Dynamic Regulation with Firm Linkages: Evidence from Texas

Matthew Leisten*
Federal Trade Commission

Nicholas Vreugdenhil
Arizona State University

January 19, 2023
Link to most recent version here.

Abstract

We evaluate the efficiency of *linked* environmental regulation. Linked regulation allows inspectors who uncover violations at one plant to increase enforcement at other plants that share a common owner. When compliance costs are correlated, regulators can then target scarce enforcement resources towards bad actors without inspecting everyone. In this paper, we develop a framework of dynamic moral hazard under linked regulation. Plants choose pollution mitigation efforts, while regulators selectively target based on scores. Our framework allows for large portfolios of plants and for choices to be interdependent within the portfolio of plants and across time. We apply the framework to the Texas Commission on Environmental Quality (TCEQ) who uses a scoring-based system of linked regulation to enforce environmental regulations. One score reflects each plants' compliance history, allowing enforcement to target past violators, while the other score reflects the compliance history for all plants owned by the same firm, allowing enforcement to be linked across co-owned plants. We evaluate this program using a novel panel of plant inspections, violations, and scores. We use our estimated model to evaluate the efficiency of linked regulation compared to alternative regimes.

*The views expressed in this article are those of the authors and do not necessarily reflect those of the Federal Trade Commission or any individual Commissioner. Email: mattleisten@gmail.com; nvreugde@asu.edu. Thank you to Wesley Blundell, Scott Holladay, and Conor Ryan, as well as audiences at AERE, FTC, IIOC, and SEA, for useful comments.

1 Introduction

Environmental regulation covers a wide range of economic activities and industries, but its effectiveness is hampered by two key complications. First, incomplete information necessitates that regulators inspect plants to uncover violations. Second, this monitoring— and sometimes the enforcement of penalties— is costly, but regulators only have limited budgets. As a result, a central question for regulators is how to target their inspections and tailor penalties efficiently.

One common form of targeting involves *linked* regulation. If the costs of pollution abatement are correlated across agents— for example, through managerial practices at commonly owned plants— inspecting one agent may be informative about other agents. Linked regulation leverages this correlation across plants and implies that past violations by one plant can lead to more intense scrutiny or harsher penalties for other linked plants in the future.[1]

The Texas Commission on Environmental Quality (TCEQ) uses linked regulation to enforce the Clean Water Act and the Resource Conservation and Recovery Act, which governs hazardous waste. They implement this using a two-dimensional scoring system that includes a plant-specific score, which reflects individual plants' compliance history, and a firm-wide score, which reflects the compliance history of all plants owned by the same firm. The plant score allows the regulator to increase its regulatory pressure on plants suspected to have difficulty complying with the laws, and the firm-wide score allows the regulator to increase pressure on co-owned plants that may also be non-compliant without having to inspect every plant. If the regulator wants to escalate its regulation of plants that are likely to be bad actors, and plants' compliance costs are correlated with co-owned plants, then some degree of linking is efficient. Linked regulation allows the regulator to increase its scrutiny of other plants that are likely to be non-compliant without necessitating every plant get inspected.

In this paper, we answer the question: how effective is linked regulation? To do this, we develop a new empirical model of moral hazard under dynamic, linked monitoring and regulation. Firms have private information about their abatement costs, which may be correlated across the plants in their

[1]The degree to which enforcement should be linked across agents is a frequently discussed aspect of regulatory design in both environmental laws and also other regulations. In terms of environmental regulation, as well as the case study discussed in this paper, examples include the EPA's MACT standards. In terms of non-environmental regulation, linked regulation is a feature of Occupational Safety and Health Administration (OSHA) regulation. Similarly, there are debates about whether linked regulation should occur in, for example, nursing home regulation.

portfolios. Plants decide how much to abate pollution. Regulators target inspections and assign penalties based on plant-specific and firm-wide scores, which evolve over time. We estimate this model using the Texas regulation as a case study and use our estimates to evaluate the effectiveness of linked regulation and to test alternative regulatory designs.

Linked regulation belongs to a class of enforcement mechanisms known as "escalation mechanisms". An escalation mechanism involves penalizing worse offenders more than non-offenders per additional violation. The efficiency of escalation mechanisms has been supported both theoretically [Mookherjee and Png [1994], Polinsky and Shavell [1998]] and empirically in the single-plant case [Blundell et al., 2020].[2] Furthermore, the notion that a principal will optimally condition its policy towards one agent on the actions of other agents when types are correlated is widespread in the mechanism design literature [Crémer and McLean [1988], McAfee and Reny [1992]].

We begin the empirical analysis by documenting that the TCEQ's scoring rules are an effective means of dynamic linked deterrence. We first show that pollution abatement costs are correlated within-firm: relatively few firms are responsible for a large share of violations, and violations by one plant are correlated with violations by other plants owned by the same firm. We then show that inspections are targeted to plants with higher plant-level and firm-level scores (note that *higher* compliance scores in this setting imply that a plant or owner has a *poorer* environmental record). Furthermore, conditional on plant fixed effects, higher probabilities of inspection are associated with fewer violations, evidence that the scores are an effective means of deterrence.

Because violations, scores, inspections, and penalties are interrelated, we must develop and estimate an equilibrium model of dynamic, linked regulation to assess counterfactual regulatory schemes. In our model, firms are endowed with a private *type* that governs how costly it is for them to reduce their pollution. This type may be correlated with the types of other plants owned by the same firm. The regulator assigns plant-specific and firm-wide scores to these plants based on what it learns by inspecting these plants. Inspections reveal violations, which are then converted into penalties for pollution. In equilibrium, firms mitigate pollution optimally given the regulators' inspection regime, and the distribution of actions and scores is stationary. We incorporate a flexible *conduct parameter* that governs the extent to which management at one plant internalizes the penalties faced

[2] We discuss in the related literature section in more detail about how we build on the insights from [Blundell et al., 2020].

by other plants owned by the same firm. This might be true if the firm promotes good communication between management at the different plants.

Intuitively, linking penalties by focusing heavily on the firm-wide score is more efficient when plants' types are more correlated within the firm. This way, an inspector can threaten to assess higher fines on plant B if its type is correlated with plant A and plant A is caught polluting, without having to inspect both plants. Linking scores has an especially strong deterrent effect if plant A internalizes the harms it places on plant B through this linkage.

We estimate the model using an unbalanced panel of 9792 plants regulated by the TCEQ either under the Resource Conservation and Recovery Act or the Clean Water Act from 2012-2020. We observe detailed information about each plant, including the firm, each environmental inspection, whether the inspection detected a violation, penalties, etc. For the analysis, we aggregate the data to the year-plant level. The average plant is linked with approximately 3 other plants through a common firm (with some firms owning large portfolios of more than 50 plants).

Estimation is complicated by the curse of dimensionality. Fully solving the dynamic model is computationally infeasible because of the size of the state space: plants that internalize dynamic, linked deterrence must keep track of their own score as well as the scores of every other plant owned by the same manager. We use *continuation value sufficiency* as a way to circumvent this problem. Under continuation value sufficiency, the plant only keeps track of its own plant-level score and the firm-level score and accounts for the continuation value of choosing action a for all other plants owned by the same firm as a simple scalar-valued function that approximates the true continuation value from the full model. We finally recover the regulator's costs of inspection and benefits of pollution reduction by leveraging the optimality of the inspection policy function given the penalty scheme.

We use our estimates to examine the efficiency of targeted and linked regulation compared to random enforcement. Motivated by the limited enforcement budget of the regulator, we simulate a 10 percentage point increase in the budget and explore how it would most optimally add value. We consider allocations of inspections under four different scenarios and predict the reduction in plant violations. In the base case, all additional inspections are performed randomly. In the other three cases, inspections are targeted based on plant scores, based on firm scores, and optimally based on a combination of plant and firm scores, respectively. We show that the first case– targeted, unlinked

escalations– outperform random inspections by 25.36%. Linked escalations do even better, outperforming random inspections by 33.14%. The optimal mix outperforms random inspections by 33.68%. Overall, our findings support the conclusion that linking inspections can improve outcomes significantly over unlinked escalations in regulation.

Contributions We make three main contributions in this paper. First, we construct a novel dataset of compliance scores, regulatory inspections, violations, and penalties, for the 9792 plants regulated under the Resource Conservation and Recovery Act or the Clean Water Act by the TCEQ from 2012-2020. A novel feature of the data in our setting (as compared to recent papers in the literature) is that we directly observe the multi-dimensional score that the regulator uses to target inspections and penalties at each plant. The availability of this scoring data— as well as a clearly defined algorithm for how the score is updated— allows us to unpack the "black box" of how the regulator makes decisions. Specifically, in our counterfactual simulations, we are not only able to test whether the regulation is effective but also explore how alternative designs of the regulation would affect environmental outcomes.

The second contribution is that we build a new empirical framework to study dynamic linked regulation. The framework combines elements of the moral hazard literature and the structural dynamic discrete choice literature. While other work has studied dynamic escalation mechanisms which utilize the compliance history of a single plant over time, the novel element of our model is that we extend this literature by additionally incorporating targeting across the portfolio of plants linked by a common manager.

The third contribution is a set of new findings about the efficacy of dynamic linked regulation, including counterfactual regulatory designs. Our findings show that linked regulation can present a significant improvement over escalation at the plant level, which is already an improvement over non-targeted, random inspections.

The remainder of the paper is as follows: Section 2 reviews the related literature. Section 3 introduces the institutional details of the TCEQ's enforcement of the Resource Conservation and Recovery Act and Clean Water Act and the data we use. Section 4 shows descriptive evidence of the main phenomena we investigate in the paper. Section 5 lays out the abstract model while Section 6 details our empirical specification and our estimation routine. Section 7 presents our estimates. Section 8

discusses our counterfactuals, and Section 9 concludes.

2 Related literature

This paper is related to three strands of literature. The first is the literature that empirically estimates models of regulation. [Blundell et al., 2020] estimate a dynamic model of EPA's enforcement of the Clean Air Act and illustrate how escalation mechanisms can add value in the single-plant case.[3] The papers [Kang and Silveira, 2021], [Abito, 2020] [Duflo et al., 2018] structurally estimate models of environmental regulation to study alternative designs. [Assunçao et al., 2022] provide a framework for optimal targeting of deforestation enforcement in the Amazon using a blacklist. The findings of this paper add to this literature by quantifying how a common feature of regulatory design— linking— can add value.

The second strand is the empirical literature that documents the importance of firm management in decision-making, and the consequences more broadly for understanding markets and policy. For example, [Bloom et al., 2019], [Goldfarb and Xiao, 2011], and [Giardili et al., Forthcoming], document from a number of perspectives the importance of firm management— as opposed to the underlying physical technology of plants— as a key driver of firm outcomes. Our paper complements these findings, and highlights the relevance of firm management for designing optimal regulation, as well as providing an estimable framework that explicitly incorporates managerial linkages across plants.

Finally, this paper is related to the theory literature on mechanism design. This literature includes papers about the optimal design of deterrence mechanisms such as [Mookherjee and Png [1994], Polinsky and Shavell [1998]], and mechanism design when agent types are correlated e.g. [Crémer and McLean [1988], McAfee and Reny [1992]]. [Varas et al., 2020] study a trade-off between random and deterministic inspections. Random inspections attenuate moral hazard, while deterministic inspections are more effective at targeting inspections when they are most informative about firms' compliance.

[3] While the major distinction of our paper is that we focus on designing regulation that incorporates linkages across plants and through time, there are also some other key differences. For example, our model views firms as facing a moral hazard problem (which we argue is the primary margin of response for firms regulated under the Clean Water Act and the Resource Conservation and Recovery Act), while [Blundell et al., 2020] focus on the decision to invest in clean technology (which is the primary margin of response for firms under the Clean Air Act). Furthermore, we estimate the regulator's objective function.

3 Context

In this paper, we focus on environmental violations subject to the Resource Conservation and Recovery Act and the Clean Water Act. The Resource Conservation and Recovery Act governs the safe "generation, transportation, treatment, storage, and disposal of hazardous wastes", where "hazardous wastes" range from nuclear material to oil from deep fat fryers. The Clean Water Act governs water pollution. Although both these laws are federal, the federal Environmental Protection Agency (EPA) delegates authority to States for their enforcement. In Texas, the Texas Commission on Environmental Quality (TCEQ) is the primary governing body that administers these regulations.

The TCEQ uses a "risk-based system" to determine inspections of firms and assess penalties (such as fines). A risk-based system means that, due to limited resources that preclude inspecting all firms regularly, the regulator needs to carefully target enforcement. For example, the 2012 Sunset Commission reports that *"the agency has implemented risk-based approaches to attempt to use its available resources wisely and on the most serious environmental concerns"*. In the next subsection, we discuss in more detail exactly how this targeting takes place.

3.1 Texas environmental enforcement

In this sub-section, we describe how exactly the TCEQ implements its risk-based approach to targeting inspections and penalties in Texas. We begin by discussing how the regulator determines the "risk" of each plant that they inspect, which is through a two-dimensional scoring system based on the compliance history of each plant and manager. We now discuss the mechanics of how the scores enter into elevated penalties and increased future inspections. In Section 4 we will document how these escalations enter into our data. We also show that firms respond to potential escalations consistent with a key overall aim of the targeting system to *"deter future noncompliance"* (Texas Sunset Commission (2012)).

The regulator uses two scores to target enforcement. The first score is a "site rating" which we refer to as the "plant score". This score captures the history of violations at a particular geographical site. The second score is a "person rating" which we refer to as the "firm score". This score aggregates the plant scores over the portfolio of plants that share a common manager. As we document later, most plants in the data share a common manager with at least one other plant, and so there is a meaningful distinction between the two scores.

Plant score The plant score is designed to take the history of compliance violations, inspections, etc at a particular site and combine it into a one-dimensional index. Data from the past five years are used to compute the score. We detail the exact scoring algorithm that the regulator uses in Appendix A.1. Overall, the score indicates how 'far from compliance' a particular site is. A plant with a clean record will have a score of 0.0. Higher scores indicate worse compliance histories.

Complexity points The regulator normalizes plant scores by "complexity points" when making enforcement decisions. Adjustments for complexity points were introduced in 2012 because there is substantial firm heterogeneity in terms of, for example, plant scale and sector. Complexity points allow different types of plants to be compared under the same rating system. For example, a medium-sized facility may mechanically generate more violations than a small facility simply because the larger size allows for more opportunities for violations; adjusting by complexity points allows for the scores to be comparable.

Firm score The firm score is the average of all the plant scores, weighted by the complexity points. Denote $\mathcal{J}_f$ as the set of plants that are managed by manager f. Denote s_j as the score of plant j and Q_j as the corresponding complexity score of plant j. Then, the firm-wide score is computed as follows:

$$s_f = \frac{\sum_{j \in \mathcal{J}_f} Q_j s_j}{\sum_{j \in \mathcal{J}_f} s_j}.$$

Penalties and Inspections The regulator escalates penalties and inspections based on the scores. The penalties are escalated using the following rule:[4]

$$\text{Penalty}(s_j, s_f) = \text{Base Penalty} \times \text{Plant Escalation}(s_j) \times \text{Firm Escalation}(s_f) \tag{1}$$

Here the base penalty is dependent on the gravity of the violation. After determining the base penalty, the total penalty is then escalated based on the compliance scores. The firm escalations are: $\text{Firm Escalation}(s_f) = 0.9$ if the firm is a 'high performer' ($c_f \in [0, 0.1)$), $\text{Firm Escalation}(s_f) = 1.0$ if the firm is a 'satisfactory performer' ($c_f \in [0.1, 55)$) and $\text{Firm Escalation}(s_f) = 1.1$ if the firm is an 'unsatisfactory performer' ($c_f \in [55, \infty)$). It is important to note that the firm score affects all of the plants in a firm's portfolio. Therefore, even small changes in firm scores are amplified

[4] More information about the penalty policy can be found here

Table 1: Summary statistics

Variable	N	Mean	Std. Dev.
log(1 + plant score)	55,175	0.64	1.04
log(1 + manager score)	55,175	0.77	1.02
inspection probability	55,175	0.28	0.45
# other co-owned plants	55,175	2.99	5.31
# violations (cond. on inspection)	15,801	2.27	5.23

across many plants and so can potentially be quite costly to the firm, resulting in large deterrence effects. The plant escalations Plant Escalation(s_j) are also increasing in the scores. We discuss how inspections escalate with scores in Section 4.

3.2 Data

Our data are an unbalanced panel of the universe of plants regulated by the TCEQ either under the Resource Conservation and Recovery Act or the Clean Water Act from 2012-2020. We observe extremely detailed information about each plant, including the manager, each environmental inspection, whether the inspection detected a violation, the nature of the violation, and the monetary penalty that was incurred. For the analysis, we aggregate the data to the year-plant level.

Table 1 contains information about key summary statistics in the data. Due to the long right tail of plant scores and manager scores, we transform these variables using a $\log(1 + x)$ transformation. Here, a value of 0 corresponds to the lowest score possible, which implies a clean compliance history with no violations in the past five years. In the summary statistics, the probability that the average plant is inspected each year is 0.28. The fact that inspections occur infrequently (typically, once every few years) is consistent with inspections being costly to the regulator. It is quite common for these inspections to uncover violations: on average, around 2.3 violations are discovered for each inspection.

We also report the number of co-owned plants in Table 1. This statistic is, for each plant, the number of other plants that it is linked with through a common manager. For instance, if a manager has four plants then this number is four. We find that the average plant is linked with approximately 3 other plants through a common manager. In the data, some managers own large portfolios of more than

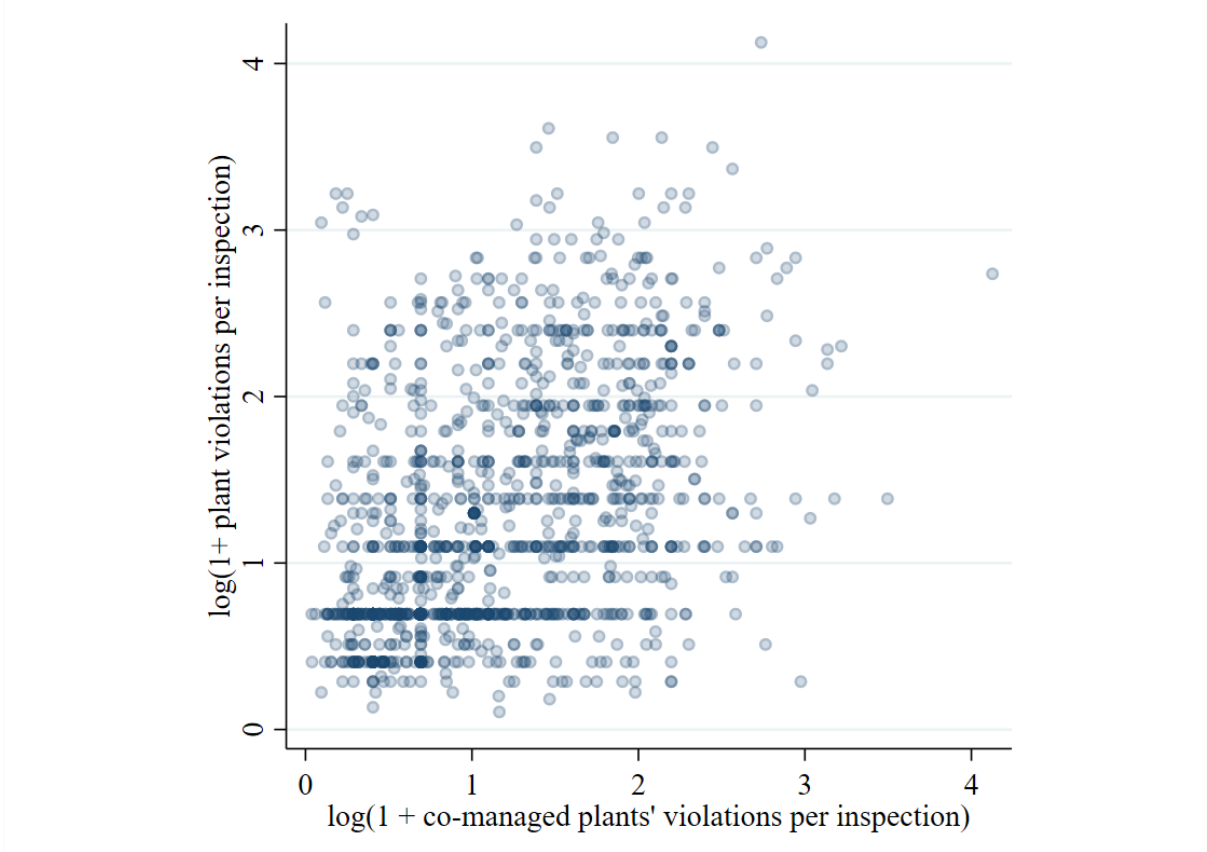
50 plants.

4 Descriptive analysis

We now provide several pieces of descriptive evidence that illustrate how linked regulation operates in practice in our setting. These pieces of evidence also motivate some key assumptions of our model. First, commonly-owned plants' cost of compliance with regulations (i.e., plant "types") must be correlated. Second, regulators must selectively target their inspections at bad actors. Third, plants must comply more with pollution regulations when their scores increase. We show each of these descriptively.

4.1 Violations are correlated within commonly-managed plants

Figure 1: Correlation between violations per inspection for plant and other plants within same firm



Note: does not include observations with zero plant or co-managed plant violations per inspection

In Figure 1 we show that plants which have more violations per inspection are jointly owned with

other plants that also have more violations per inspection. The horizontal axis of the figure is the average number of violations per inspection at plant j over the course of the panel. The vertical access shows the average number of violations per inspection at all other plants owned by the same firm as plant j, leaving out plant j itself.

The positive correlation in Figure 1 suggests that plants' pollution mitigation actions are correlated with the actions of other co-owned plants. Actions are endogenous, however. We would ideally show that plants' pollution mitigation costs are correlated with other co-owned plants. The correlation in costs is likely to be *higher* than correlation in actions under the TCEQ regulatory scheme. If commonly-owned plants A and B both have high costs of abatement, and plant A commits violations causing its score to increase, then plant B will be deterred from committing violations. This will attenuate the correlation in violations at plant A and plant B.

4.2 The regulator targets inspections (as well as penalties) based on linkages

Next, we show that inspections are targeted towards bad actors. Across all plant-year observations in the panel, only 28% are inspected, suggesting that inspection is costly. To identify which plant characteristics predict inspection, we report the results from logit regressions where the binary outcome is whether a plant is inspected in Table 2. Regardless of whether we condition on industry category-level or plant-level fixed effects, plants are more likely to be inspected when they have higher scores, which is consistent with greater scrutiny of plants that are suspected to be bad actors. This is true of both the plant-specific and firm-level scores.

4.3 Evidence of moral hazard and deterrence

Next, we demonstrate that higher scores do have the desired effect of deterring plants from polluting. In Table 3 we report estimates from regressions where the dependent variable is number of violations uncovered at a plant. Specification (1) shows that the probability of inspection implied by plant scores does not significantly predict changes in violations. However, this reflects two opposing effects: on one hand, plants with higher difficulty complying have higher scores and violations, while on the other hand, higher scores deter plants from violating. Specification (2) eliminates the first effect by conditioning on plant fixed effects. When a plant's score increases, its violations tend to decrease, consistent with higher scores having a deterrence effect. The estimate on (2) is still attenuated by the fact that plants with higher scores are more likely to be inspected, and therefore

Table 2: Inspection probability regressions

	(1)	(2)
Dependent Variable	Inspection	Inspection
Log(1+firm score)	.047	.063
	(.017)	(.018)
Log(1+plant score)	.153	.123
	(.017)	(.017)
Year FEs	Yes	Yes
Unit FEs	None	Category
N	55175	55175
Adjusted R^2	.020	.028

plants with higher scores are more likely to have violations *uncovered* in any particular year. We address this in specification (3) by estimating the same specification as in (2) but using only plant-years where an inspection occurred. Here, the coefficient on plant scores becomes larger and more significant, suggesting that plants' pollution-reducing actions are even more responsive to scores than specification (2) suggests.

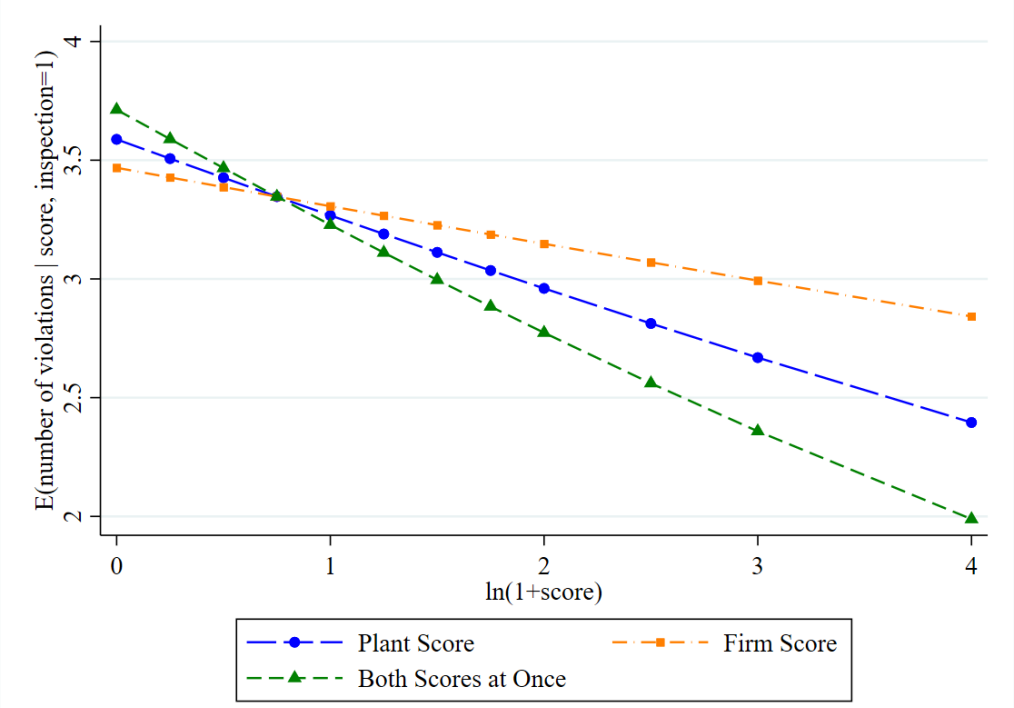
Using our estimates in this section it is possible to compute the implied effect of higher scores on deterring inspections via a higher probability of inspection. We do this using the regression (1) from Table 2 to compute the probability of inspection at different scores, and then plugging these probabilities into the regressions in Table 3. We do this separately for the firm-wide score and the manager-specific score. Because the regressions include fixed effects, we normalize the fixed effect so that the probability of inspection, and the expected number of violations conditional on inspection, both equal the sample means when the score equals 0.75. The results are plotted in Figure 2. Because the plant score has a larger estimated effect on the probability of being inspected, the deterrence effect of plant scores is bigger. Both, however, yield strong deterrence effects: an increase in log scores from 0 to 1.5, which corresponds roughly to the interquartile range, decreases the expected number of violations by roughly 0.5.

Although these results suggest that a combination of dynamic scoring and targeted inspections deter plants from polluting, we cannot ascertain a causal relationship from this analysis. The exact

Table 3: Deterrence regressions

Dependent Variable	(1) Violations	(2) Violations	(3) Violations
Pr(Inspection)	-.231	-3.196	-3.882
	(.589)	(.602)	(.678)
Year FEs	Yes	Yes	Yes
Unit FEs	Category, Firm	Plant	Plant
Only Inspected	No	No	Yes
N	40999	34840	10496
Pseudo R^2	.307	.35	.413

Figure 2: Implied deterrence effect of scores



Note: Implied effect of a change in plant score, firm score, or both at once, on expected number of violations, conditional on inspection, for the average firm in our dataset. The fixed effect of inspection regression is normalized so that probability of inspection equals the mean of 0.28 when all log scores equal 0.75, approximately the mean score.

relationship between scores, inspections, actions, and penalties is more complicated. Violations are only recorded if a plant is inspected, the expectation of inspection drives a plant's actions, and inspections are targeted based on scores. These relationships become even more complicated for multi-plant firms, where managers trade off the costs of mitigation at one plant versus the effects of harsher regulation across several plants. Pollution abatement may be a strategic substitute: if one plant pollutes less, other plants may be able to pollute more without facing steeper penalties. In equilibrium, higher manager-level scores may induce plants to increase or decrease their pollution abatement efforts.

If a plant's score increases, it not only faces greater scrutiny today, it faces greater scrutiny in the future, and other jointly-owned plants face higher scrutiny in the future. The extent to which each of these effects matters will depend on the scores and abatement costs of each plant. In order to disentangle all of these effects, as well as to measure the welfare implications of different regulatory schemes, we must estimate a model that explicitly accounts for the objectives of managers and regulators.

5 Model

5.1 Outline of model

Our model is a dynamic, discrete-time game that involves two main sets of agents: a set of firms $\mathcal{F}$ indexed by f and a regulator. Each firm owns a set of plants $\mathcal{J}_f$, with each plant indexed by j. Since scores are updated once a year, in our model one period is one year.

Plants differ in their industry sector g (e.g. manufacturing, resources, transportation, etc). Let $g(j)$ be the industry of plant j. Each plant is also endowed with a type θ_j. Higher types correspond to higher costs of pollution abatement. Plant types of co-owned firms may be affiliated, i.e., $E(\theta_j|\theta_{j'})$ is weakly increasing in $\theta_{j'}$, where $j, j' \in \mathcal{J}_f$. We let the joint distribution over the vector of all plants' types, from which the realized plants' types θ are drawn, be $G_\theta(.)$. Affiliated types within plant portfolios in our model could in principle be generated from the presence of 'bad managers', or alternatively they could be generated by managers who specialize in managing plants with high compliance costs. For the policy experiments that we consider in this paper, the results only depend on the degree of correlation in types, and so we do not need to take a stand on the exact cause of the

correlation.

In each period t, the timing is as follows. It is also outlined as a schematic in Figure 3.

1. Each plant enters the period with a score s_{jt} and each firm has a firm-level score s_{ft}.

2. The firm chooses a vector of pollution actions, one for each plant, a_{jt}. At the same time, inspectors decide which plants to inspect. We denote whether or not plant j is inspected at time t by z_{jt}, and model it as a draw from a Bernoulli distribution with a sector-specific mean $\bar{z}_{g(j)}(s_{jt}, s_{ft})$.

3. Conditional on inspecting a plant j, the inspector observes a number of violations v_{jt} which is drawn from a Poisson random variable with mean a_{jt}. If the plant is not inspected then no violations are detected, and so we skip this step and the next, to the final step where scores are updated.

4. Penalties are assessed which depend on the industry sector. These penalties scale with the number of detected violations. Total penalties per violation are given by $x_{g(j)}(s_{jt}, s_{ft})$.

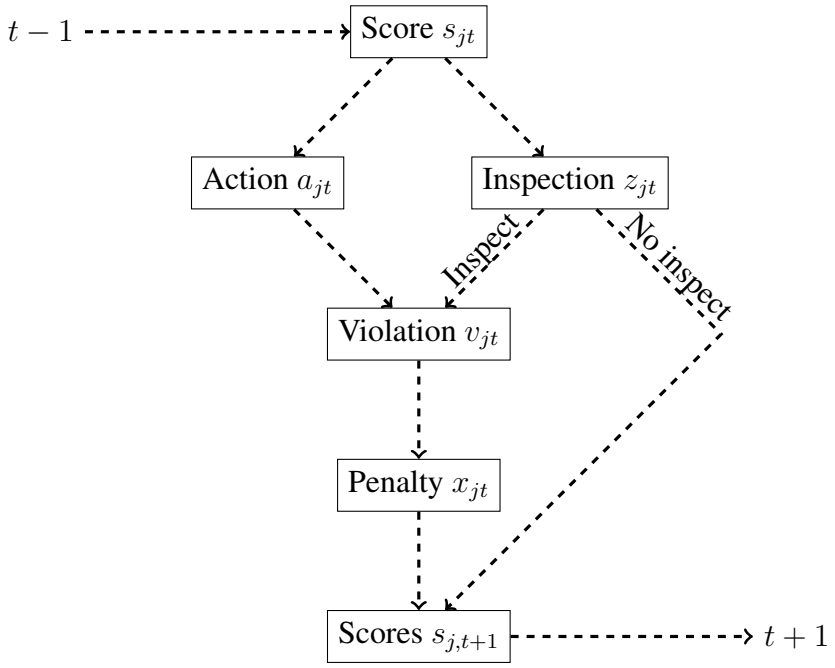
5. Plant scores are updated to the next period ($s_{j,t+1}$) as a function of the current scores and the number of detected violations. Firm-level scores are also updated using the aggregation rule (i.e. firm scores are the weighted average of the individual plants, where the weights are given by the complexity score of each plant: $s_{f,t+1} = \sum_{j \in \mathcal{J}_f} Q_{g(j)} s_{j,t+1} / \sum_{j \in \mathcal{J}_f} s_{j,t+1}$). We move to the next period.

The firm must trade off the flow benefits of higher pollution (or, equivalently, less pollution abatement) against the flow cost of higher expected violations and penalties, as well as the dynamic cost of increased scrutiny in future periods.

5.2 The firm's problem

We begin with the firm's decision problem. To keep notation concise we drop firm-specific subscripts in this section, but each firm will need to solve this problem given its particular portfolio of plants. Similarly, we drop time t subscripts. The expected flow payoff of plant j, given action a_j,

Figure 3: Timing of model for each plant j



with the expectation taken over stochastic violations and inspections, is:[5]

$$\pi_j(a_j; s_j, s_f) = \underbrace{\theta_j b(a_j)}_{\substack{\text{Flow benefit} \\ \text{from negligent} \\ \text{actions}}} - \underbrace{\bar{z}_{g(j)}(s_j, s_f)}_{\text{Inspection}} \times \underbrace{\mathbb{E}_{v_j}(v_j x_{g(j)}(s_j, s_f)|a_j)}_{\substack{\text{Expected penalty} \\ (=a_j x_{g(j)}(s_j, s_f))}} \tag{2}$$

The first component of Equation 2 is the plant's flow benefit from negligent actions. We call this a 'benefit' because $b(.)$ is increasing in the negligent action a_j, but it accrues due to cost savings. For example, a higher negligent action a_j provides a higher benefit to a firm manager because it requires less attention and effort to implement the proper pollution mitigation procedures. The second component in Equation 2 comprises the costs to the firm of taking a negligent action a_j. Here, the regulator— who we can think of as a 'regulatory machine'— first inspects the plant with probability $\bar{z}_{g(j)}(s_j, s_f)$. If the plant is inspected then the regulator will observe the true number of violations v_j which is stochastic and drawn from a Poisson distribution with mean a_j. Finally,

[5]We work with the expected value of the flow payoff conditional on the action because there is an analytical expression and it simplifies the exposition. Note that the state transitions— and therefore the expectation of the value functions— will also be a function of the stochastic realizations of violations and inspections, and so the entire problem cannot be written analytically.

after observing violations, the regulator administers penalties $x_{g(j)}(s_j, s_f)$, which we assume scale multiplicitively with the number of violations. Note that $\mathbb{E}_{v_j}(v_j x_{g(j)}(s_j, s_f)|a_j) = a_j x_{g(j)}(s_j, s_f)$.

Using these flow payoffs, the firm's decision problem at each period is to choose the (continuous) level of negligence at each plant in its portfolio a_j. To make this choice, the firm will need to account for the fact that the action at each plant will affect the payoffs *across the entire portfolio of plants* $\mathcal{J}$. Denoting the vector of plant specific variables in bold, e.g., $\mathbf{a} = \{a_j\}_{j\in\mathcal{J}}$, the firms' problem is:

$$V(\mathbf{s}) = \max_{\{a_j\}_{j\in\mathcal{J}}} \sum_{j\in\mathcal{J}} \pi_j(a_j; s_j, s_f) + \beta \mathbb{E}_{\mathbf{z},\mathbf{v}} \left[V(\mathbf{s}') \middle| \mathbf{s}, \mathbf{a} \right] \tag{3}$$

Simplifying the state space. The above portfolio problem is extremely complex to solve. Intuitively, in order to determine the optimal action at each plant j a firm needs to think about how this will affect future payoffs at other plants, since the regulation in linked. This generates a curse of dimensionality when solving for the optimal action at each plant: the firm manager needs to think about the entire state space s which is of dimension $n_{plant} + 1$.[6] In our application many firms have large portfolios (that is, n_{plant} can be large) hence solving Equation 3 is essentially infeasible. Therefore, we need to an employ an alternative approach. Our approach begins from the observation that the optimal actions a_j^* that satisfy Equation 3 are equivalent to the optimal actions that jointly satisfy the following n_{plant} equations (one for each plant):

$$\max_{a_j} \underbrace{\pi_j(a_j; s_j, s_f)}_{\text{Plant flow payoff}} + \beta \mathbb{E}_{\mathbf{z},\mathbf{v}} \left[\underbrace{V_j(\mathbf{s}')}_{\substack{\text{Plant} \\ \text{continuation} \\ \text{value}}} + \underbrace{\sum_{k\in\mathcal{J}/j} V_k(\mathbf{s}')}_{\substack{\text{Other plants'} \\ \text{continuation values}}} \middle| \mathbf{s}, \mathbf{a} \right] \tag{4}$$

where $V_j(\mathbf{s})$ denotes the value function for an individual plant j (i.e. the lifetime discounted stream of payoffs to plant j) and therefore the firm's value function $V(\mathbf{s}) = \sum_{j\in\mathcal{J}} V_j(\mathbf{s})$.

Inspecting Equation 4 highlights where the state space complexity bites in the portfolio choice problem. Concretely, it is through the third term in Equation 4 which governs how firm j's action affects the future payoffs of other commonly-owned plants. If that third term were eliminated, the value function for plant j only— V_j— would only depend on s_j and s_f because flow payoffs only depend

[6] Because the firm score, s_f is implied by the n_{plant} plant-level scores s_j, the state space is technically only of dimension n_{plant}, but this is still insufficient to solve the curse of dimensionality for large firms.

on s_j and s_f. Therefore, to solve the curse of dimensionality, we make an assumption in the spirit of Gowrisankaran and Rysman [2012], which we call *continuation value sufficiency*. The intuition underlying continuation value sufficiency is that the firm makes decisions at each plant in its portfolio one at a time, and that the firm uses only three variables to make each decision: the plant's score s_j, the aggregate firm score s_f, and a third scalar, $W_j = \sum_{k \in \mathcal{J}/j} V_k(\mathbf{s})$, that summarizes the continuation values of all the other plants. While s_j and s_f enter the flow payoffs for plant j directly, W_j acts as a heuristic to account for the cross-plant effects of a_j. Note that all cross-plant effects are dynamic, since the flow payoffs π_j are separable across plants' actions. We assume that the transitions over these three states are governed by a separate AR(1) process for each plant; we discuss more about how we specify and compute this AR(1) process in the estimation section.

Defining the resulting firm-level problem. The above assumptions turn the computationally intractable problem in Equation 4 that requires solving a value function over a $n_{plants} + 1$ dimensional state space, into n_{plant} separate value function computations each over three states, which is computationally feasible. Denote $\hat{\mathbf{s}}_j = (s_j, s_f, W_j)$. Then— leaving a formal derivation to the Appendix— the following three equations fully characterize the firm's problem. First, the optimal action $a_j^*(\hat{\mathbf{s}}_j)$ is given by:

$$a_j^*(\hat{\mathbf{s}}_j) = \arg\max_{a_j} \ \pi_j(a_j; s_j, s_f) + \beta \mathbb{E}_{z_j, v_j}\left[V_j(\hat{\mathbf{s}}_j') + W_j' \middle| \hat{\mathbf{s}}_j, a_j\right] \tag{5}$$

where plant j's value function can be defined with respect to this optimal action $a_j^*(\hat{\mathbf{s}}_j)$:

$$V_j(\hat{\mathbf{s}}_j) = \pi_j(a_j^*(\hat{\mathbf{s}}_j); s_j, s_f) + \beta \mathbb{E}_{z_j, v_j}\left[V_j(\hat{\mathbf{s}}_j') \middle| \hat{\mathbf{s}}_j, a_j\right] \tag{6}$$

and the state transition beliefs are governed by the AR(1) process (where $\hat{v}_j$ denotes detected violations):

$$\hat{\mathbf{s}}_j' = R_{0,j} + R_{1,j} \times [s_j, s_f, W_j, \hat{v}_j]' \tag{7}$$

Note the subscripts in Equation 7 indicate that the transition matrices $R_{0,j}$ and $R_{1,j}$ need to be computed separately for each plant as part of the solution algorithm. Concretely, they need to be consistent with the larger portfolio of plants to which j is a member, as well as the updating rules of the scores defined by the regulator. Also, note that we do not include an additive error term in

our AR(1) process for the state belief evolution. However, the state evolution is not deterministic because detected violations comprise a random draw from a Poisson distribution with mean a_j and inspections are also stochastic. We detail the full solution algorithm— which comprises an outer loop where we solve for the beliefs in Equation 7 and an inner loop where we compute actions and value functions that satisfy Equation 5 and Equation 6— in the estimation section.

5.3 Equilibrium: firms

Define a steady state equilibrium as the joint distribution of actions and scores given an inspection policy Z and a vector of types for each plant θ, $F(a, s; Z, \theta)$, that satisfies the following conditions:

1. Given the inspection policy function Z, the penalty escalations, and beliefs about state transitions, each firm chooses actions optimally at each state according to Equation 5 and Equation 6.

2. Beliefs about state transitions in the form of Equation 7 are consistent with the equilibrium.

3. The distribution of scores is in steady state. Given the law of motion for scores and optimal actions at each state $\mathbf{s}$, $F_s(\mathbf{s}') = F_s(\mathbf{s})$.

Intuitively, equilibrium satisfies three conditions. Firms behave optimally given their beliefs about how the state evolves; their beliefs are a good approximation to how the state actually evolves given their optimal actions; and the distribution of scores, after applying optimal actions and the law of motion for scores, is self-generating.

5.4 The regulator's problem

The regulator needs to choose an inspection rule from a set of allowable rules ($Z_g \in \mathcal{Z}$) to minimize the social cost of violations subject to an inspection budget constraint B. We expand in much more depth about how we parameterize the inspection policy function in the estimation section, but it might, for example, feature inspections that are entirely random, or inspections that are responsive to individual plant scores ('unlinked regulation'), or inspections that are responsive to firm scores ('linked regulation'). We allow for the social harm of a violation to vary across industry sectors, h_g, and in addition the inspection probability function is sector-specific, Z_g. The regulator does

not know the type of each plant, but does know the distribution of plant types G_θ. Formally, the regulator's problem is:

$$\min_{Z \in \mathcal{Z}} V^R = \min_{Z \in \mathcal{Z}} \int \int \sum_{f \in \mathcal{F}} \sum_{j \in \mathcal{J}_f} h_{g(j)} a_j \; dF(a, s; Z, \theta) dG_\theta(\theta) \tag{8}$$

$$\text{subject to:} \quad \int \int \sum_{f \in \mathcal{F}} \sum_{j \in \mathcal{J}_f} z_{g(j)}(s_j) dF(a, s; Z, \theta) dG_\theta(\theta) \leq B$$

In the above equations we are minimizing actions a_j directly and not the resulting violations. We set up the problem this way because the firms are in a steady state— and the violations are given by a Poisson distribution with mean parameter a_j— and so the expected number of violations for plant j is given by the action a_j.

Discussion In setting up the regulators' problem in Equation 8 we are making two implicit assumptions, which we now justify. First, we assume that the regulator does not place a weight on the private compliance costs of individual managers, but is simply trying to minimize the total costs of pollution. Our justification for this assumption is that scores are normalized by complexity points, and the inspection policy functions are defined over these normalized scores. An objective of complexity points— and a key reason they were introduced— is to ensure that enforcement is adjusted for the compliance costs of different types of firms. Of course, since compliance costs are private information, there is still dispersion in compliance costs even after normalizing for complexity points. However, in official reports the TCEQ tends to refer to the complexity points as the way that it incorporates variation in individual compliance costs into its regulation, rather than through ad-hoc adjustments to inspections, which is consistent with how we are estimating their objective function in Equation 8.

The second implicit assumption is that we are not optimizing for penalties. Rather, we are finding optimal inspections taking the penalty schedule as a given. Although the regulation is national, it is enforced by states, and the regulator in Texas is extremely prescriptive in its bylaws about how violations are determined (using a multi-page checklist), and also how violations translate into penalties. On the other hand, Texas allows for much more choice for the regulator as to whom to inspect. Therefore, we assume that the regulator takes the penalty schedule as fixed and its main choice is how to allocate inspections.

6 Estimation and identification

6.1 Parameterization

In order to take the model to the data, and also to leverage some institutional features, we make some parametric assumptions.

Plant types We assume plant j's type is $\theta_j = |\bar{\theta}_{g(j)} + \varepsilon_{f(j)} + \varepsilon_j|$. Here, $\bar{\theta}_g$ is the mean type of a plant in sector g, $\varepsilon_{f(j)} \sim_{i.i.d.} \mathcal{N}(0, \sigma_F^2)$ is a firm-wide type draw, and $\varepsilon_j \sim_{i.i.d.} \mathcal{N}(0, \sigma_J^2)$ is a plant-specific type draw. The larger σ_F^2 is relative to σ_J^2, the more within-firm correlation in types, all else equal.

Payoffs We also parameterize the plant's benefit of taking a negligent action $b(a) = a^y$, so a firm with type θ has marginal benefit of pollution equal to $\theta y a^{y-1}$. While θ governs different firms' relative private benefits of pollution, y captures the extent to which there are diminishing marginal returns to polluting more.

Law of motion for scores Recall that we assume that firms make decisions based on a reduced state space, with an estimated AR(1) process given by Equation 7. The plant-specific transition matrices $R_{0,j}$, $R_{1,j}$ in this law of motion are given by:

$$R_{0,j} = \begin{bmatrix} 0 \\ 0 \\ r_w^{0,j} \end{bmatrix}, \quad R_{1,j} = \begin{bmatrix} r_{s_j}^{s_j,g(j)} & 0 & 0 & r_v^{s_j,g(j)} \\ 0 & r_{s_f}^{s_f,j} & 0 & r_v^{s_f,j} \\ 0 & 0 & r_w^{w,j} & r_v^{w,j} \end{bmatrix} \tag{9}$$

We set a number of the coefficients in these matrices to 0.0. First, in the matrix of intercepts $R_{0,j}$ we set the intercept for the plant score updating and also the firm score updating equal to 0.0. This restriction follows directly from how scores work in our setting: if no violations are observed then the score will converge to 0.0.

The regulator only updates the plant score s_j using plant-specific information. Therefore, we switch the dependencies s_j on other states in $R_{1,j}$ to 0.0. An additional implication of using only plant-specific information to update this score is that even though agents' use a 'heuristic' to compute the state process, the remaining parameters in the first row in $R_{1,j}$ can be computed directly from the data (subject to the specific industry of j, $g(j)$).

The remaining coefficients— which enter into the updating rules for the firm state and also the value functions of the other states— need to be computed to be consistent with the equilibrium choices of the other plants in the portfolio, as well as how the regulator aggregates individual scores to the firm-level score. Since these coefficients need to be computed individually for each plant j, and there are thousands of individual plants in the data, we set four coefficients in the second and third rows of $R_{1,j}$ to 0.0 to ease the computational burden. These four coefficients relate to second-order interactions between the state variables; we still allow for states to be serially correlated with past realizations of themselves (through the diagonal terms) and detected violations (the final column).

Inspection functions We assume probability that the regulator conducts an inspection of plant j at time t is of the logit form, consistent with our regression specification (1) in Table 2:

$$z_{jt} = \frac{\exp\left\{\beta^z_{0g(j)} + \beta^z_1 \ln(1+s_{jt}) + \beta^z_2 \ln(1+s_{ft})\right\}}{1+\exp\left\{\beta^z_{0g(j)} + \beta^z_1 \ln(1+s_{jt}) + \beta^z_2 \ln(1+s_{ft})\right\}}. \tag{10}$$

When solving the regulator's problem, we limit $\mathcal{Z}$ to be logit functions measurable with respect to this same set of variables. We also do not assume, in estimation, that the regulator has chosen β^z_1 or β^z_2 optimally. Finally, in estimation, we assume the regulator has a fixed (and binding) inspection budget so that $\sum_j z_{jt} = \bar{z}$ in equilibrium. Therefore, our estimation imposes that the regulator has chosen $\beta^0_{g(j)}$ to minimize the total social cost of violations subject to this budget constraint.

Penalties Penalties are assessed according to a base penalty, plant-level scores, and firm-level scores, as defined in Equation 1. While we can calibrate the escalations as a function of the states directly from the TCEQ bylaws, we still need to estimate the base penalty per violation. To do this we take data on observed penalties. Since we also observe scores and know the escalation rule, we then deflate penalties by the escalations and divide the total penalty by the number of violations. We then estimate the base penalty per violation as the industry-specific mean of these deflated penalties.

6.2 Estimation routine

The primitives we estimate are: those of the inspection function β^z, those of the penalty function β^x, those of the score updating rule r, the complexity weights Q_g, those governing the distribution of firm types $(\sigma^F, \sigma^2_J, \bar{\theta}_g)$, the curvature parameter y, and the marginal social cost of pollution h_g.

We estimate the model in three steps.

First step: We estimate the offline parameters (all denoted by β or Q) through ordinary least-squares regression (in the case of β^s and β^x), logit regression (in the case of β^z), and nonlinear least squares (in the case of Q). In addition, we can estimate the parameters on the plant score evolution.

Second step: We estimate $(\sigma_F^2, \sigma_J^2, \bar{\theta}_g, y)$ through a method of simulated moments procedure that contains two nested fixed point procedures: one to obtain the value function given a guess of the parameters, and another to obtain the CVS transition coefficients given a value function and a guess of the other parameters. The estimation routine is described in more detail in the Appendix. In addition, as part of this second step and also as part of computing the equilibrium, we estimate the parameters on the firm score evolution and also on the evolution of the W_j state.

Third step: Finally, we estimate h_g by leveraging the assumption that the regulator chooses the intercepts in the logit inspection probability function optimally. Given our estimates of the other parameters, we can approximate the derivative of V^R with respect to each of the g intercept terms, β_g^Z, by simulation and taking finite differences. We can use the functional form of the regulator's objective function V^R to analytically invert the set of first order conditions with respect to each intercept. This allows us to recover $h(g)$ directly from our linear approximation. Details are in Appendix E.

6.3 Identification and moments

We briefly and informally discuss the moments selected and how the moments selected identify the parameters of the model that are recovered in steps two and three of the above estimation routine. The parameters recovered in step two are: the mean plant type, $\bar{\theta}_g$, for each g of the six NAICS categories; the variances of the firm- and plant-level type draws σ_F^2 and σ_J^2, and the curvature parameter y.

To exposit identification, consider a highly simplified setting in which a plant of type θ_j considers only the present period and does not have any other co-owned plants, and suppose that uncovered violations are deterministic: $v_j = a_j$. The firm's problem is then: $a_j^* = \arg\max_a \theta_j a^y - za$. Taking

a first-order condition, rearranging, and taking logs yields the solution

$$a_j^* = \frac{1}{1-y}\log(\theta_j) + \frac{1}{1-y}\log(z). \tag{11}$$

From this equation, it is clear that differences in *levels* of violations across different plants identifies their types θ, while the overall responsiveness of recovered violations to changes in the probability of inspection (i.e., deterrence) identifies y.

Though our empirical model is more complicated, these comparative statics inform our choice of moments. To recover the average type for each NAICS sector, $\{\theta_g\}$, we use the average number of violations per year, conditional on inspection occurring, for all plants in that sector. Formally these moments are $\bar{v}_g \equiv \frac{\sum_{t,j:g(j)=g} v_{jt}}{\sum_{t,g(j)=g} \tilde{z}_{jt}}$ and the simulated moments, which are computed analogously. This yields $|g|$ moments. All else equal, a higher mean type for sector g corresponds to more violations on average for plants in that sector, conditional on inspection.

To identify σ_F^2 and σ_J^2, we examine the distribution of equilibrium violations— conditional on inspection— across firms. Suppressing the notation that conditions on inspection occurring, we decompose the total variance of violations as follows:

$$\text{Var}[v] = \underbrace{\mathbb{E}_j\left[\text{Var}_t(v_{jt}|j)\right]}_{\substack{\text{Within plant,}\\ \text{over time}}} + \underbrace{\mathbb{E}_f\left[\text{Var}_j[\mathbb{E}_t(v_{jt}|j)j \in \mathcal{J}_f]\right]}_{\substack{\text{Across plants,}\\ \text{within firm}}} + \underbrace{\text{Var}_f\left[\mathbb{E}_{j,t}(v_{jt}|j \in \mathcal{J}_f)\right]}_{\text{Across firms}} \tag{12}$$

We divide the second and third terms on the right hand side of this equation by the total variance Var$[v]$ to identify σ_F^2 and σ_J^2. All else equal, higher type draws will choose higher actions and commit more violations upon inspection. Therefore, greater variance in violations across plants indicates larger variances of the type draws. The share of this variance that occurs within a firm but across plants identifies σ_J^2, while the share of this variance that occurs across plants identifies σ_F^2.

As evident in Equation 11, curvature parameter y is identified by the responsiveness of violations to changes in deterrence. We construct the relevant moment in two steps: first, we compute the predicted probability of inspection using our first-stage estimates (also the fitted values from regression (1) in Table 2). Next, we regress violations on this predicted inspection probability and plant fixed effects, conditional on inspection occurring (specification (3) in Table 3). We compute the same moment analogously in our simualted data.

Table 4: Second-stage estimates

Category	Estimate	Std. Err
Mean type		
Utility	0.287	-
Services	0.019	-
Manufacturing	0.229	-
Resources	0.027	-
Transportation	0.270	-
Trade	0.411	-
Type variances		
Plant-level, σ_J^2	0.192	-
Firm-level, σ_J^2	0.241	-
Shape parameter		
y	0.376	-

Notes: Standard errors TBA.

The discussion in Appendix E elucidates identification of $\{h_g\}$ from the inspector's first order condition with respect to each of the sector-specific intercepts in the estimated inspection function. Generally, the greater the intercept for sector g, the larger the implied harms of pollution from that sector relative to the (normalized to 1 across all firms) cost of inspection c. However, because we impose a binding budget constraint, we do not know if these first order conditions hold exactly. Instead, we only know that the regulator allocates its budget of inspections efficiently across sectors in equilibrium. This allows us to identify the *relative* marginal harms of pollution. Therefore, we normalize $h_{\text{services}} = 1$ and estimate marginal harms of pollution in other industries relative to this.

7 Results

The first-stage estimates are in Appendix Table 7. The estimates that were recovered "structurally" are in Table 4. Notably, the variances of the plant-level and firm-level draws are similar, indicating a moderate degree of within-firm correlation conditional on a firm's portfolio. Differences across firms in their portfolio may induce even more correlation by firm: if firms that own one "trade" plant also own other "trade" plants, these firms are likely to have high types across all their plants.

Our estimates show that trade plants tend to have higher types, and services and resources plants tend to have lower types. As shown in Appendix Figure 5 these differences are not clearly reflected in heterogeneity in actions across NAICS sectors. In part, this is because plants' actions are determined by the level of scrutiny they receive, and, for instance, resources plants are not inspected as frequently (see Appendix Figure 6). Actions are also affected by the portfolio of co-owned plants because of cross-plant deterrence effects.

7.1 Model fit

We first assess model fit by reporting the fit of each moment used in estimation in Table 5. Broadly, we match the moments well.

In Figure 4 we plot the density of the joint distribution of plants' scores with the leave-one-out average of other co-owned plants' scores, similar to the scatterplot in Figure 1 but using simulated data evaluated at our estimates. Then we accurately capture the positive correlation between plants' average number of violations, conditional on being inspected, and the leave-one-out firm-wide average number of violations conditional on being inspected. The correlation coefficient in our simulations between these two measure is 0.32, while the correlation in the data is 0.26. We consider this good model fit, given this correlation pattern was not a directly targeted moment in our estimation.

8 Counterfactuals

8.1 Overview

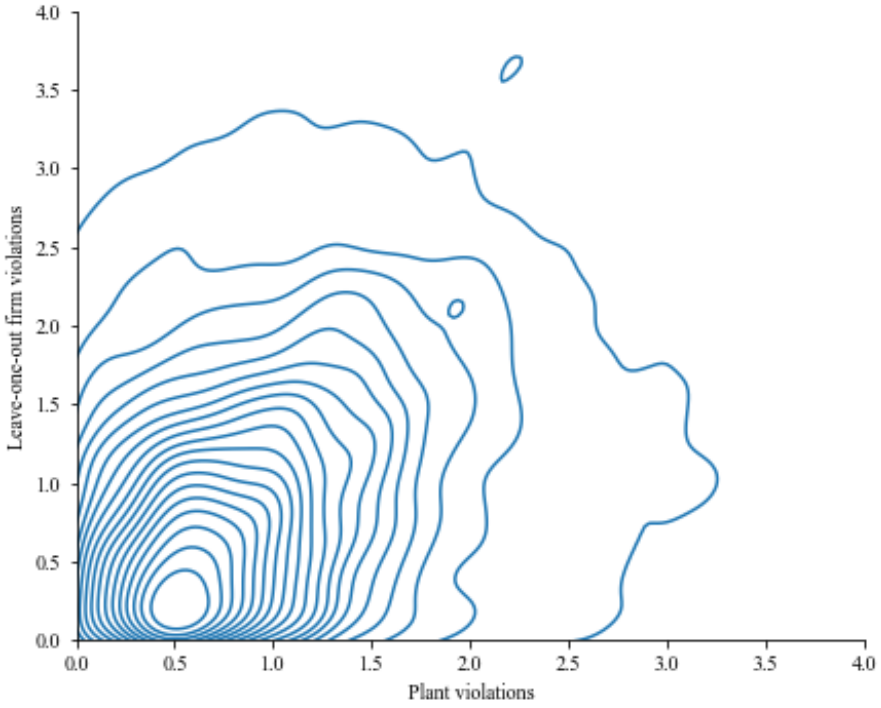
In order to evaluate the efficiency of linked regulation we perform several counterfactuals. These counterfactual experiments center around how to most effectively spend a 30% increase in the inspections budget (which would increase the probability of inspecting a plant by approximately 10%). We measure effectiveness in terms of the decrease in total violations. Specifically, we investigate how much more effective this budget increase would be compared to random inspections if: (i) the budget was completely spent on linked inspections; (ii) the budget was completely spent on unlinked inspections; (iii) the budget was split between linked and unlinked inspections optimally to minimize violations.

We also decompose each counterfactual into three components in order to illustrate how each form

Table 5: Model fit: moments

Moment	Empirical	Simulated
Mean violations by category		
Utility	0.201	0.184
Services	0.111	0.106
Manufacturing	0.164	0.139
Resources	0.108	0.088
Transportation	0.179	0.136
Trade	0.252	0.236
Share of variance in violations		
Within-firm	0.205	0.248
Across-firm	0.199	0.333
Responsiveness	-6.398	-6.095

Figure 4: Correlation of simulated firm and plant violations



of regulation works. These components are a static effect, a dynamic own-plant deterrence effect that comes from increased future scrutiny at a plant, and a dynamic cross-plant deterrence effect that comes from increased future scrutiny at plants co-owned with a violating plant. To evaluate cross-plant deterrence separately from own-plant deterrence, we generalize our baseline model to include a "conduct" parameter, $\omega \in [0, 1]$, that governs the degree to which the actions at one plant internalize the effects of that action on other co-owned plants. Formally, we incorporate ω into Equation 5 as follows, with ω multiplying W_j':

$$a_j^*(\hat{\mathbf{s}}_j) = \arg\max_{a_j} \ \pi_j(a_j; s_j, s_f) + \beta \mathbb{E}_{z_j, v_j}\left[V_j(\hat{\mathbf{s}}_j') + \omega W_j' \middle| \ \hat{\mathbf{s}}_j, a_j\right]. \tag{13}$$

When $\omega = 0$, plants do not internalize spillovers to other co-owned plants. For any counterfactual scenario, we the compute the components as follows, denoting $v_j^{\beta,\omega}$ as the number of violations at plant j with a discount factor of β and a conduct parameter of ω:

$$\begin{aligned}\sum_j v_j = &\underbrace{\sum_j v_j^{\beta=0.0,\omega=0.0}}_{\text{Static effects}} \\ &+ \underbrace{\sum_j \left(v_j^{\beta=0.95,\omega=0.0} - v_j^{\beta=0.0,\omega=0.0}\right)}_{\text{Own-plant effects}} \\ &+ \underbrace{\sum_j \left(v_j^{\beta=0.95,\omega=1.0} - v_j^{\beta=0.95,\omega=0.0}\right)}_{\text{Cross-plant effects}}\end{aligned} \tag{14}$$

The intuition behind Equation (14) is as follows. We begin with the static effects of the regulation, which we measure by the number of violations that occur when switching off the dynamic effects by setting the discount parameter $\beta = 0$ and also the conduct parameter $\omega = 0$. Note that agents will still be responsive to changes in the probability of inspections due to static deterrence: for example, if the probability of inspection increases this will increase the marginal cost of a negligent action, resulting in fewer violations. However, agents will not internalize the fact that their actions will affect future punishments of their own plant and also of co-owned plants.

The second component in Equation (14) captures the own-plant effects of the regulation. We measure this by switching on the discount factor $\beta = 0.95$ but keeping the conduct parameter $\omega = 0$.

The resulting additional decrease in violations is due to own-plant dynamic effects: agents reduce their negligent actions even further because if they are caught violating then future punishments will escalate at their own plant. Conceptually, this is the same economic force that is at work in Blundell et al. [2020].

The third component in Equation (14) captures the cross-plant effects of the regulation. We measure this by switching on the conduct parameter $\omega = 1$. The extra decrease in violations that occurs is due to cross-plant dynamic effects: plants internalize that if they are caught violating this will increase scrutiny not just at their own plant (through the plant score s_j) but also at all commonly-owned plants (through an increase in the manager score $s_{f(j)}$). Linked regulation will be most successful when types are highly correlated between commonly managed plants. For example, if there are 'bad managers' then linked regulation will result in the discovery of more violations at other plants within the same portfolio, resulting in a greater deterrence effect of the regulation.

Computational details Overall, we model changes to random inspections, unlinked inspections, and linked inspections, by modifying the parameters in the regulator's inspection function: $\beta^z_{0g(j)}, \beta^z_1, \beta^z_2$. Specifically, we use a root-finding algorithm to search for the coefficient where the total probability of inspection equals the budget increase. We interpret random inspections as an increase in the constant of the inspection function $\beta^z_{0g(j)}$. Similarly, changing the level of unlinked inspections corresponds to changing the coefficient on plant scores (β^z_1) and changing the level of linked inspections corresponds to increasing the sensitivity to manager scores (β^z_2). Note that to compute the counterfactuals we completely recompute the new equilibrium, which includes recomputing the equilibrium beliefs' of agents. We also search for the optimal escalations mix by searching over combinations of (β^z_1, β^z_2) that yield the target probability of inspection.

8.2 Discussion

We present the counterfactual results in Table 6. Recall that the objective of the regulator is to *minimize* violations and so fewer violations corresponds to 'more effective' regulation. Without the budget increase the average number of violations per plant (both discovered and not discovered by the inspector) is 0.88. To produce the numbers in Table 6 we first consider a random allocation of the inspection budget increase, which we find would reduce violations per plant by -0.28, and then compare how much more effective alternative types of regulation would be compared to this. In

Table 6: Summary of Counterfactuals

	↑ Inspections budget by 10%. Spent on:		
	Unlinked escalations	Linked escalations	Optimal escalations mix
%Δ Total violations vs random	-25.4%	-33.1%	-33.7%
Decomposition: Total			
= Static effect	2.5%	3.9%	2.8%
+ Own-plant effect	-29.7%	-22.1%	-23.8%
+ Cross-plant effect	1.8%	-14.9%	-12.7%
%Δ Detected violations vs random	-16.2%	-22.5%	-20.5%

Notes: In the baseline model the average number of violations per plant (both discovered and not discovered by the inspector) is 0.88. Allocating these extra inspections to random inspections reduces violations per plant by -0.28. The row %Δ Violations shows the increase in efficacy of the regulation compared to random inspections. To compute this row we find how much more effective the regulation would be— in terms of reduced violations— if the inspections are spent on different types of regulation (e.g. unlinked escalations, linked escalations). In the decomposition we split this total change into cross-plant deterrence, own-plant deterrence, and a static effect.

other words, the numbers in Table 6 are the change in violations compared to spending the additional inspections randomly.

Results: unlinked escalations only We first consider the results for the counterfactual world where inspections are entirely used for unlinked escalations (that is, escalations that only respond to plant scores). Consistent with Blundell et al. [2020] we find overall that these inspections add value and reduce violations by -25.4% compared to purely random inspections.

The decomposition illustrates how unlinked escalations add value compared to random inspections. First, the static effect is actually slightly positive; this implies that the random inspections actually do better when plants are completely myopic optimizers. The intuition behind this somewhat surprising result is that there is a diminishing marginal deterrence benefit of inspections, which implies that - when agents are completely myopic - it is more effective to spend the budget with a uniform probability of inspection across agents (i.e. random inspections).

As might be expected, the own-plant effect of unlinked escalations is to reduce violations. Con-

cretely, it reduces violations by -29.7% compared to random escalations. The cross-plant effect is also slightly positive; this may be surprising since unlinked inspections only directly affect own-plant deterrence. However, only considering this direct effect misses the complex interactions of the different types of regulation in this setting. Specifically, unlinked inspections directly improve the revelation of information at individual plants which then spills over to the existing regulation. For example, suppose that one of these extra unlinked inspections uncovers violations at a particular plant. The plant score will increase, which leads to increased scrutiny at this plant, reallocating existing inspections away from linked towards unlinked regulation. Furthermore, the firm internalizes these effects when choosing its optimal negligence level, leading to a decrease in the cross-plant effect.

Results: linked escalations only Next, in the second column of Table 6, we investigate the effects of spending the budget entirely on linked escalations. We find that these inspections tend to do better than unlinked escalations, reducing violations by -33.1% in total. The decomposition shows why: the own-plant effect is lower in magnitude than in unlinked escalations, but the cross-plant deterrence effect is greater much greater, and the second effect dominates. Finally, note that the percentage of detected violations is actually lower than under unlinked escalations; a key reason for this is that linked regulation is more effective and so there are fewer violations to uncover.

Results: optimal escalations mix Finally, we compute the optimal mix of linked and unlinked escalations in the third column. Note that our optimal escalations mix does not allow for additional inspections to be random. We find that the optimal mix is an interior solution and that the two types of regulation are somewhat complementary: allocate approximately 0.4 of the new inspections towards own-plant escalations, and the rest toward linked escalations. We find that the optimal mix does slightly better than pure linked escalations and significantly better than unlinked escalations. The decomposition illustrates why: the optimal has a higher own-plant effect in magnitude than pure linked escalations, and a higher cross-plant effect in magnitude that pure unlinked escalations.

9 Conclusion

Regulators often face incomplete information about which firms are bad actors, while only having limited enforcement budgets to detect and deter violations. Therefore a central question for regulators is how to efficiently target scarce enforcement resources. In this paper we develop an new

empirical framework to study dynamic linked regulation, which is a common form of targeting that is present in many real-world enforcement regimes. Linked regulation exploits the correlation of types within firms— driven potentially by the presence of 'bad managers' who are negligent across the plants in their portfolio— to improve targeting and the effectiveness of regulation. In the framework firms face a dynamic moral hazard problem of choosing negligent actions over time subject to scores, penalties, and the probability of inspection. The regulator chooses how to allocate a limited inspections budget to minimize the social cost of violations, observing only a set of scores for each plant. Our framework allows for firms with large portfolios of plants which they need to keep track of over time, while retaining computational tractability. We apply it to a novel dataset of scores, violations, inspections, and ownership structures, for the universe of firms in Texas who are subject to the Hazardous Wastes Act and the Clean Water Act. Our results show that linked regulation adds value and illustrates the complementarities behind unlinked targeting, linked targeting, and random targeting.

References

Jose Miguel Abito. Measuring the welfare gains from optimal incentive regulation. *The Review of Economic Studies*, 87(5):2019–2048, 2020.

Juliano Assunçao, Robert McMillan, Joshua Murphy, and Eduardo Souza-Rodrigues. Optimal environmental targeting in the amazon rainforest. 2022.

Nicholas Bloom, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John Van Reenan. What drives differences in management practices? *American Economic Review*, 109(5):1648–1683, 2019.

Wesley Blundell, Gautam Gowrisankaran, and Ashley Langer. Escalation of scrutiny: The gains from dynamic enforcement of environmental regulations. *American Economic Review*, 110(8): 2558–85, 2020.

Jacques Crémer and Richard P McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, pages 1247–1257, 1988.

Esther Duflo, Michael Greenstone, Rohini Pande, and Nicholas Ryan. The value of regulatory discretion: Estimates from environmental inspections in india. *Econometrica: Journal of the Econometric Society*, 86(6):2123–2160, 2018.

Soldedad Giardili, Kamalini Ramdas, and Jonathan Williams. Leadership and productivity: A study of us automobile assembly plants. *Review of Economics and Statistics*, Forthcoming.

Avi Goldfarb and Mo Xiao. Managerial ability and strategic entry in us local telephone markets. *American Economic Review*, 101(7):3130–3161, 2011.

Gautam Gowrisankaran and Marc Rysman. Dynamics of consumer demand for new durable goods. *Journal of Political Economy*, 120(6):1173–1219, 2012.

Karam Kang and Bernardo S. Silveira. Understanding disparities in punishment: Regulator preferences and expertise. *Journal of Political Economy*, 129(10):2947–2992, October 2021.

R Preston McAfee and Philip J Reny. Correlated information and mecanism design. *Econometrica: Journal of the Econometric Society*, pages 395–421, 1992.

Dilip Mookherjee and Ivan PL Png. Marginal deterrence in enforcement of law. *Journal of Political Economy*, 102(5):1039–1066, 1994.

A Mitchell Polinsky and Steven Shavell. On offense history and the theory of deterrence. *International Review of Law and Economics*, 18(3):305–324, 1998.

Felipe Varas, Iván Marinovic, and Andrzej Skrzypacz. Random inspections and periodic reviews: Optimal dynamic monitoring. *The Review of Economic Studies*, 87(6):2893–2937, 2020.

A Additional information about the regulation and context

A.1 Plant score construction

In this section we detail the exact algorithm used to compute the plant score. The score of plant j (s_j) is given as follows:[7]

$$s_j = \frac{\Big[(\text{Violation Points}) + (\text{Chronic Excessive Emission Points}) + (\text{Repeat Violator Points}) - (\text{Self Audit Points})\Big]}{(\text{No. of investigations} \times 0.1) + (\text{Complexity Points})} \tag{15}$$

The components in the above equation are as follows. The term 'violation points' corresponds to the number of violations at a particular plant (these may be weighted by the magnitude of the violation). The terms 'chronic excessive emission points' and "repeat violator points'' correspond to adjustments if the plant has a history of the same violation. The term "self audit points" correspond to minor score adjustments if plants have a formal environmental audit system. In the denominator, 'no. of investigations' corresponds to the number of recent inspections.

[7] This score may also be adjusted in minor ways if the plant participates in a 'voluntary program'.

B Additional tables and figures

Table 7: First-stage estimates

Variable	Estimate	Std. err.
Inspection probability function		
Plant rating	0.122	-
Firm rating	0.061	-
Intercept		
Utility	−0.848	-
Services	−1.270	-
Manufacturing	−0.949	-
Resources	−1.369	-
Transportation	−1.382	-
Trade	−1.459	-
Score law of motion		
Lagged rating	0.785	-
No. violations	0.012	-
Intercept		
Utility	0.086	-
Services	0.048	-
Manufacturing	0.061	-
Resources	0.061	-
Transportation	0.021	-
Trade	0.029	-
Complexity weights		
Utility	1.008	-
Services	1.193	-
Manufacturing	1.702	-
Resources	0.707	-
Transportation	0.293	-
Trade	1.085	-

Figure 5: Heterogeneity: Actions by NAICS sector

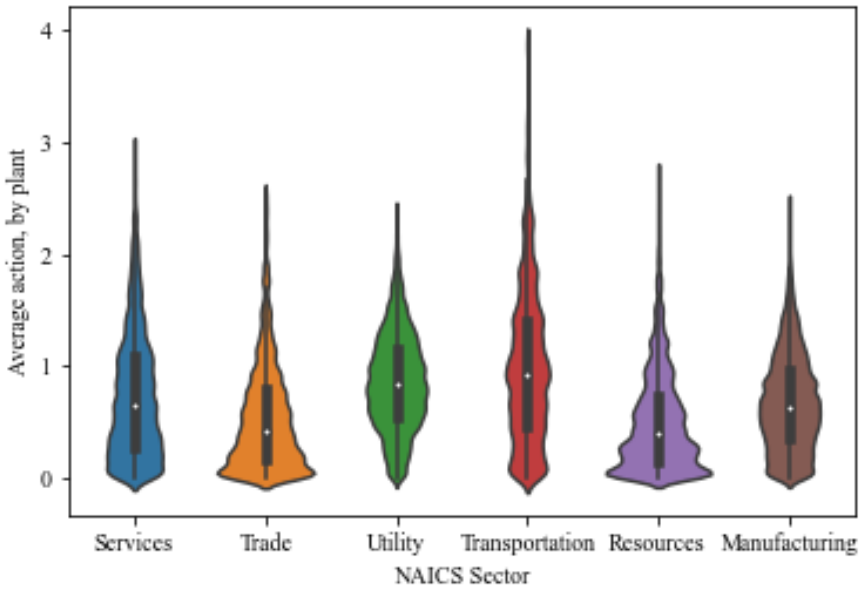


Figure 6: Heterogeneity: Inspection probability by NAICS sector

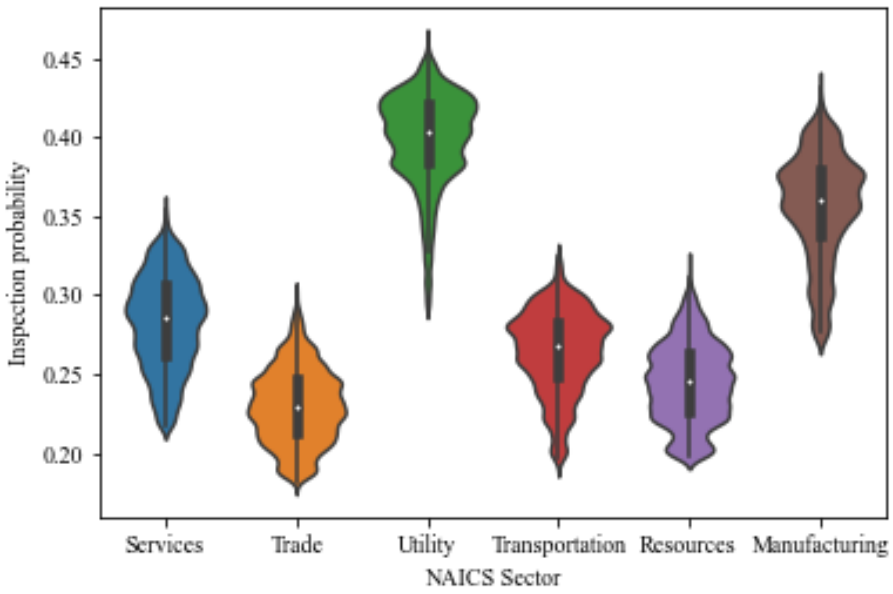


Figure 7: Types vs inspection probability, by plant

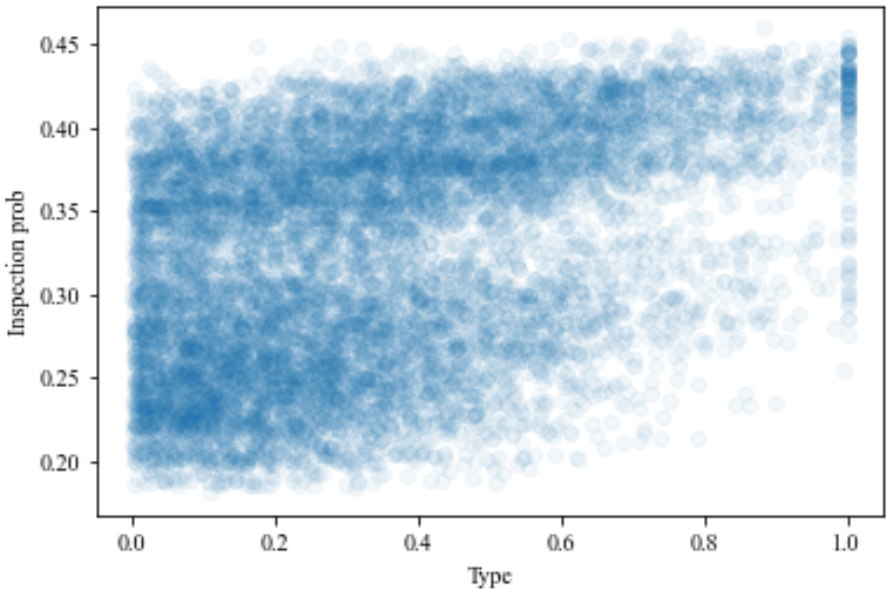
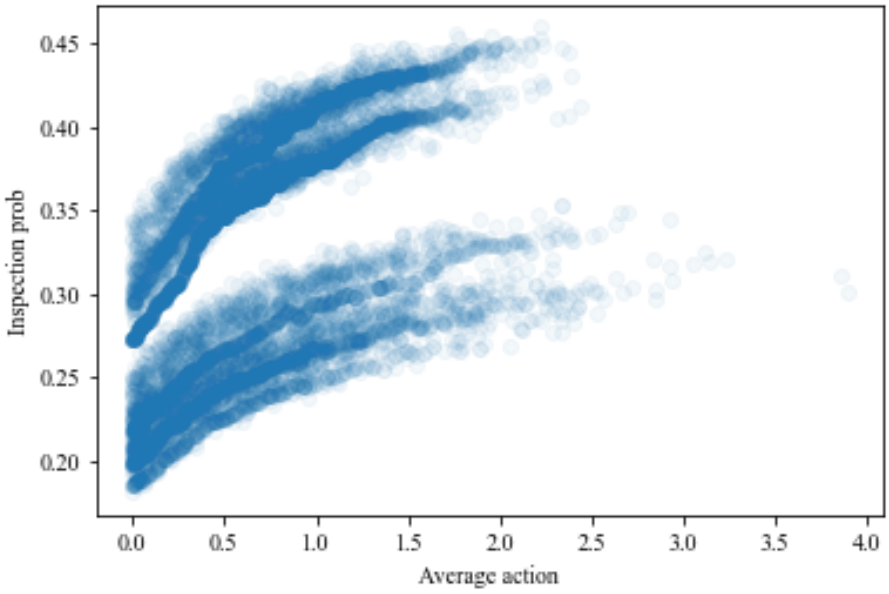


Figure 8: Average actions vs inspection probability, by plant



C More detail about the derivations and proofs

C.1 Derivation of the firm-level problem

Here we provide more detail about how the complete firm portfolio problem, detailed in Equation 4, results in Equation 5, Equation 6, and Equation 7, under our assumption of continuation value sufficiency. Rewriting our starting point, which is Equation 4:

$$\max_{a_j} \ \pi_j(a_j; s_j, s_f) + \beta \mathbb{E}_{\mathbf{z},\mathbf{v}} \left[V_j(\mathbf{s}') + \sum_{k \in \mathcal{J}/j} V_k(\mathbf{s}') \middle| \ \mathbf{s}, \mathbf{a}^*_{-j} \right] \tag{16}$$

Then, applying the continuation value sufficiency assumption:

$$\max_{a_j} \ \pi_j(a_j; s_j, s_f) + \beta \mathbb{E}_{\mathbf{z},\mathbf{v}} \left[V_j(\mathbf{s}') + W_{-j} \middle| \ \mathbf{s}, \mathbf{a}^*_{-j} \right] \tag{17}$$

At this point the optimal action could still potentially depend on the entire state space, and so we denote the optimal action $a^*_j(\mathbf{s})$. Note that the value function for plant j is given recursively in terms of the optimal action by:

$$V_j(\mathbf{s}) = \pi_j(a^*_j(\mathbf{s}); s_j, s_f) + \beta \mathbb{E}_{\mathbf{z},\mathbf{v}} \left[V_j(\mathbf{s}') \middle| \mathbf{s}, \mathbf{a}^*_{-j} \right] \tag{18}$$

Note that the state transitions, as well as Equation 17 and Equation 18, are only a function of v_j, W_{-j}, s_j, s_f. Therefore, the remaining states are not payoff-relevant and so we can replace $\mathbf{s}$ with $\hat{\mathbf{s}}_j = (s_j, s_f, W_j)$. Furthermore, the expectations can be written as simply over $\mathbf{z}_j$ and $\mathbf{v}_j$, which completes the derivation.

D Estimation routine

Below is the estimation routine for recovering the dynamic parameters that govern firms' decisions to abate pollution.

1. Compute moments from the data. We discuss our choice of moments more in the identification section.

2. Guess $\sigma^2_F, \sigma^2_J, \bar{\theta}_g$ and y.

(a) Guess the nuisance transition coefficients ρ^s and ρ^W.

i. Guess the value function.

ii. Given the current guesses, solve for the manager's policy function by differentiating and solving the Bellman Equation 5.

iii. Update the value function using the policy function and iterate until convergence of the value function.

(b) Using the converged value function and policy function, simulate actions, scores, and violations for T periods. Compute W_j at each simulated data point using the current iteration of the value function.

(c) Obtain a new guess of the ρ^W and ρ^s transition function by estimating an AR(1) process in the form in Equation 7 on the simulated data. Return to 2(a)(i). Iterate until convergence of the transition coefficients.

3. Use the final version of the simulated data to compute the simulated moments.

4. Search over σ_F^2, σ_J^2, $\bar{\theta}_g$, and y to minimize the distance between the simulated moments and the data moments computed in step (1).

E The regulator's problem: detail

Consider the regulator's problem, copied below for convenience:

$$\min_{Z \in \mathcal{Z}} \int \int \sum_{f \in \mathcal{F}} \sum_{j \in \mathcal{J}_f} h_{g(j)} a_j \; dF(a, s; Z, \theta) dG_\theta(\theta)$$

$$\text{subject to:} \quad \int \int \sum_{f \in \mathcal{F}} \sum_{j \in \mathcal{J}_f} z_{g(j)}(s_j) dF(a, s; Z, \theta) dG_\theta(\theta) \leq B$$

Simplifying notation and rearranging, the Lagrangian can be written as

$$\min_{z \in \mathcal{Z}} \mathbb{E}_{a,s,\theta} \left[\sum_j h_{g(j)} a_j - \lambda z_j \middle| Z(.) \right] + \lambda B.$$

When we assume that $\mathcal{Z}$ is the family of logit fucntions of the form in equation 10 and that the intercepts for each NAICS sector are chosen optimally, the first order conditions with respect to the intercepts β_g^Z

$$0 = \sum_j \left[h_{g(j)} \frac{\partial \mathbb{E}[a_j|Z]}{\partial \beta_g^Z} - \lambda \frac{\partial \mathbb{E}[z_j|Z]}{\partial \beta_g^Z} \right]$$

for each NAICS sector g. Stacking these $|g|$ first order conditions yields the equation

$$\mathbf{0} = \mathbf{A}\mathbf{h} - \lambda \mathbf{z}$$

where $\mathbf{A}$ is a $|g| \times |g|$ matrix with entry $A_{ik} = \sum_{j:g(j)=k} \frac{\partial \mathbb{E}[a_j|Z]}{\partial \beta_i^Z}$, $\mathbf{h}$ is a $|g| \times 1$ vector of $h(g)$, and $\mathbf{z}$ is a $|g| \times 1$ vector with ith element $Z_i = \sum_j \frac{\mathbb{E}[z_j|Z]}{\partial \beta_i^z}$. Provided $\mathbf{A}$ is nonsingular, inverting this equation yields

$$\mathbf{h} = \lambda \mathbf{A}^{-1} \mathbf{z}.$$

The scalar Lagrange multiplier λ is unknown and, in general, is not identified without more information, which implies that the entire vector $\mathbf{h}$ is not identified. However, the direction of vector $\mathbf{h}$ is pinned down by this system of equations, implying that h_g is identified up to the normalization that $h_{trade} = 1$.