% of periods), suggesting that the heterogeneity in the types of plants that invest and the timing of their investment is important. Finally, given the much higher level of plants in HPV status, we also find much higher levels of air pollution damages. Specifically, damages from criteria air pollutants rise from $1.5 million per plant / quarter to $4.0 million per plant / quarter, an increase of 164%. This provides strong evidence that dynamic fines are effective in lowering pollution damages, conditioning on the fine level.

Column (4) of Table 4 also removes the escalation of fines with regulatory state, but now holds pollution damages within region, industry, and gravity state constant while allowing fines to vary. We find a slightly higher share of plants in HPV status (2.8% vs 1.4%) with a related slight increase in the inspection and violation rates and a slight increase in the investment rate relative to the baseline. What is striking, however, is that mean fines increase by 519%, from \$320 per plant / quarter to \$1,980 per plant / quarter.[27] To the extent that the regulator bears costs from imposing fines, this result shows that the regulator would find it quite costly to have fine policies that do not escalate across regulatory states.

Finally, column (5) of Table 4 doubles the fines for plants in HPV status from their baseline level. This decreases the share of plants in HPV status from 1.42% to 1.01%, while simultaneously decreasing the inspection and violation rates and increasing the investment rate slightly. With this fine policy, average pollution damages drop from \$1.53 million to \$1.48 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.

We replicate these counterfactuals for the quasi-likelihood coefficient estimates in On-Line Appendix A4. The effects of dynamic enforcement are larger under the random coefficients model, demonstrating that heterogeneous plant types (with an inability to contract on plant type) adds to the value of dynamic enforcement.

Our second set of counterfactuals evaluates how escalation mechanisms relate to policies that charge each plant in regular or high priority violator status for its additional pollution damages relative to compliance, much like a Pigouvian tax (Pigou, 1947).[28] Charging plants according to their pollution damages is efficient in a world where the regulator does not care about inspection costs or imposing fines.[29] These Pigou-style policies have two fundamental differences with current EPA policies. First, to increase the marginal deterrence of HPV status, existing fines escalate much more steeply with regulatory state than pollution damages,[30] while Pigou-style policies do not escalate in this way. Second, Pigou-style policies

[27] The 90% confidence interval is [\$1,465, \$6,750], well above the baseline level.

[28] As with the counterfactuals that removed dynamic enforcement, these counterfactuals assume that the regulator never fines plants when they are in compliance and that plants face no direct cost of HPV status.

[29] Note also that the EPA's mandate is not to achieve the efficient level of pollution but rather to enforce the CAAA. Explicitly, the EPA may assess civil and administrative penalties for violations under Section 113(b) of the Clean Air Act Amendments. Since the CAAA set specific definitions of a violation, this enforcement behavior can differ substantially from a Pigouvian tax even apart from a disutility on fines.

[30] Actual fines are approximately 13 times higher in high priority violator status than in regular violator status (Table 1) while damages are only 2.1 times higher (the weighted mean from Table 2).

lower pollution damages by allowing for higher fines for industries that are more polluting. Because we believe that some of the cost to plants of fines could be non-monetary, we conduct this experiment in two ways: (1) where the fine cost to plants is entirely monetary, so the efficient fine is the full damages, and (2) where the fine cost to plants is three times the imposed fine (following our discussion of Becker, 2005), so the efficient fine is one third of the damages. Finally, our third counterfactual escalates fines at the same rate as pollution damages, but scales them to keep aggregate pollution damages the same as the baseline.

Table 5 presents the results of these experiments. Focusing on column (2), Pigouvian fines where the fine cost to plants is entirely monetary are extremely large: 173 times higher than in the baseline at $55,240 per plant / quarter. Even with this massive increase in fines, the share of plants in HPV status actually increases from 1.4% to 1.7%. Importantly, the share of plants in regular violator status drops substantially, from 3.5% to 1.6%. This is consistent with the theory on escalation mechanisms (Mookherjee and Png, 1994): dynamic enforcement "underdeters" one-time violations in order to increase the marginal deterrence for repeat violations. Further, Pigouvian fines lead to a 13.7% reduction in pollution damages (from $1.53 million to $1.32 million per plant / quarter), so the dynamic enforcement approach leads to inefficiently high pollution damages if it were costless for the regulator to impose fines and the fine cost to plants was entirely monetary. Column (3) reports analogous figures where Pigouvian fines are scaled by one-third. It shows similar results to column (2).

Finally, column (4) of Table 5 displays the outcome if we set fines so that they escalate from regular violator to HPV at the same rate as damages, but are scaled so that total pollution damages across all regions, industries, and gravity states is unchanged from the baseline.[31] These results demonstrate the value of dynamic enforcement: with scaled Pigouvian fines, average fines are 394% higher,[32] the share of plants in HPV status is 934% higher, and inspections increase by 51%, relative to the baseline.

In order to evaluate the impact of our counterfactuals across industries, Table 6 shows how four of our counterfactual fine structures affect fines, pollution damages, and regulatory status

[31] In order to recover pollution damages that are the same as the baseline but with fines escalating at the same rate as Pigouvian fines, we divide the Pigouvian fines by 168.

[32] The 90% confidence interval is [$445, $3,950], which is above the baseline level.

Table 5: Counterfactual Results: Scaled Pigouvian Fines

	(1) Baseline	(2) Pigouvian fines	(3) Pigouvian fines scaled by 1/3	(4) Pigouvian fines scaled to yield base pollution damages
Compliance (%)	95.11 (0.22)	96.69 (1.05)	95.38 (1.78)	82.44 (4.60)
Regular violator (%)	3.47 (0.25)	1.60 (0.30)	2.09 (0.30)	2.88 (0.37)
HPV (%)	1.42 (0.05)	1.72 (1.02)	2.52 (1.80)	14.68 (4.89)
Investment rate (%)	0.54 (0.05)	0.86 (0.05)	0.79 (0.06)	0.53 (0.06)
Inspection rate (%)	9.41 (0.06)	9.34 (0.33)	9.60 (0.58)	14.18 (1.72)
Fines (thousands $)	0.32 (0.03)	55.24 (1.81)	19.06 (0.69)	1.58 (1.67)
Violations (%)	0.54 (0.01)	0.52 (0.12)	0.60 (0.21)	2.31 (0.60)
Plant utility	0.006 (0.034)	−0.349 (0.047)	−0.117 (0.038)	0.032 (0.042)
Pollution damages (mil. $)	1.53 (0.03)	1.32 (0.02)	1.32 (0.02)	1.53 (0.03)

Note: each statistic is the long-run equilibrium mean, weighting by the number of plants by region, industry, and gravity state in our data. Plant utility reports the average flow utility across types and states including ε except for Euler's constant. Column (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. All values are per plant / quarter. Bootstrapped standard errors are in parentheses.

for three representative industries: mining, utilities, and metal (and related) manufacturing. Column (1) recreates our baseline results, this time separately for each of the three industries, with the other columns replicating columns (3) and (4) of Table 4 and columns (3) and (4) of Table 5. Focusing on column (2)—which removes escalation, holding fines constant—the increase in HPV status relative to the baseline varies across industries. While the fraction of plants in HPV status increases by a factor of 9 for utilities, it increases more than 20 times for the other two industries. This suggests that there are substantial differences across industries in the gains from dynamic enforcement. Column (3) shows that the increase in fines that is required to hold pollution damages constant without dynamic enforcement also varies substantially across industries. For utilities, average fines only increase by 284%, whereas for mining and extraction, they increased by nearly 11 times their original level (1,094%).

Columns (4) and (5) make clear the benefits of Pigouvian fines: since utilities have substantially higher pollution damages than any other industry, their fines also increase more. Column (5) shows that, when holding pollution damages the same as the baseline, scaled Pigouvian fines reallocate pollution damages from utilities to other industries with lower marginal pollution damages. However, this column also highlights the cost of Pigouvian fines:

Table 6: Counterfactual Results: By Industry

	(1) Baseline	(2) All violators same fines; fines constant	(3) All violators same fines; pollution damages constant	(4) Pigouvian fines scaled by 1/3	(5) Pigouvian fines scaled for base pollution damages
Mining & extraction (NAICS 21)					
Fines (thousands $)	0.17	0.17	2.03	6.10	0.69
Pollution damages (mil. $)	0.58	2.34	0.58	0.53	0.62
Regular violator (%)	4.86	3.71	3.58	3.36	4.16
HPV (%)	0.76	26.23	1.16	1.93	13.81
Utilities (NAICS 22)					
Fines (thousands $)	0.88	0.88	3.38	260.83	5.82
Pollution damages (mil. $)	18.78	41.69	18.78	15.81	16.00
Regular violator (%)	4.11	2.82	3.43	1.68	2.54
HPV (%)	3.93	35.31	5.89	3.51	7.41
Manufacturing: metal (NAICS 33)					
Fines (thousands $)	0.25	0.25	1.51	5.10	1.39
Pollution damages (mil. $)	0.40	1.50	0.40	0.33	0.55
Regular violator (%)	2.58	1.83	2.18	1.50	2.13
HPV (%)	1.48	31.95	2.87	2.64	15.55

Note: each statistic is the long-run equilibrium mean, weighting by the number of plants by region, industry, and gravity state in our data. All columns use the GMM random coefficient estimates. Column (1) presents the results of our model given the estimated coefficients and the existing regulatory actions and outcomes. Other columns change the plants' fines and cost of HPV status. All values are per plant / quarter.

the fine level required to achieve the same total amount of pollution damages is substantially higher than with dynamic enforcement, and this burden falls particularly on utilities,where fines increase by 561%. We take this as suggestive that the EPA finds imposing fines on utilities that are commensurate with their pollution damage levels to be relatively costly.

6 Conclusion

This paper measures the value of dynamic enforcement in the context of the Clean Air Act Amendments. We build and estimate a dynamic model of a plant which is faced with a regulator and must choose when to invest in pollution abatement. We estimate a non-parametric random coefficients specification that is computationally tractable and that allows for wide heterogeneity in plants' costs from regulatory scrutiny.

We find that there are substantial and heterogeneous costs to plants of investing in pollution abatement and of facing regulator enforcement actions, particularly fines and designation as a high priority violator. For 95.7% of plants, the mean investment costs are equivalent to between \$218,000 and \$450,000 in fine costs and the relative heterogeneity in plants' regulatory compliance cost is even larger.

Our counterfactuals yield three main takeaways. First, we find that dynamic enforcement is valuable when fines are costly to the regulator: removing dynamic enforcement would increase pollution damages by 164% if fines were held constant or raise fines by 519% if pollution damages were held constant. These high benefits derive in part from the heterogeneous plant types (and an inability to contract on type). Second, increasing the extent to which fines escalate with the regulatory state would add little additional value: a doubling of fines for plants in HPV status would increase assessed fines by 13% but only lower pollution damages by 3.3%. Third, while scaled Pigouvian fines optimally reallocate enforcement to sectors with high marginal pollution damages—specifically utilities—they do not exploit marginal deterrence. Pigouvian fines scaled to have the same level of pollution damages as in the baseline lead to more plants in HPV status and fewer in regular violator status, which further leads to a 394% increase in assessed fines. Our Pigouvian counterfactuals demonstrate empirically the theoretical point that dynamic enforcement can add value by underdeterring first-time violators relative to repeat offenders, in order to increase marginal deterrence.

While we believe that this analysis provides substantial evidence that dynamic enforcement is valuable, our approach is limited in certain ways. First, we lack detailed pollution data for the majority of observations in our data and can only use more aggregate pollution information. Relatedly, our measure of plant investment in regulatory compliance is imprecise in that it is derived from regulator responses and generally does not include smaller process improvements that may improve plant regulatory compliance. In addition, identification of our model relies on a series of assumptions, including that plants' perceptions of regulatory actions match our regulatory conditional choice probabilities. Finally, 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. Future research could extend our approach by modeling regulator decisions.

Overall, this analysis provides the first empirical estimates of the plants' responses to the dynamic environmental regulations used around the world. Our modeling framework and results on dynamic enforcement for the CAAA may allow 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.

Bentham, Jeremy, "An Introduction to the Principles of Morals and Legislation," *London: Athlone*, 1789.

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," *Journal of Environmental Economics and Management*, 2020.

Clay, Karen, Akshaya Jha, Nicholas Muller, and Randall Walsh, "External Costs of Transporting Petroleum Products: Evidence from Shipments of Crude Oil from North Dakota by Pipelines and Rail.," *Energy Journal*, 2019, *40* (1).

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.

Heiss, Florian, Stephan Hetzenecker, and Maximilian Osterhaus, "Nonparametric estimation of the random coefficients model: An elastic net approach," *arXiv preprint arXiv:1909.08434*, 2019.

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

Hollingsworth, Alex and Ivan Rudik, "The social cost of leaded gasoline: Evidence from regulatory exemptions," 2019.

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.

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 Illustrative Simple Case of our Model

To illustrate the value of dynamic enforcement, we present a simple, special case of our general model. For this special case, we assume a two-period model with $\beta = 1$. In both periods, a single new violation occurs with probability p. Violations are costly to the regulator in that they may result in emissions, but costless to the plant. Inspections occur with probability $\mathcal{I}(\Omega) = 1$, are costless to the regulator and plant, and perfectly reveal the presence of a violation. The regulator assigns the plant to compliance if it has 0 outstanding violations; regular violator status if it has 1 outstanding violation; and HPV status if it has 2 outstanding violations. We assume that the signals for violations and fines are the same: $e_t^1 = e_t^2$, and indicate the number of outstanding violations. Note that $e_1^1 \in \{0, 1\}$ and $e_2^1 \in \{0, 1, 2\}$. The remaining signals follow from the above description of the status transitions, with $e_t^3 = \mathbb{1}\{e_t^1 = 0\}$, $e_t^4 = \mathbb{1}\{e_t^1 = 1\}$, and $e_t^5 = \mathbb{1}\{e_t^1 = 2\}$.

A period 1 investment, $X_1 = 1$, clears a period 1 violation with probability q; violations are never cleared without investment. The pollution cost to the regulator is $c_E e_t^1$ 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 $\tilde{\Omega}_2 = (X_1, e_1^1, e_2^1)$. Finally, the per-period objective function to the plant is $-\theta^X X_t - Fine(\Omega_t, e_t^1)$, where θ^X is the cost of investment. The regulator minimizes the sum over the two periods of $c_E e_t^1$, $\theta^X X$, and its cost of assessing fines.

We allow the regulator to pre-commit to an enforcement strategy and focus on the case with a period 1 violation—so $e_1^1 = 1$—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 $c_F e_t^1$. 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 $e_t^1 = 2$, which would remove the cost of fining when $e_t^1 = 1$ 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. Although the model is dynamic, this escalation mechanism is not explicitly dynamic (since it does not depend on the regulatory state Ω_t, but only on the current number of outstanding violations, e_t^1). It 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 $e_{t-1}^1 > 0$, $X_{t-1} = 0$, 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 of a single θ^X 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 to invest. For simple investment cost type distributions, the regulator will incentivize investment for more values of θ^X with dynamic enforcement than with a static escalation mechanism or with linear fines.

Overall, our illustrative simple case shows that escalation mechanisms add value by increasing the marginal deterrence for two violations relative to one. Dynamic escalation mechanisms add more value by reducing equilibrium fines and by increasing the set of actions over which the regulator can condition.

A2 Data Construction Details

ECHO Database Overview

The ECHO database is divided into a number of components. 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 each plant. Our analyses control for the EPA region, the first two digits of the NAICS code, and the expected gravity of violations based on industry and county. We keep seven industries with high pollution damages that we believe to have plants of broadly comparable costs of investment and enforcement: the three manufacturing industries, mining and extraction, transportation, educational services (which includes school buses), and utilities.

Second, the *Air Facility System Actions* (AFS) dataset (or *Actions* dataset for short) records the history of regulatory actions taken by state, regional, and federal environmental regulators, from Q4:2006 through the Q4:2014.[33] 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. One potential issue with our data is that some states were not reporting non-HPV violations prior to 2010.[34] The EPA customer support staff were not sure if the data had been corrected and suggested we review the data for anomalous changes in the violation rate. We examined the data for changes in the prevalence of violations in 2010 by performing a series of regressions of reported violations on state or region dummies interacted with a dummy for post-2010. We found no systematic evidence of an increase in reported violations, suggesting that this error had been corrected ex post.

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

[34] See `https://echo.epa.gov/system/files/FRVMemoandAppxFinal3.22.10.pdf`.

Each record in this dataset details a regulator 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[35] and the penalty amount when the action is a fine.[36] For each plant, we combine EPA actions across all EPA programs to capture completely its regulatory enforcement status.

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.[37] We use this dataset to determine whether a plant is in compliance or a violator in any quarter. 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.

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 pro-

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

[36]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.

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

gram. 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] 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 these cases, 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.

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

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

[38] Although this list indicates both plants that are regular violators and HPVs, we determined HPV status from the *High Priority Violator History* dataset, for greater accuracy.

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

issues a Section 114 letter for gathering information from the plant. 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).[40] 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.

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 quarter with this status and all subsequent quarters.

Investment

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

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

First, 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,[41] "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."[42]

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." Of the investments that are determined by a resolution code, we observe "VR" for the overwhelming majority (77%). An additional 14% of these investments are from "C7", and the remaining 10% are from "OT."

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 as well as all outstanding violations.

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

[41] Downloaded September 2014.

[42] Note that we do not always observe "VR" or other investment codes when plants return to compliance from regular violator status. Thus, we allow for the possibility that plants can return to compliance from regular violator status without an investment.

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

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. 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 also that not every plant covered by this regulation will have an abatement device and 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 four 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 within four quarters gives a coefficient of 0.031 with a t-stat of 16.9.

We also used the TCEQ abatement device measure to determine 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 quarter 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.

Pollution and Damages Data

National Emissions Inventory data. We match 59% of observations in the ECHO data for 2008 and 2011 to the NEI data. The imperfect match is consistent with other studies that use the NEI data; e.g., Shapiro and Walker (2015) achieve a 77.4% match rate between the NEI and the Census of Manufacturing. We measure smokestack emissions for six pollutants: PM2.5, NO_X, SO_2, volatile organic compounds, NH_3, and Pb. For our counterfactuals, we need the expected level of pollution by regulatory state. To obtain this, we aggregated the matched NEI data to the region, industry, gravity state, and compliance / regular violator / HPV status level. We then calculated the mean pollution for each of these states, imputing missing values. We did not use the full regulatory state here given the limited number of matching observations in the NEI data for some states.

AP3 data. The AP3 data come 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. While we consider the damages from criteria air pollutants—ozone (O_3), particulate matter (PM), carbon monoxide (CO), nitrogen oxides (NO_X), sulfur dioxide (SO_2), and Pb—the AP3 data include damages from smokestack emissions that can lead to criteria air pollutants—PM2.5, NO_X, SO_2, volatile organic compounds (a precursor to ozone), and NH_3 (a precursor to PM).

National Ambient Air Quality Standards attainment data. We consider NAAQS attainment status for each pollutant covered 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." We do not include information on the 1979 1-hour ozone standard because it was revoked 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, after the end of our sample period.

A3 Details on Empirical Framework

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.[44] We first exposit

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

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

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

where $dP(e|\Omega, \mathcal{I}, i)$ is the integral over the density of the environmental compliance signal e given the plant state, the inspection policy, and the inspection decision. 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 policies and e.

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

$$\begin{aligned} \tilde{V}(\tilde{\Omega}) = Com(\tilde{\Omega}) \times \textstyle\int [\beta V(\tilde{\Omega}, \theta) + \varepsilon_0] dF(\varepsilon_0) + (1 - Com(\tilde{\Omega})) \times \\ \textstyle\int \int \max\{\beta V(\Omega|\tilde{\Omega}, X = 0) + \varepsilon_0, -\theta^X + \beta V(\Omega|\tilde{\Omega}, X = 1) + \varepsilon_1\} dF(\varepsilon_0) dF(\varepsilon_1) \\ = Com(\tilde{\Omega})[\beta V(\tilde{\Omega}, \theta) + \gamma] + (1 - Com(\tilde{\Omega})) \times \\ [\ln(\exp(\beta V(\Omega|\tilde{\Omega}, X = 0)) + \exp(-\theta^X + \beta V(\Omega|\tilde{\Omega}, X = 1)) + \gamma], \end{aligned} \tag{A2}$$

where $dF(\cdot)$ is the integral over 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 continuation value plus 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 from the current state, because it updates both lagged investments and depreciated accumulated violations.

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

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

Since the probability in (A3) is used in our estimators, we have written it as a function of the structural parameter vector θ.

Computation of 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 (A1) and (A2). 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 $R^{-1}(0.025)$ to $R^{-1}(0.975)$. 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^{-9} between subsequent iterations. Following Assumption 1, when we compute Bellman equations under counterfactual policy environments, the state-contingent inspection, violation, and transition probabilities remain the same as in the base computations.

Empirical Implementation of Homogeneous Coefficients Model

In addition to our main random coefficient model, we estimate a model with homogeneous coefficients θ using a quasi-likelihood nested fixed point estimator. We calculate a quasi-likelihood (and not a likelihood) because we use the regulator's estimated CCPs in the plant's dynamic optimization process. In this model, there are no serially correlated unobservables for a plant over time, and hence, we can treat each plant i and quarter t as an independent observation. The quasi-log-likelihood of a parameter vector θ is:

$$\log L(\theta) = \sum_i \sum_t \log\left(\left[X_{it} Pr(X=1|\tilde{\Omega}_{it},\theta) + (1-X_{it})(1-Pr(X=1|\tilde{\Omega}_{it},\theta)) \right] \right), \quad \text{(A4)}$$

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

Our nested fixed point estimator is 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.[45] 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.

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

[45] 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.

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 by again following Fox et al. (2016) and using co-prime Halton sequences for each parameter, using the first five prime numbers, since 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 4.3, we have three sets of moments. In order to explain our moments, order the states $1, \ldots, K$ and let ω_k^1 denote the fixed component of state k and ω_k^2 denote the variable component of state k. Then, let $\pi_k(\theta)$ be the steady state share of plants at ω_k^2 given ω_k^1. For a given ω_k^1, we recover the associated $\pi(\theta)$ values by solving the Bellman equation for ω_k^1, generating the transition matrix between variable states, and finding the vector that is invariant when transformed by this matrix.

As in (3), each moment is constructed from some m_k^d and $m_k(\theta_j)$. We now denote these terms m_k^1, m_k^2, and m_k^3, and m_k^{d1}, m_k^{d2}, and m_k^{d3}, corresponding to our three sets of moments. Our first set of moments indicates differences in the steady state share of plants π_k between the model and the data. Specifically, for any moment $G_k(\eta) = m_k^{d1} - \sum_{j=1}^{J} \eta_j m_k^1(\theta)$, we let:

$$m_k^1(\theta_j) = \pi_k(\theta_j), \tag{A5}$$

and

$$m_k^{d1} = \frac{\sum_i \sum_t \mathbb{1}\{\tilde{\Omega}^2_{it} = \omega^2_k, \tilde{\Omega}^1_{it} = \omega^1_k\}}{\sum_i \sum_t \mathbb{1}\{\tilde{\Omega}^1_{it} = \omega^1_k\}}. \tag{A6}$$

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 $\tilde{\Omega}$, this would be difficult in practice given that we have over 50,000 states. Hence, we create moments for the 5,000 states which have the highest expected number of steady state observations at our estimated quasi-likelihood parameters and given our data on $\tilde{\Omega}^1$.

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 steady state share of plants at any variable state times the conditional share having an investment at that state:

$$m_k^2(\theta_j) = \pi_k(\theta_j) \times Share[X = 1|\tilde{\Omega}, \theta_j], \tag{A7}$$

and

$$m_k^{d2} = \frac{\sum_i \sum_t \mathbb{1}\{\tilde{\Omega}^2_{it} = \omega^2_k, \tilde{\Omega}^1_{it} = \omega^1_k, X_{it} = 1\}}{\sum_i \sum_t \mathbb{1}\{\tilde{\Omega}^1_{it} = \omega^1_k\}}. \tag{A8}$$

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 steady state share of plants at any variable state times the conditional share having an investment at that state times the sum from 1 to 6 of the product of the number itself and the conditional share with that many investments in the next six periods:

$$\begin{gathered} m_k^3(\theta_j) = \pi_k(\theta_j) \times Share[X = 1|\tilde{\Omega}, \theta_j] \times \\ \left(\textstyle\sum_{s=1}^6 s \times Share[s \text{ investments within 6 periods}|X = 1, \tilde{\Omega}, \theta_j] \right), \end{gathered} \tag{A9}$$

and

$$m_k^{d3} = \frac{\sum_i \sum_t \left[\mathbb{1}\{\tilde{\Omega}_{it}^2 = \omega_k^2, \tilde{\Omega}_{it}^1 = \omega_k^1, X_{it} = 1\} \times \left(\sum_{s=1}^{6} X_{i,t+s} \right) \right]}{\sum_i \sum_t \mathbb{1}\{\tilde{\Omega}_{it}^1 = \omega_k^1\}}. \tag{A10}$$

These moments seek to match the extent of repeated investments by plants in the data 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 (A9), 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 for our base specification has 14,374 moments, composed of 5,000 of the first set and 4,687 each of the second and third set. Our computation of $m_k(\theta_j)$ results in a 14,374 $\times$ 10,001 matrix and takes approximately eight days on an iMacPro with eight processors, with code written in C with MPI, or two days on the University of Arizona high performance cluster, using 28 processors.

Weighting Matrix and Estimation of GMM Parameters η_j

We follow the standard approach in GMM estimation of weighting by an estimate of the inverse of the variance-covariance matrix to improve the efficiency of our estimates.[47] We proceed in two stages. In stage 1, we estimate the model with a weighting matrix that does

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

[47] Our GMM estimator is non-standard in that it includes the constraints in (2), which limits our ability to prove asymptotic efficiency of this estimator.

not reflect an asymptotic approximation to the variance-covariance matrix. Then, we use our stage 1 estimates to compute an approximation to the variance-covariance matrix.[48] In stage 2, we reestimate our parameters using this 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.[49]

We calculate the diagonal elements of this matrix as:

$$Var(m_k(\theta_Q)) = \frac{E[m_k(\theta_Q)m_k(\theta_Q)] - E[m_k(\theta_Q)]^2}{N_k}, \tag{A11}$$

where N_k is the number of plant / quarter observations from the region, industry, and gravity state 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 with different values of $\tilde{\Omega}^1$. We can write the covariance between moments k and l from the same $\tilde{\Omega}^1$ as:

$$Cov(m_k(\theta_Q), m_l(\theta_Q)) = \frac{E[m_k(\theta_Q)m_l(\theta_Q)] - E[m_k(\theta_Q)]E[m_l(\theta_Q)]}{N_k}. \tag{A12}$$

The first term in (A12) 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 (A12) is simply the product of the means.

In stage 1, we invert and take a Cholesky decomposition of this estimated variance-

[48]We base our approximation on the stage 1 parameters with weights of 0.01 or greater.

[49]For some robustness specifications, we had collinearity issues with inverting this variance-covariance matrix. We dropped moments with zero variance in one specification and used the diagonal of the matrix for another specification.

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 (3) subject to the constraints in (2), via constrained least squares. We use the Matlab package lsqlin to perform this minimization process, which takes approximately 10 minutes on an iMacPro. The process generates consistent estimates of η that we use to construct a weighting matrix.

We then estimate the variance-covariance matrix of $G(\eta)$ using our stage 1 GMM estimates of η. From (3), 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 by 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 those based on the 5,000 states which have the highest expected number of steady state observations at our estimated quasi-likelihood parameter. Note

that the exact number of moments, m_k, varies across iterations of the bootstrapping procedure, depending on how many of those 5,000 states are in compliance.

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 second stage 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 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 and performs the counterfactual computation with the new structural parameters, regulator CCPs, and dataset (analogous to step 6).

A4 Extra Figures and Tables

Table A1: 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 region, industry, and gravity state dummies. Regression uses the estimation sample restricted to plants not in compliance at the start of the period. Standard errors, which are clustered at the plant level, are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Table A2: 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	0.48^{***}	(.12)
Two quarters lag of investment	1.11^{***}	(.17)

Note: table shows estimates from a multinomial logit regression. Regression includes region, industry, and gravity state dummies. Regression uses the estimation sample restricted to plants in compliance at the start of the period. Standard errors, which are clustered at the plant level, are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Table A3: 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)
HPV status at start of period	0.115^{***}	0.032^{***}	0.006^{***}
	(0.006)	(0.002)	(0.001)
Number of observations	103,338	103,338	103,338

Note: regressions include region, industry, and gravity state dummies. Regression uses the estimation sample restricted to plants not in compliance at the start of the period. Standard errors, which are clustered at the plant level, are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Figure A1: Mean Inspection Probabilities and Fines by EPA Region

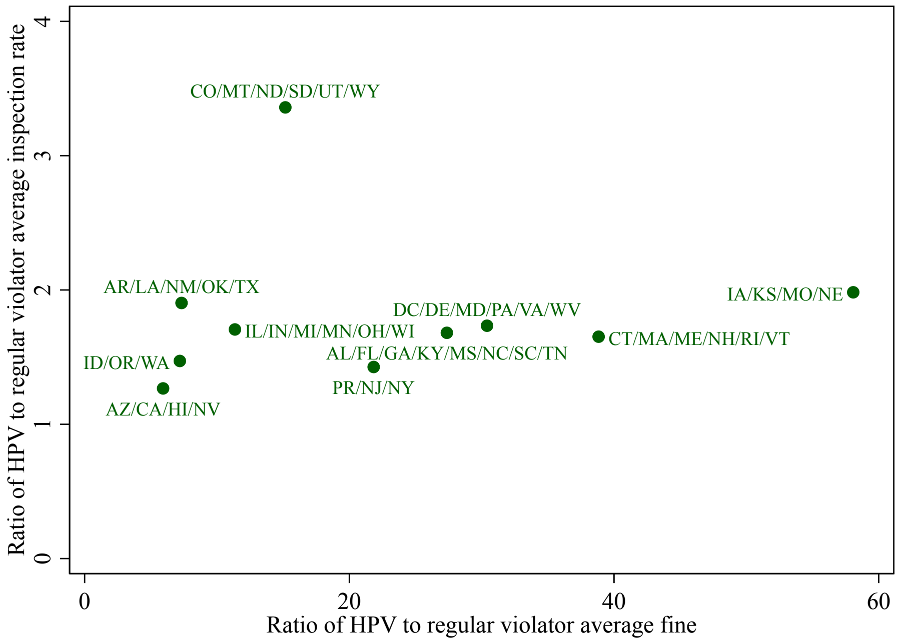

Note: authors' calculations based on estimation sample. States in each EPA region are indicated next to value.

Table A4: 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 start of period regulatory status.

Figure A2: Mean Inspection Probabilities and Fines by Industrial Sector

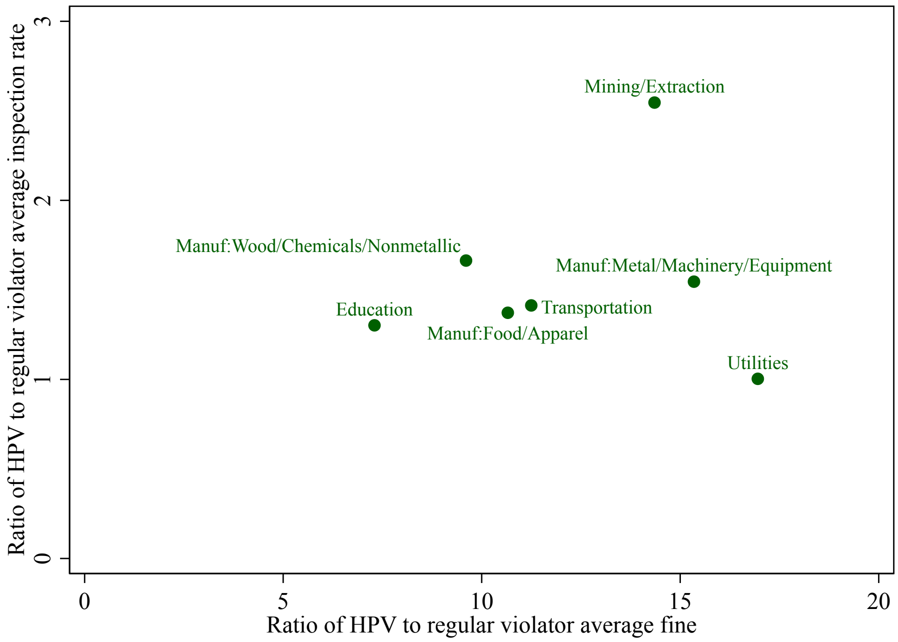

Note: authors' calculations based on estimation sample. Industrial sector measured by 2-digit NAICS code.

Figure A3: Depreciated Accumulated Violations and Monitoring and Enforcement

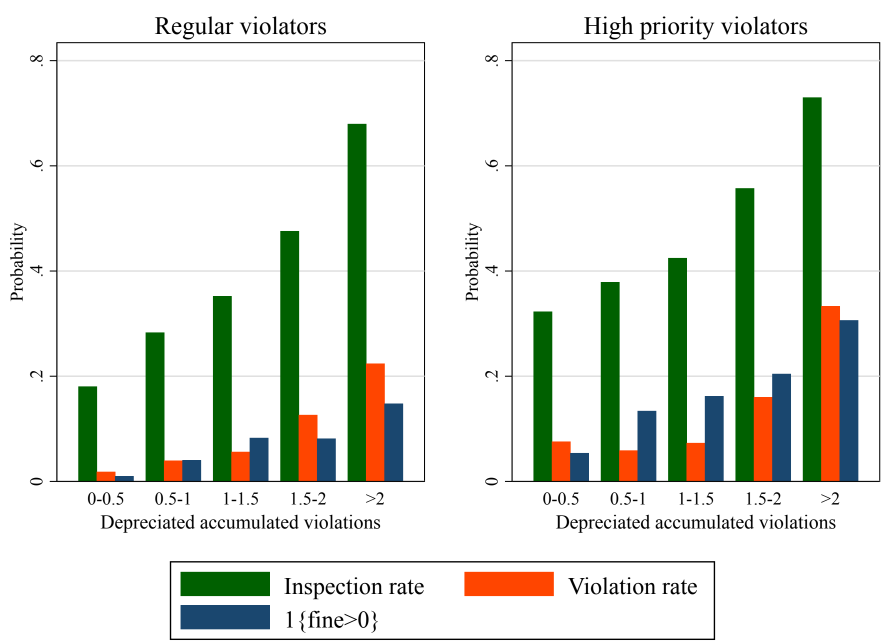

Note: authors' calculations based on estimation sample.

Table A5: Regulatory CCPs Marginal Effects: Inspections

	In compliance	Regular violator	HPV
Plant time-varying state			
Lag investment (0 to 1)	—	0.050	0.012
2nd lag investment (0 to 1)	—	0.100	0.043
Deprec. accum. vio. (mean to mean + 1)	—	0.126	0.110
Plant fixed state			
Non-attainment (given highest gravity)	−0.028	−0.022	0.006
Highest gravity and attainment (versus lowest)	−0.000	−0.022	−0.022
SE EPA region (versus SW)	−0.101	−0.026	0.040
Utility sector (versus manuf. food)	0.107	0.193	0.134
Mean	0.086	0.272	0.428
Pseudo R^2	0.085	0.091	0.075

Note: table shows marginal effects from probit regressions. Regressions include region, industry, and gravity state dummies. We run each regression separately by start of period regulatory status (compliance, a regular violator, or HPV). Each entry reports a marginal effect as described in the table.

Table A6: Regulatory CCPs Marginal Effects: Violations

	In compliance	Regular violator	HPV
Regulator actions			
Inspection (0 to 1)	0.021	0.063	0.085
Plant time-varying state			
Lag investment (0 to 1)	—	−0.007	−0.026
2nd lag investment (0 to 1)	—	−0.001	0.029
Deprec. accum. vio. (mean to mean + 1)	—	0.026	0.041
Plant fixed state			
Non-attainment (given highest gravity)	0.001	0.001	0.010
Highest gravity and attainment (versus lowest)	−0.000	0.006	−0.010
SE EPA region (versus SW)	−0.002	−0.010	−0.026
Utility sector (versus manuf. food)	−0.001	−0.003	−0.013
Mean	0.000	0.102	0.156
Pseudo R^2	0.182	0.152	0.099

Note: table shows marginal effects from probit regressions. Regressions include region, industry, and gravity state dummies. Most regressions also include inspection × gravity state interactions. We run each regression separately by start of period regulatory status (compliance, a regular violator, or HPV). Each entry reports a marginal effect as described in the table.

Table A7: Regulatory CCPs Marginal Effects: Fines

	In compliance	Regular violator	HPV
Regulator actions			
Violation (0 to 1)	0.000	0.020	0.279
Inspection (0 to 1)	0.000	0.024	0.176
Plant time-varying state			
Lag investment (0 to 1)	—	0.002	−0.592
2nd lag investment (0 to 1)	—	0.002	0.139
Deprec. accum. vio. (mean to mean + 1)	—	0.000	0.000
Plant fixed state			
Non-attainment (given highest gravity)	0.000	0.005	0.196
Highest gravity and attainment (versus lowest)	0.000	−0.001	−0.117
SE EPA region (versus SW)	0.000	−0.150	0.125
Utility sector (versus manuf. food)	0.000	−0.005	0.025
Mean	0.035	0.637	8.268
Pseudo R^2	0.187	0.245	0.108

Note: table shows marginal effects from tobit regressions. Regressions include region, industry, and gravity state dummies. Most regressions also include inspection $\times$ gravity state interactions. We run each regression separately by start of period regulatory status (compliance, a regular violator, or HPV). Each entry reports a marginal effect as described in the table.

Figure A4: Model Fit: Further Investments in the Six Periods After Initial Investment

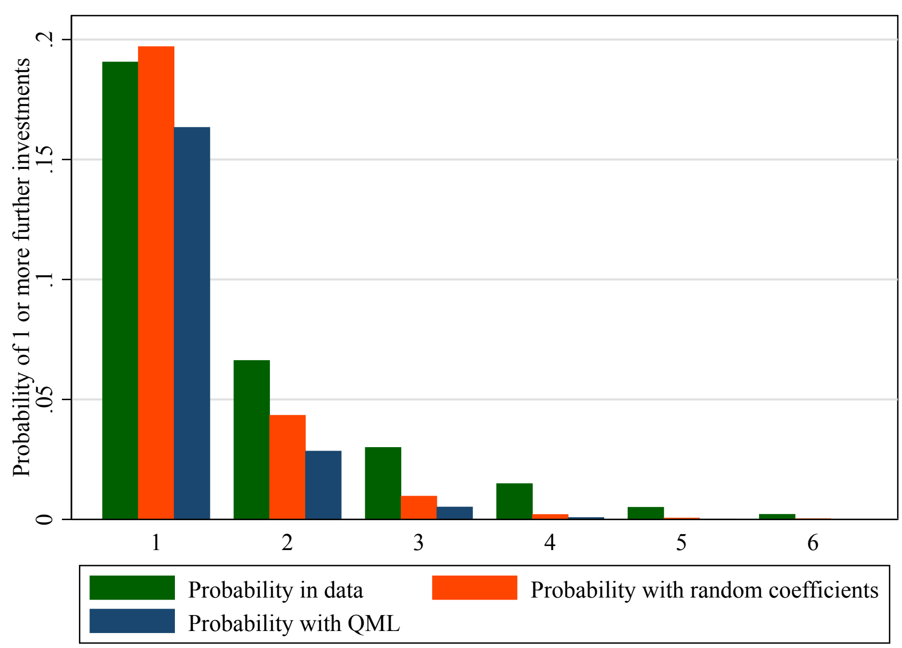

Note: authors' calculations based on estimation sample and estimated models evaluated at steady state.

Table A8: 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 (mean to mean + std. dev.)	0.000	0.000	−0.048	0.001	−0.018	−0.001
Violation (0 to 1)	0.676	0.166	−0.123	0.132	−0.118	−0.017
Inspection (0 to 1)	0.006	0.004	−0.007	0.013	−0.013	−0.002
Plant time-varying state						
Lag investment (0 to 1)	—	—	0.313	−0.004	0.461	0.248
2nd lag investment (0 to 1)	—	—	0.136	0.007	−0.046	−0.008
Deprec. accum. vio. (mean to mean + 1)	—	—	0.032	0.004	−0.030	0.013
Plant fixed state						
Non-attainment (given highest gravity)	0.000	0.000	0.004	0.002	0.007	−0.005
Highest gravity and attainment (versus lowest)	−0.000	−0.000	−0.012	−0.000	0.000	−0.001
SE EPA region (versus SW)	0.002	−0.004	0.186	−0.152	−0.044	0.044
Utility sector (versus manuf. food)	−0.000	0.000	−0.011	0.011	−0.004	−0.006
Pseudo R^2	0.502		0.175		0.307	

Note: table shows marginal effects from multinomial logit regressions. Regressions include region, industry, and gravity state dummies. Most regressions also include inspection × gravity state interactions. We run each regression separately by start of period regulatory status (compliance, a regular violator, or HPV). Each entry reports a marginal effect as described in the table.

Table A9: Estimates of Plants' Structural Parameters: More Interactions in CCPs

	Quasi-likelihood estimates	GMM random coefficient estimates (1)	(2)	(3)	(4)	(5)	(6)
Negative of investment cost ($-\theta^X$)	-2.856^{***}	−2.856	−2.318	−2.482	−1.906	−1.778	4.404
	(0.023)						
Inspection utility (θ^I)	-0.083^{***}	−0.083	−0.228	−0.130	0.106	−2.553	−2.323
	(0.028)						
Violation utility (θ^V)	0.039	0.039	0.260	0.767	−0.362	−1.356	−0.870
	(0.074)						
Fine utility (millions \$, θ^F)	-5.328^{***}	−5.328	−4.529	−6.114	−5.993	−7.055	−7.238
	(0.225)						
HPV status utility (θ^H)	-0.081^{***}	−0.081	−0.045	−0.094	−0.168	−2.564	0.377
	(0.007)						
Weight on parameter vector	1	0.273	0.265	0.213	0.175	0.049	0.008

Note: standard errors for quasi-likelihood estimates, which we calculate via an outer product formula, are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. GMM estimates are for a one-step estimator, unlike main results. For GMM estimates, we report the 6 parameter vectors with the highest weight. The CCPs used in these estimates include region-by-industry fixed effects instead of region and industry fixed effects.

Table A10: Estimates of Plants' Structural Parameters for Mining and Extraction Only

	Quasi-likelihood estimates	GMM random coefficient estimates (1)	(2)	(3)	(4)	(5)	(6)
Negative of investment cost ($-\theta^X$)	-2.316^{***}	-1.175	-2.219	-2.189	-0.964	-5.324	-8.918
	(0.074)						
Inspection utility (θ^I)	-0.129	-1.111	-0.993	-0.938	-0.201	2.320	-0.496
	(0.121)						
Violation utility (θ^V)	-0.218	-1.490	-2.481	-2.225	-1.449	-1.609	-2.616
	(0.657)						
Fine utility (millions \$, θ^F)	-5.891^{***}	-3.505	-6.039	-4.307	-3.728	-7.091	-8.272
	(1.155)						
HPV status utility (θ^H)	-0.058^{***}	-0.205	-0.074	-0.333	-0.341	-0.821	0.215
	(0.018)						
Weight on parameter vector	1	0.603	0.209	0.131	0.022	0.012	0.010

Note: standard errors for quasi-likelihood estimates, which we calculate via an outer product formula, are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. GMM estimates are for a one-step estimator, unlike main results. For GMM estimates, we report the 6 parameter vectors with the highest weight. Estimation uses only data from mining and extraction (2-digit NAICS code 21). Within this, the estimation uses the 6-digit NAICS codes with the most plant / quarters, 211111, 211112, 212312, and 212321, and EPA regions 3-8. Estimation replaces 2-digit NAICS code fixed effects in the CCPs with 6-digit NAICS code fixed effects.

Table A11: Estimates of Plants' Structural Parameters for 10 Most Populous States

	Quasi-likelihood estimates	GMM random coefficient estimates (1)	(2)	(3)	(4)	(5)	(6)
Negative of investment cost ($-\theta^X$)	-3.354^{***}	-2.843	-3.856	-6.458	-3.689	-0.813	-3.838
	(0.036)						
Inspection utility (θ^I)	-0.038	0.572	-1.195	-0.070	-0.286	-1.539	0.311
	(0.042)						
Violation utility (θ^V)	0.827^{***}	0.041	1.209	-0.359	0.467	3.314	-0.447
	(0.076)						
Fine utility (millions \$, θ^F)	-7.139^{***}	-8.967	-9.615	-8.258	-5.384	-4.934	-5.670
	(0.271)						
HPV status utility (θ^H)	-0.184^{***}	-0.181	-0.129	-0.155	-0.020	-2.466	-0.257
	(0.009)						
Weight on parameter vector	1	0.417	0.222	0.144	0.053	0.051	0.048

Note: standard errors for quasi-likelihood estimates, which we calculate via an outer product formula, are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. GMM estimates are for a one-step estimator, unlike main results. For GMM estimates, we report the 6 parameter vectors with the highest weight. Estimation uses only data from CA, TX, NY, FL, IL, PA, OH, MI, GA, and NC and replaces region fixed effects in the CCPs with state fixed effects.

Table A12: Counterfactual Results With the Quasi-Likelihood Estimates: Changing the Escalation Rate of Fines

	(1) Data	(2) Baseline	(3) Same fines for all violators; fines constant	(4) Same fines for all violators; pollution damages constant	(5) Fines for HPVs doubled relative to baseline
Quasi-likelihood estimates					
Compliance (%)	95.62	94.66 (0.12)	91.45 (2.84)	94.81 (0.15)	95.06 (0.12)
Regular violator (%)	2.88	3.91 (0.11)	3.78 (0.13)	3.49 (0.11)	3.91 (0.11)
HPV (%)	1.50	1.43 (0.04)	4.77 (2.91)	1.70 (0.14)	1.03 (0.03)
Investment rate (%)	0.40	0.44 (0.01)	0.43 (0.02)	0.51 (0.02)	0.45 (0.01)
Inspection rate (%)	9.65	9.43 (0.06)	10.60 (1.36)	9.52 (0.09)	9.31 (0.05)
Fines (thousands $)	0.18	0.32 (0.04)	0.32 (0.04)	1.51 (0.29)	0.38 (0.05)
Violations (%)	0.55	0.54 (0.01)	1.08 (0.85)	0.60 (0.06)	0.50 (0.01)
Plant utility	—	−0.007 (0.004)	−0.003 (0.006)	−0.013 (0.004)	−0.008 (0.004)
Pollution damages (mil. $)	1.65	1.54 (0.02)	1.87 (0.26)	1.54 (0.02)	1.50 (0.02)

Note: each statistic is the long-run equilibrium mean, weighting by the number of plants by region, industry, and gravity state in our data. Plant utility reports the average flow utility across types and states including ε except for Euler's constant. Column (1) presents the value of each statistic in our data. Column (2) presents the results of our model given the estimated coefficients and the existing regulatory actions and outcomes. Other columns change the state-contingent fines and HPV cost faced by plants. Columns (3) and (4) impose the same fines for all regular and high-priority violators for a given fixed state. Column (5) doubles the fines for plants in HPV status. All values are per plant / quarter. Bootstrapped standard errors are in parentheses.