Review of Economic Studies (2026) **00**, 1–29 doi:10.1093/restud/rdag013

Advance access publication 30 January 2026

Dynamic Regulation with Firm Linkages: Evidence from Texas

Matthew Leisten
Federal Trade Commission, USA

and

Nicholas Vreugdenhil
Arizona State University, USA

First version received June 2023*; Accepted June* 2025 *(Eds.)*

We evaluate the efficiency of dynamic *linked* environmental regulation. Linked regulation allows inspectors who uncover violations at one plant to increase future 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. We develop an empirical framework of dynamic moral hazard under linked regulation that allows for large portfolios of plants and for choices to be interdependent within the portfolio of plants and across time. Using the framework we evaluate a linked regulation scheme in Texas and find that linked regulation performs substantially better than both unlinked regulation and untargeted regulation. We test two alternative theoretical mechanisms that underpin the benefit—a "firm-wide moral hazard mechanism" and a "correlated targeting mechanism"—and find that a large share of the value of linked regulation is due to the former.

Key words: Regulation, Monitoring, Dynamic Estimation, Clean Water Act, RCRA

JEL codes: C57, L51, Q53

1. INTRODUCTION

Environmental regulation covers a wide swath of industrial activity, 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, and regulators only have limited budgets. As a result, a central question for regulators is how to target their enforcement efficiently. This question is particularly important because of the sheer scale of environmental enforcement in the U.S.: over 60,000 plants are subject to federal hazardous waste regulation alone, with 80% of the U.S. population living within three miles of a regulated facility (EPA, 2015). More generally, inspections are widely used in regulatory enforcement *e.g.* OSHA plant safety (OSHA, 2024), SEC enforcement actions

The editor in charge of this paper was Bård Harstad.

(Government Accountability Office, 2018), accounting violations (PCAOB, n.d.), and restaurant hygiene (Jin and Leslie, 2009).

One common form of targeting involves *linked* regulation, where regulators dynamically target plants that are co-owned with other plants that have a history of committing violations. If the costs of pollution abatement are correlated within a multi-plant firm—for example, through managerial practices at commonly owned plants—inspecting one plant may be informative about the firm's other plants. Linked regulation leverages this correlation through two theoretical mechanisms. The first is a "correlated targeting mechanism", where the links provide useful information about where next to investigate. The second is a "firm-wide moral hazard mechanism", where the regulator can punish a common owner for bad behaviour.

Although linked regulation is a widespread feature of enforcement design, little is known about whether it adds value in practice.[1] This paper aims to fill that gap. Overall, we answer the question: how effective is linked regulation? To do this, we build a new empirical framework of dynamic, linked regulation and estimate it using data from a linked enforcement regime in Texas. Our framework allows us to go beyond a simple quantification of the benefits of linked regulation; we are also able to test the extent to which each of the two theoretical mechanisms account for its effectiveness. Doing so is crucial to understanding how linked regulation would fare in other settings. For example, the firm-wide moral hazard mechanism relies partially on the regulator's ability to commit to a policy; the correlated targeting effect does not rely on commitment.

In our framework, firms have private information about abatement costs, which may be correlated across the plants in their portfolios, and they choose negligent polluting actions that are hidden from the regulator, *i.e.* there is moral hazard. Our framework accommodates large portfolios of plants and exploits novel data that includes compliance scores, inspections, violations, and pollution, for the universe of plants subject to environmental regulation in Texas. Our estimates show that linking significantly outperforms unlinked regulation (targeting at the plant level), as well as untargeted regulation.

We find that the dominant mechanism driving the value of linked regulation is the firm-wide moral hazard effect, with the correlated targeting effect playing a smaller role. This finding has implications beyond the setting of this paper since many group incentive structures in environmental protection and natural resource management rely on similar mechanisms to the firm-wide moral hazard effect (Segerson, 2022). Specifically, it suggests that these group incentive structures may be successful more generally.

The Texas Commission on Environmental Quality provides an ideal research setting as it uses linked regulation to enforce the Clean Water Act and the Resource Conservation and Recovery Act, which governs hazardous waste. Texas's regulation regime involves 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 be non-compliant 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.

1. Other linked environmental regulations include the EPA's MACT standards (EPA, 2005). Non-environmental linked regulations include Occupational Safety and Health Administration (OSHA) regulation (OSHA, n.d.) There are debates about whether linked regulation should occur in, for example, nursing home regulation (Abramo and Lehman, 2015).

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).

Our data are an unbalanced panel of 7,379 plants regulated by the Texas Commission on Environmental Quality under the Resource Conservation and Recovery Act and/or the Clean Water Act from 2012–2020. We observe detailed information about each plant, including the firm ownership network, scores, each environmental inspection, whether the inspection detected a violation, penalties incurred, and pollution. For the analysis, we aggregate the data to the year-plant level. The average plant is linked to a portfolio of approximately three other plants, but this average masks substantial heterogeneity in portfolio size: some firms own large portfolios of more than 40 plants.

We begin the empirical analysis by documenting that Texas's scoring rules are an effective means of dynamic linked deterrence. We first show suggestive evidence that pollution abatement costs are correlated within-firm: violations at one plant are correlated with violations by other plants owned by the same firm. We then show that inspections target plants with higher plant-level and firm-level scores (note that *higher* 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.

We next develop and estimate a model of dynamic, linked regulation to assess counterfactual regulatory schemes. In our model, plants 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 due to the presence of bad managers.[3] Our model also allows for observed heterogeneity (*i.e.* plants in different industry sectors).

The regulator assigns plant-specific and firm-wide scores to these plants based on what it learns by inspecting these plants. At the outset, the regulator publicly commits to a policy that maps scores and plant characteristics to a probability of inspection in a given period. Inspections reveal violations, which translate into penalties and update scores. Firms choose polluting actions optimally across each plant by trading off present benefits of negligence against the cost of present-period penalties, future escalations, and cross-plant linkages. In equilibrium, optimal firm behaviour gives rise to a steady-state distribution of actions and scores. On the regulator's side, different sectors have different perceived marginal social harms of pollution (as well as actual social harms), and regulators allocate their inspections across sectors accordingly.

Estimation is complicated by the curse of dimensionality. Fully solving the dynamic model is computationally infeasible: due to the scoring rule, optimal actions at each plant are a function of the state at every other commonly-owned plant, and some firms have large portfolios leading to a large state space. We use *continuation value sufficiency*, an assumption similar to inclusive value sufficiency in Gowrisankaran and Rysman (2012), as a way to circumvent this problem. Under continuation value sufficiency, the plant only keeps track of its own plant-level score and

2. We discuss in the related literature section in more detail about how we build on the insights from Blundell *et al*. (2020).

3. Types may be correlated for other reasons. Our analysis remains valid so long as this correlation structure is unaffected by the regulator's policy rule.

the firm-level score and accounts for the continuation value of choosing an action for all other plants owned by the same firm as a scalar-valued function that approximates the true continuation value from the full model. We finally recover the relative social costs of pollution, by sector, by leveraging the optimality of the allocation of inspections across sectors.

We use our estimates to examine the efficiency of linked regulation. Motivated by the limited enforcement budget of the regulator, we simulate an increase in the budget that would increase the probability of inspecting each plant by an average of 10 percentage points. We explore how this increase would add value. We consider allocations of inspections under four different scenarios. In the base case, all additional inspections are not targeted based on scores. In the other three cases, inspections are targeted based on plant scores, based on firm scores, and a 50/50 mix of plant and firm scores.

We show that linked escalations perform best, followed by a 50/50 mix, then unlinked escalations, with untargeted inspections worst. Our main results focus on the effects on multi-plant firms, which are the primary target of linked regulation. For these firms, a 10 percentage point increase in the budget devoted to unlinked escalations lowers perceived social costs of pollution by 55.3% compared to a 10 percentage point increase in untargeted inspections. Linked escalations do substantially better, lowering perceived social costs of pollution by 76.2% compared to untargeted inspections. Our results are robust to using more direct measures of pollution, rather than using the regulator's estimated "perceived" social costs.

Decomposing the total effects into the two channels, we find that the "correlated targeting" effect is significant, and reduces the perceived social costs of pollution by 8.7%. However, a large share of the benefits to linked regulation is due to the "firm-wide moral hazard" channel which reduces the perceived social costs of pollution by 67.5%.

We also provide results including single-plant firms and show that the overall results are not due to a reallocation of enforcement resources from single-plant to multi-plant firms. Additional counterfactuals evaluate the efficiency of random crackdowns and alternative scoring rules that record longer histories of violations.

1.1. *Contributions*

We make three main contributions in this paper. First, we construct a novel dataset of compliance scores, regulatory inspections, violations, penalties, and pollution, for the 7,379 plants regulated under the Resource Conservation and Recovery Act or the Clean Water Act in Texas 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 estimate the model without having to infer the score evolution from inspection and violation histories, which would require strong assumptions to maintain tractability.

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 dynamic discrete choice literature. While other work has studied dynamic escalation mechanisms based on the compliance history of a single plant over time (*e.g.* Blundell *et al*., 2020), our framework allows for large portfolios of plants and for choices to be interdependent within the portfolio of plants and across time.

The third contribution is a set of new findings about the efficiency 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 untargeted inspections.

As well as the previously discussed cases where linked regulation is used, there are also numerous examples where firm-based risk scoring is *not* used to regulate multi-plant firms but where implementing it could add substantial value. A prime example is where firm ownership crosses state borders. For instance, expanding linked regulation by using information on co-owned plants outside Texas could potentially improve Texas's scoring regime. More generally, in the US, firms have polluting facilities located across 10 EPA regions and 50 states, but enforcement is rarely coordinated, as has been studied extensively in the literature on environmental federalism (Millimet, 2014). Overall, the results in this paper illustrate another downside of delegating legislation like the CWA and RCRA to states: the lost precision in targeting enforcement resources due to dispersed information about firm types compared to a more centralized linked regime.

1.2. *Related literature*

This paper is related to several strands of literature. The first is the literature that empirically estimates structural models of environmental regulation, to study alternative policy designs. Closest to this paper is Blundell *et al*. (2020), who 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. Our model nests the idea that single-plant escalations can add value, but the major distinction of our paper is that we focus on designing regulation that incorporates dynamic linkages that leverage correlated types within commonly-owned portfolios.[4]

The second strand is the broader literature on environmental regulation *e.g.* Colmer *et al*. (2025), Shimshack and Ward (2022), and Blundell *et al*. (2021). Both Gibson (2019) and Rijal and Khanna (2020) document pollution substitution within multi-plant firms from plants regulated under the Clean Air Act. This channel would tend to result in additional benefits to linked regulation. However, we do not model this channel in our application which focuses on different regulation; instead, we argue that reallocation does not seem to be empirically important in Online Appendix 3.1.

The third 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. Scur *et al*. (2021) provide a non-technical summary of an influential series of papers in this literature. Giardili *et al*. (2023), Bloom *et al*. (2019), and Goldfarb and Xiao (2011) document 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 ownership linkages across plants.

Finally, this paper is related to the theory literature on deterrence. This literature includes papers about the optimal design of deterrence mechanisms such as (Mookherjee and Png, 1994; Polinsky and Shavell, 1998). Many of these papers find that escalation mechanisms—whereby the marginal penalty faced by the most egregious or repeat offenders is higher—are optimal. The reasons for escalation mechanisms are varied and include limited liability constraints (a feature of Mookherjee and Png, 1994), costly enforcement (as in Polinsky and Shavell, 1998), and a positive relationship between agents' cost of compliance and the responsiveness to greater

4. There are also some other key differences. For example, our model views firms as facing a hidden action/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.

penalties. There is also some theoretical work on how to optimally inspect in the presence of firm dynamics, *e.g.* Varas *et al*. (2020).

2. 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 (RCRA) 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 laws are federal, the Environmental Protection Agency delegates enforcement to states. In Texas, the Texas Commission on Environmental Quality (TCEQ) is the primary governing body that administers these regulations.

2.1. *Details about regulated agents, violations, and inspections*

2.1.1. What is a plant? We use the term "plant" to describe what the regulator calls "sites": specific geographical locations where regulated behaviour takes place. We define a site according to the TCEQ's "RN" identifier. In our data, examples of plants are dry cleaners, gas stations, and convenience stores. For example, a gas station may have an underground petroleum tank which would be regulated under RCRA. There is substantial overlap between facilities that are regulated under RCRA and Clean Water Act (and the regulator aggregates violations under both laws to determine enforcement decisions) so we consider plants regulated under both laws together.

We proxy for observable plant heterogeneity using the industry according to its NAICS code (concretely, we classify plants into utility, services, manufacturing, resources, transportation, and trade). To foreshadow, later in the paper and the model we also incorporate unobserved plant heterogeneity by allowing for different compliance costs within a sector which are private information to firms. One may ask whether plants might be differentiated on other observable dimensions (for example Blundell *et al*., 2020 allow for plant types to depend on the geographical region). However, in Section 3.2 we empirically investigate this by testing for whether geographical regions and other potentially observable characteristics are relevant to how the regulator targets inspections. Ultimately, we do not find that the regulator uses these characteristics to target regulation.

2.1.2. What is a firm? A firm is the owner of a portfolio of plants and we define a firm using the TCEQ's "CN" identifier. Sometimes the "firm" in our data is just the name of the person who owns the business; when this is not the case we often see an individual person responsible for compliance in the firm listed across the permits for the plants. As a result, we view the organizational structure of a typical firm as quite simple: they have a flat structure and operate a collection of plants with different compliance costs.[5] Note that there is substantial heterogeneity in these portfolios of plants. For example, firm portfolios can contain varying numbers of plants, and within a portfolio, plants can differ in their industry sector as well as their (unobserved) plant-specific compliance costs.

5. Consistent with this, the structure of a "firm" is quite simple in our model as well, abstracting away from delegation, etc.

2.1.3. What is an inspection? Inspections are typically on-site, although it is possible for the regulator to conduct inspections virtually. Inspections can be untargeted (*i.e.* a plant with a clean record may still be inspected) or through the history of violations at the plant or co-owned plants. The regulator also allows for inspections to be triggered by public complaints, but the regulator allocates <5% of its total inspections as a response to complaints (TCEQ, 2022). We therefore abstract away from this channel later in the model. One may ask whether distance/remoteness plays a role in which plants get inspected. However, the TCEQ structure guards against this: it is split into 16 operating regions with inspectors drawn from local offices, so inspectors do not need to travel great distances.

2.1.4. What is a violation? Most of the violations of the regulations that we study are "process violations" such as deviations from the correct management process for storing toxic waste. As a result, we often do not see in the data an associated specific environmental harm with each violation (like a spill, or a leak into groundwater). Rather, the regulator penalizes and corrects these violations because they run a potential risk of a high-magnitude environmental harm event occurring.

These violations usually cannot be remedied through, for example, a technological investment. Instead, firms must maintain continuous effort and attention to ensure that these management practices are followed properly. We proxy for the gravity of bad management practices through multiple violations discovered at the same facility during an inspection; the rationale behind this is that empirically total penalties in the data scale approximately linearly in the number of detected violations, as we document in the Online Appendix in Table A-13.

The regulator can also detect violations by means other than inspections, for example through continuous monitoring or self-reported violations (see Shimshack and Ward (2022) for more details about these violations). We allow for this possibility in the model and therefore the overall results.

2.2 *How does Texas perform environmental enforcement?*

Texas 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"*.

The regulator determines the "risk" of each plant through a two-dimensional scoring system based on the compliance history of each plant and firm. 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 are owned by the same firm.

2.2.1. Plant score The plant score combines the past five years of violations and inspections at a particular site into a one-dimensional index. We detail the exact scoring algorithm that the regulator uses in Online Appendix Section 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. Note that older violations are discounted and erased from the score after five years. Therefore, if no violations are detected at a plant then over time the score eventually resets to 0.0.[6]

2.2.2. Complexity points The regulator normalizes plant scores by "complexity points" when making enforcement decisions. Adjustments for complexity points were introduced because there is heterogeneity in terms of, for example, the sector to which the plant belongs. Complexity points allow different types of plants to be compared under the same rating system. For example, a facility in one sector may mechanically generate more violations than another sector; adjusting by complexity points allows for the scores to be comparable.

2.2.3. 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 part of firm *f*. Denote s_j as the score of plant *j* and Q_j as the corresponding complexity points of plant *j*. Then, the firm-wide score is computed as follows:

$$s_f = \frac{\sum_{j\in\mathcal{J}_f} Q_j s_j}{\sum_{j\in\mathcal{J}_f} Q_j} \tag{1}$$

Note that for a single-plant firm, the firm score will be equal to the plant score by construction. However, complexity points are still incorporated, since violations are pre-adjusted by dividing by complexity points before they enter into the plant score. This is illustrated in the TCEQ's algorithm in Online Appendix Section 1.

2.2.4. Inspections and penalties The regulator escalates inspections and penalties based on the scores s_j and s_f. We discuss how inspections escalate with scores in Section 3. The penalties are escalated using the following rule:

$$\text{Penalty} = \text{Base Penalty} \times \text{Plant Escalation} \times \text{Firm Escalation} \tag{2}$$

where the Plant Escalation is a function of s_j and the Firm Escalation is a function of s_f. The base penalty depends on the gravity of the violation. After determining the base penalty, the total penalty is then escalated based on the compliance scores.[7] 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 across many plants and so can potentially be quite costly to the firm, resulting in large deterrence effects. The plant escalations are also increasing in the scores.

2.3 *Data*

Our data are an unbalanced panel of the universe of plants regulated in Texas under RCRA and the Clean Water Act from 2012–2020. There are 7,379 plants across 5,345 firms in our dataset. We observe extremely detailed information about each plant, including the firm, each

6. Although the focus of the scores is past violations, the scoring rule detailed in Online Appendix Section 1 shows that—in principle—scores can also be affected by other considerations like whether a plant "participated in a self-audit". In practice, as we explain further in Online Appendix Section 1.1, it is extremely rare for firms in our dataset to take these actions. As a result, we abstract away from their effect on scores later in our modelling framework to keep the empirical model tractable.

7. The firm escalations are: Firm Escalation $= 0.9$ if the firm is a "high performer" ($s_f \in [0, 0.1)$), Firm Escalation $= 1.0$ if the firm is a "satisfactory performer" ($s_f \in [0.1, 55)$) and Firm Escalation $= 1.1$ if the firm is an "unsatisfactory performer" ($s_f \in [55, \infty)$).

TABLE 1
Summary statistics

Variable	N	Mean	Std. dev.	Percentiles 1	Percentiles 99
Log(1 + plant score)	50,864	0.66	1.04	0	4.23
Log(1 + firm score)	50,864	0.79	1.01	0	4.19
Inspection	50,864	0.31	0.46	0	1
# Violations	50,864	0.71	3.10	0	16
# Plants in portfolio	50,864	3.25	5.89	1	40
Env. justice score	50,864	0.44	0.18	0.09	0.87

environmental inspection, whether the inspection detected a violation, the nature of the violation, and the monetary penalty that was incurred. We map violations to pollution using an index; we detail the steps used to create this index in Online Appendix 2.2. For the analysis, we aggregate the data to the year-plant level.[8]

Table 1 contains information about key summary statistics in the data. Due to the long right tail of plant scores and firm-wide 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.

Overall, the probability that the average plant is inspected each year is 0.31. The fact that inspections occur infrequently is consistent with inspections being costly to the regulator. It is quite common for these inspections to uncover violations: on average, around 2 violations are discovered for each inspection. Note that a very small percentage (3.5 percent) of inspections are conducted by the federal EPA rather than the TCEQ. We choose to retain these inspections in the dataset, since these inspections may be relevant to firms' actions. We discuss in greater detail in Online Appendix 3.3 that this decision has minimal effects on the descriptive results, estimation, and counterfactual findings.

The average portfolio size is 3.25.[9] This average masks substantial heterogeneity in portfolio size: while the largest portfolio consists of 43 plants, single-plant firms are also present in the data.

Finally, the environmental justice index is an EPA measure that averages a census block's percent people of colour and low-income percentage in the census tract where the plant is located.

3. DESCRIPTIVE ANALYSIS

We now provide several pieces of descriptive evidence illustrating how linked regulation operates in practice in our setting. First, commonly-owned plants' violations per inspection are correlated, suggesting that costs of compliance with regulations (*i.e.* plant "types") are correlated. Second, regulators selectively target their inspections towards bad actors. Third, we show evidence of moral hazard: firms respond to increased regulation by decreasing their pollution.

3.1 *Violations are correlated within commonly-managed plants*

In Figure 1 we show that violations per inspection of a given plant are correlated with violations per inspection of other plants of the same firm. The horizontal axis of the figure is the logarithm

8. Scores are updated in November, so we define the start/end to each "year" as November 1st.
9. Note that the percentage of plants in multi-plant firms is 41.2%. Furthermore, we observe multi-plant firms for 7.1 years on average, and single-plant firms for 6.7 years on average.

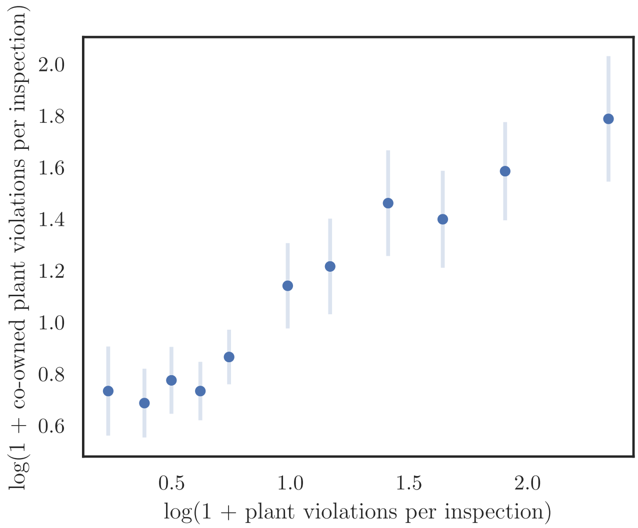

FIGURE 1
Binscatter of violations per inspection at one plant vs other co-owned plants

Note: 95% confidence intervals plotted around each bin

of the average number of violations per inspection at plant *j* over the course of the panel. The vertical axis shows the logarithm of 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 (note that the exact correlation is 0.504). As we show in Table A-14 of the Online Appendix, this positive correlation is robust to controls for observable plant characteristics. Actions are endogenous, so we would ideally show that plants' pollution mitigation costs are correlated with other co-owned plants' costs, but these costs are not observed. The correlation in costs is likely to be *higher* than the correlation in actions under this 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 between violations at plant A and plant B.

3.2 *The regulator targets inspections based on linkages*

Next, we show that inspections are targeted towards bad actors. 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. We condition on year and NAICS sector fixed effects. Plants with higher scores, reflecting more past violations, are more likely to be inspected. This is true of both plant-level and firm-wide scores. Specification (1) implies the average plant with a clean record is inspected 28% of the time. For the same plant with a "bad" record (at the top percentile of firm and plant scores) the probability of inspection is escalated to 44%.

The magnitude of these effects does not change substantially when we include the environmental justice index for each plant (specification (2)) or when we include fixed effects for each of the sixteen enforcement regions of TCEQ (specification (3)). Though the log-likelihood is higher, by construction, when we include these additional controls, the gain is small. We therefore use the more parsimonious specification (1) as our main specification of the inspector's policy function.

TABLE 2
Inspection probability regressions

Dependent Variable	(1) Inspection	(2) Inspection	(3) Inspection
Log(1 + firm score)	0.056	0.056	0.049
	(0.018)	(0.018)	(0.018)
Log(1 + plant score)	0.109	0.110	0.111
	(0.017)	(0.017)	(0.017)
Env. justice index	–	0.088	–
	(–)	(0.054)	(–)
Year FEs	Yes	Yes	Yes
NAICS category FEs	Yes	Yes	Yes
Region FEs	No	No	Yes
N	50,864	50,864	50,864
Log-likelihood	−30,706	−30,705	−30,545

TABLE 3
Deterrence regressions

Dependent Variable	(1) Violations	(2) Violations	(3) Violations
Pr(Inspection)	−0.246	−3.079	−4.229
	(0.703)	(0.765)	(0.750)
Year FEs	Yes	Yes	Yes
Unit FEs	Category, Firm	Plant	Plant
Only Inspected	No	No	Yes
N	38,869	33,527	10,531

3.3 *Evidence of moral hazard and deterrence*

Next, we demonstrate that higher scores have the desired effect of deterring plants from polluting. In Table 3 we report estimates from regressions where the dependent variable is the number of violations uncovered at a plant and the main independent variable is the predicted probability of an inspection occurring as determined by specification (1) of Table 2. Specification (1) of Table 3 shows that the probability of inspection does not significantly predict changes in violations. However, this reflects two opposing effects: on one hand there is a sorting effect where plants with higher compliance costs have both higher scores and violations, while on the other hand, higher scores deter plants from violating. Specification (2) eliminates the first effect by including plant fixed effects. When a plant's score increases, its violations tend to decrease, consistent with higher scores having a deterrence effect. The estimate in (2) is still attenuated by the fact that plants with higher scores are more likely to be inspected and therefore are more likely to have violations *uncovered*. 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 in magnitude. In Table A-12 of the Online Appendix we also show plants that have high violations, on average, are more responsive to increases in the probability of inspection. This supports the assumption that—even in a dynamic setting—targeting bad actors is useful because they are more responsive to regulation.

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 plant-specific score. The results are plotted in Figure 2. Because the plant

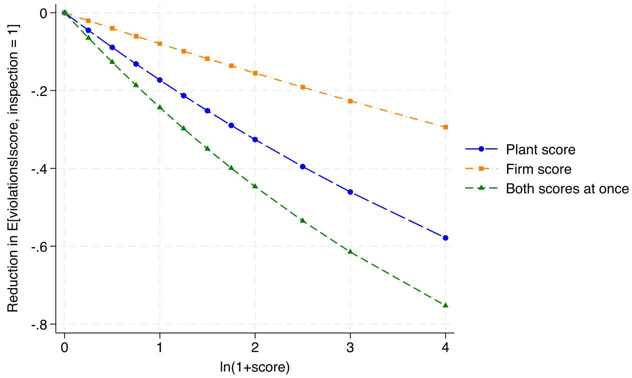

FIGURE 2
Implied deterrence effect of scores

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

score has a larger estimated effect on the probability of being inspected, the deterrence effect of plant scores is larger. Both, however, yield strong deterrence effects. An increase in scores from a "good" plant with a clean record at both the plant and firm level, to a "bad" plant with a plant score and firm score in the top percentile, decreases the expected number of violations by roughly 0.8.

Although these results suggest that a combination of dynamic scoring and targeted inspections deters plants from polluting, we cannot ascertain a causal relationship from this analysis. For example, violations are only recorded if a plant is inspected, the expectation of inspection implied by scores drives a plant's actions, and inspections are targeted based on scores. These relationships become even more complicated for multi-plant firms, where firms dynamically trade off the costs of mitigation at one plant versus the effects of future harsher regulation across several plants. In order to disentangle all of these effects, as well as to evaluate the welfare implications of different regulatory schemes, we estimate a model that explicitly accounts for the objectives of firms and regulators.

3.3.1. Robustness In the Online Appendix we perform two robustness tests for alternative ways that firms might respond to the regulation. First, in Online Appendix 3.2 we test for strategic avoidance. Specifically, we test the hypothesis that a devious firm might use the knowledge of what an inspector looked for in one plant to somehow hide violations at other plants. We focus on firms that have plants in multiple TCEQ regions, and we exploit that TCEQ inspectors are drawn from the local regional office, generating plausible variation in the identities of who is inspecting each plant. Ultimately, we find no evidence of strategic avoidance. Second, we test for production reallocation in Online Appendix Section 3.1 but we also find no evidence for this.

4. MODEL

4.1 *Outline of model*

Our model is a discrete-time dynamic model, with time indexed by t.

4.1.1. Agents There is a set of firms, $\mathcal{F}$, and a regulator. Each firm owns a portfolio of plants $\mathcal{J}_f$, with each plant indexed by j. Each plant belongs to industry $g(j)$ and is assigned a type,

θ_j, from joint distribution $G_\theta(.)$ which can include correlated types by firm and industry. The parameter θ_j affects a plant's marginal cost of abatement or, equivalently, its marginal benefit of polluting.

4.1.2. Actions The regulator chooses an inspection policy at time $t = 0$ which is fixed over time (committed) and public.[10] We model inspections z_{jt} as a draw from a Bernoulli distribution, and the regulator chooses the Bernoulli parameter $\bar{z}_{g(j)}(s_{jt}, s_{ft})$. Here, $g(j)$ is the firm industry, s_{ft} is the firm score, and s_{jt} is the plant score. The Bernoulli draws z_{jt} are independent across plants conditional on industries and scores.

At time t, each firm firm chooses a negligence action a_{jt} for each plant in its portfolio. Higher actions incur higher pollution and, conditional on plant j being inspected, more violations v_{jt}.

4.1.3. Information structure The inspection policy $\bar{z}_{g(j)}(s_{jt}, s_{ft})$ is common knowledge. Violations are common knowledge once they are uncovered. Scores, and the updating rule for scores, are also common knowledge. Plant types are not observed by the regulator.[11] Because the scores are common knowledge, and the inspection policy is committed, there is no regulatory hold-up or moral hazard on the regulator's side, and therefore no double moral hazard as in, *e.g.* Bhattacharyya and Lafontaine (1995).

4.1.4. Timing Since scores are updated once a year, in our model one period is one year. We summarize the timing of the model below and in the diagram in Figure 3.

$t = 0$	Nature draws private plant types, θ_j, for each j. Regulator chooses public inspection policy $\bar{z}_{g(j)}(s_{jt}, s_{ft})$.
$t = 1, 2, \ldots$	
(i)	Plants enter period t with score s_{jt} and firm score s_{ft}.
(ii)	Firms choose pollution levels a_{jt} for each j in $\mathcal{J}_f$.
(iii)	Regulator inspects each plant independently with Bernoulli probability $\bar{z}_{g(j)}(s_{jt}, s_{ft})$. Violations are drawn $v_{jt} \sim \text{Poisson}(a_{jt})$, and they are uncovered with probability 1 if an inspection occurs and with probability γ if no inspection occurs.[12] Denote detected violations by $\hat{v}_{jt}$.
(iv)	Penalties are assessed. These penalties scale with the number of detected violations.[13] Total penalties per violation are given by $x(s_{jt}, s_{ft})$.
(v)	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 points of each plant: $s_{f,t+1} = \sum_{j \in \mathcal{J}_f} Q_{g(j)} s_{j,t+1} / \sum_{j \in \mathcal{J}_f} Q_{g(j)}$). The updating and aggregation rules are common knowledge.

We next characterize the payoffs governing the two key decisions in the game: the firms' choice of a_{jt} for their plants, and the regulator's design of $\bar{z}_{g(j)}(s_{jt}, s_{ft})$. We begin our discussion with the periods where, given $\bar{z}_{g(j)}(s_{jt}, s_{ft})$, firms choose a_{jt}.

10. In the empirical application of this paper, it is arguably reasonable to assume that the regulator can commit to a policy. For example, the scoring rule, penalties, and many other components are codified in the Texas Administrative Code (2002) and are unchanging over the period of our study.

11. Equivalently, plant types and actions could be observed by the regulator, but the regulator cannot contract directly on plant types and actions and must use scores instead.

12. This "alternative detection" can include self-reported violations and continuous monitoring.

13. Online Appendix Table A-13 shows that this linear scaling is a reasonable assumption.

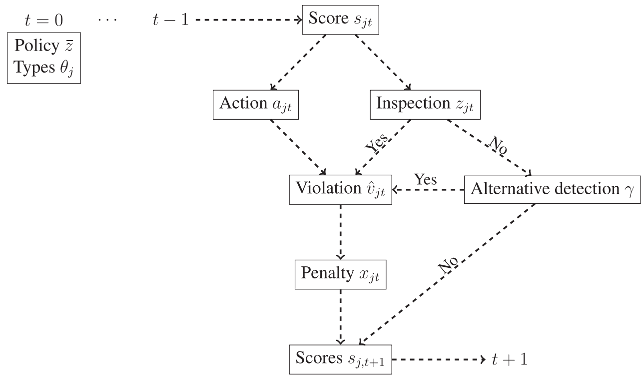

FIGURE 3
Timing of model

4.2 *The firm's problem*

Firms are individually small compared to the overall set of 5,345 firms in our dataset, with no firm accounting for $>1\%$ of total violations in our dataset. Similarly, no firm accounts for $>0.8\%$ of total inspections in our dataset. Furthermore, the inspection policy only depends on the scores and characteristics of the firm's own plants. Therefore, there is no strategic interaction across firms. Because each firm's problem can be considered independently of the other firms, we drop firm-specific subscripts in this section. We also drop time t subscripts for notational simplicity. The expected flow payoff of plant j, given action a_j, with the expectation taken over stochastic violations and inspections, is:

$$\pi_j(a_j;s_j,s_f)=\underbrace{\theta_j b(a_j)}_{\substack{\text{Flow benefit}\\ \text{from negligent}\\ \text{actions}}}-(\underbrace{\bar{z}_{g(j)}(s_j,s_f)}_{\text{Inspection}}+\underbrace{\gamma\,(1-\bar{z}_{g(j)}(s_j,s_f)))}_{\substack{\text{Alternative methods}\\ \text{to detect violations}}}$$
$$\times\underbrace{\mathbb{E}_{v_j}(v_j x(s_j,s_f)\mid s_j,s_f,a_j)}_{\substack{\text{Expected penalty}\\ (=a_j x(s_j,s_f))}} \qquad (3)$$

The first component of equation (3) 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 (3) comprises the costs to the firm of taking a negligent action a_j. Here, the regulator first inspects the plant with probability $\bar{z}_{g(j)}(s_j,s_f)$. Violations may also be detected by alternative methods than inspections (for example, by an automatic monitoring technology), and this occurs as a draw from a Bernouilli distribution with parameter γ. If the plant is inspected, or the alternative detection draw occurs, then the regulator will observe the count of violations v_j (which is drawn from a Poisson distribution with mean a_j). Finally, after observing violations, the regulator administers penalties per violation $x(s_j,s_f)$. Note that $\mathbb{E}_{v_j}(v_j x(s_j,s_f)\mid s_j,s_f,a_j)=a_j x(s_j,s_f)$.

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.* $\mathbf{a} = \{a_j\}_{j \in \mathcal{J}}$, and the state in the next period as $\mathbf{s}'$, the firms' problem is:

$$V(\mathbf{s}) = \max_{\mathbf{a}} \sum_{j \in \mathcal{J}} \pi_j(a_j; s_j, s_f) + \beta \mathbb{E}_{\mathbf{z},\mathbf{v}} \left[V(\mathbf{s}') \mid \mathbf{s}, \mathbf{a} \right] \tag{4}$$

4.2.1. Discussion of assumptions The above setup contains two implicit assumptions. The first is that we assume that flow payoffs are linked through correlated types and linked regulation, but are otherwise independent. In particular, we are ruling out linkages that occur through production shifting within a firm when one plant is under more scrutiny (like in Gibson (2019) and Rijal and Khanna (2020) which study the Clean Air Act or Chen *et al.* (2025) for conglomerates in China). As previously discussed, we run robustness tests for reallocation in Online Appendix 3.1 and do not find any evidence. One potential reason for this result is that plants in our application (such as gas stations) are smaller than those studied in these papers and so do not have the ability to easily scale production.[14] We also test for whether firms strategically avoid the regulation in other ways in the Online Appendix, and similarly do not find evidence of this.

Second, we assume that firms internalize the effects of their actions on co-owned plants. This assumption may be questionable if firms have complex organizational structures (*e.g.* regional managers at different plants). However, as we mention in the context section of the paper, the typical firm in our setting has quite a flat structure, usually with an individual person as an owner/operator or responsible for environmental compliance.

4.2.2. Simplifying the state space The portfolio problem defined in equation (4) is extremely complex to solve. 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 because the regulation is linked. This generates a curse of dimensionality: the firm manager needs to think about the entire state space s which is of dimension $n_{plant} + 1$.[15] In our application many firms have large portfolios (that is, n_{plant} can be large), so solving equation (4) is essentially infeasible. Therefore, we need to employ an alternative approach. Our approach begins from the observation that the optimal actions a_j^* that satisfy equation (4) are equivalent to the optimal actions that jointly satisfy the following n_{plant} equations (one for each plant):[16]

$$\max_{a_j} \underbrace{\pi_j(a_j; s_j, s_f)}_{\text{Plant flow payoff}} + \beta \mathbb{E}_{\mathbf{z},\mathbf{v}} \left[\underbrace{V_j(\mathbf{s}')}_{\substack{\text{Plant} \\ \text{continuation} \\ \text{value}}} + \underbrace{\sum_{k \in \mathcal{J}/j} V_k(\mathbf{s}')}_{\substack{\text{Other plants'} \\ \text{continuation values}}} \mid \mathbf{s}, \mathbf{a}^*_{-j}, a_j \right] \tag{5}$$

14. Even if they could, some degree of linked regulation would tend to guard against this behaviour since scrutiny increases at all plants with a common owner, and so our results about the efficacy of linked regulation would be conservative.

15. 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.

16. This can be verified by observing that the first order conditions resulting from equations (4) and (5) are identical.

where $V_j(\mathbf{s})$ denotes the value function for an individual plant j (*i.e.* the lifetime discounted stream of flow 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 (5) highlights where the state space complexity bites in the portfolio choice problem. Concretely, it is through the third term in equation (5) which governs how firm j's action affects the future payoffs of other commonly-owned plants, *i.e.* cross-plant effects. 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 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.

4.2.3. Defining the resulting firm-level problem The above assumptions turn the computationally intractable problem in equation (5) 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 Online Appendix 4.1—the following three equations fully characterize the firm's problem. First, the optimal action $a_j^*(\hat{\mathbf{s}}_j)$ is given by:

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

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

$$V_j(\hat{\mathbf{s}}_j) = \pi_j(a_j^*(\hat{\mathbf{s}}_j); s_j, s_f) + \beta\mathbb{E}_{z_j,\upsilon_j}\left[V_j(\hat{\mathbf{s}}_j') \mid \hat{\mathbf{s}}_j, a_j\right] \tag{7}$$

and the state transition beliefs are governed by the AR(1) process (where $\hat{\upsilon}_j$ denotes detected violations that are normalized by the complexity score of plant j):

$$\hat{\mathbf{s}}_j' = R_{0,j} + R_{1,j} \times [s_j, s_f, W_j, \hat{\upsilon}_j]' \tag{8}$$

Note the subscripts in equation (8) 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 equilibrium behaviour across the larger portfolio of plants for firm f, as well as the updating rules of the scores defined by the regulator. We detail the full solution algorithm—which comprises an outer loop where we solve for the beliefs in equation (8) and an inner loop where we compute actions and value functions that satisfy equations (6) and (7)—in the Online Appendix. Note that the firm-level problem needs to be computed separately for each individual firm that we observe in the data, since each firm has a different portfolio of plants.

4.2.4. Equilibrium: firms Define a steady state equilibrium as the joint distribution of actions and scores given an inspection policy, which we shorten to $\bar{z}$, and a vector of types for each plant $\boldsymbol{\theta}$, $F(\mathbf{a},\mathbf{s};\bar{z},\boldsymbol{\theta})$, that satisfies the following conditions:

1. Each firm chooses actions optimally according to equations (6) and (7).
2. Beliefs about state transitions in equation (8) are consistent with the equilibrium.
3. Given the law of motion for scores, the regulator's policy function, and the optimal actions at each state, the distribution of scores is in steady state.

4.3 *The regulator's problem*

At $t = 0$, the regulator commits to a stationary and observable inspection rule, $\bar{z}$, from a set of allowable rules $\mathcal{Z}$ to minimize the "perceived" social cost of equilibrium violations subject to an inspection budget B and an implementability constraint that ensures violations are part of a steady state equilibrium for $\bar{z}$. We expand on how we parameterize the inspection policy function—and therefore the set $\mathcal{Z}$—in the estimation section, but it might, for example, feature inspections that do not depend on plant scores ("untargeted regulation"), inspections that are responsive to individual plant scores ("unlinked regulation"), or inspections that are responsive to firm scores ("linked regulation"). The perceived social harm of a violation, h_g, can vary across sectors. $\mathcal{Z}$ also allows inspection rules to depend on sectors. The regulator does not know the type of each plant but does know the distribution of plant types $G_{\boldsymbol{\theta}}$. Formally, the regulator's problem is:

$$\begin{aligned}
&\min_{\bar{z}\in\mathcal{Z}} V^R = \min_{\bar{z}\in\mathcal{Z}} \int\int \sum_{f\in\mathcal{F}} \sum_{j\in\mathcal{J}_f} h_{g(j)} a_j \, dF(\mathbf{a},\mathbf{s};\bar{z},\boldsymbol{\theta})\, dG_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \\
&\text{s.t.: } \int\int \sum_{f\in\mathcal{F}} \sum_{j\in\mathcal{J}_f} \bar{z}_{g(j)}(s_j, s_f)\, dF(\mathbf{a},\mathbf{s};\bar{z},\boldsymbol{\theta})\, dG_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \leq B \quad \text{[Budget balance]} \\
&F(\mathbf{a},\mathbf{s};\bar{z},\boldsymbol{\theta}) \text{ from equilibrium firm behaviour in Section 4.2.4} \quad \text{[Implementability]}
\end{aligned} \tag{9}$$

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

4.3.1. Discussion of assumptions We now justify three assumptions we are making in the regulator's problem in equation (9). 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. We verify robustness of our results to this assumption in Online Appendix Section 5.6. 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.[17] 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 (9).

The second implicit assumption is that we are taking the penalty structure as given and we are not taking a stand on whether the penalty schedule is chosen optimally by the regulator. We

17. Of course, since compliance costs are private information, there is still dispersion in compliance costs even after normalizing for complexity points.

choose to not do this for two reasons. First, while it is clear that the TCEQ has short-run discretion over who to inspect (as documented by their previously discussed "risk-based approach" to targeting), penalties are much more rigid. Keeping the penalty policy fixed reflects this inflexibility. For example, penalties are explicitly spelled out in Texas Administrative Code (2002). Although these rules can be modified by the TCEQ in the long-run, these rules are likely to be inflexible in the short-run since changes may be subject to a more elaborate rule-making procedure.[18] The second reason for not explicitly modelling optimal penalties is that the inspection moments—which result from assumptions that the regulator spends their inspection budget optimally across industries—are sufficient for identification. Therefore, in terms of estimating the model, we do not need to make additional assumptions on how these more rigid penalties are determined.

The third assumption is that the perceived social harms of violations vary across industry sectors, but not other plant characteristics such as location or environmental justice considerations. This is broadly consistent with the descriptive findings in Table 2, which showed that the environmental justice score and region fixed effects had little predictive power for the regulator's inspection decisions. Still, we include results in the counterfactuals where social costs are based on pollution measures instead of inferred from the regulator's behaviour as robustness to this assumption.

Finally, we explain how assumptions on the regulator's problem in this section affect the estimates and the counterfactual results. Overall, we only use the regulator's problem to estimate parameters relating to the perceived social cost of violations $h_{g(j)}$. These social costs are then used in the counterfactual section to aggregate violations from heterogeneous industry sectors into a single social cost number. Therefore, to illustrate robustness in the counterfactuals section, we also provide disaggregated results by industry sector in Online Appendix 5.3; these results are identified, estimated, and computed, independently from any assumptions on how the regulator makes decisions. In addition, in the counterfactual results, we also use an alternative measure constructed directly using pollution data.

5. ESTIMATION AND IDENTIFICATION

5.1 *Parameterization*

5.1.1. Plant types and payoffs 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. 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. This multiplicative-in-types specification implies increasing differences and is standard in the literature on marginal deterrence (Mookherjee and Png, 1994; Kang and Silveira, 2021).

The specification $b(a) = a^y$ implies that marginal static deterrence—the responsiveness of a myopic firm's optimal action to an increase in z—is larger in magnitude for high types than for low types. This is supported by the regressions in Online Appendix Table A-12, which shows that plants that violate more are more responsive to regulation. We also provide further discussion in

18. A useful guide to these procedures is Texas Environmental Public Participation Guide (2017). Note that we do not see any changes to the penalty rule in our sample period 2012–2020.

Online Appendix 4.6 about why this may imply that it is optimal to target inspections on firms with a worse history of violations.

5.1.2. Law of motion for scores Firms make decisions based on a reduced state space, with an estimated AR(1) process given by equation (8). The plant-specific transition matrices $R_{0,j}$, $R_{1,j}$ in this law of motion are given by:

$$R_{0,j} = \begin{bmatrix} 0 \\ 0 \\ r_0^{w,j} \end{bmatrix}, \quad R_{1,j} = \begin{bmatrix} r_{s_j}^{s_j} & 0 & 0 & r_v^{s_j} \\ 0 & r_{s_f}^{s_f,j} & 0 & r_v^{s_f,j} \\ 0 & 0 & r_w^{w,j} & r_v^{w,j} \end{bmatrix} \tag{10}$$

We set certain coefficients in these matrices to 0. First, we set the intercepts for the plant score updating and the firm score updating equal to 0. This restriction follows directly from how scores work in our setting: if no violations are observed then the score will converge to 0.[19] The regulator only updates the plant score s_j using plant-specific information. Therefore, we switch the dependence of s_j on other states in $R_{1,j}$ to 0. An additional implication of using only plant-specific information to update this score is that the remaining parameters in the first row in $R_{1,j}$ can be computed directly from the data.

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 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).

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

$$\bar{z}_{g(j)}(s_{jt}, s_{ft}) = \frac{\exp\left\{\rho_{0g(j)}^z + \rho_1^z \log(1+s_{jt}) + \rho_2^z \log(1+s_{ft})\right\}}{1+\exp\left\{\rho_{0g(j)}^z + \rho_1^z \log(1+s_{jt}) + \rho_2^z \log(1+s_{ft})\right\}} \tag{11}$$

When solving the regulator's problem, we limit $\mathcal{Z}$ to the family of logit functions in the above form.

5.1.4. Penalties Penalties are assessed according to a base penalty, plant-level scores, and firm-level scores, as defined in equation (2). We calibrate the escalations as a function of the states directly from the regulator's bylaws. We normalize the base penalty per violation to 1.0, which is without loss of generality when computing the firms' problem because we do not normalize the scale of the plant types θ.

19. The AR(1) process we estimate closely resembles the score updating process in the Texas Administrative Code, 30 Tx. Admin. Code § 60.2: "The total number of points assigned for all resolved violations [...] over two years old will be multiplied by 0.75 [...] over three years old will be multiplied by 0.5 [...] over four years old will be multiplied by 0.25".

5.2 *Estimation routine*

The primitives we estimate correspond to: the inspection function ρ^z, the complexity weights Q_g, the state updating process $R_{0,j}$ and $R_{1,j}$, the distribution of firm types σ_F, σ_J, and $\bar{\theta}_g$ for each industry sector g, the curvature parameter y, and the marginal social costs of pollution by sector h_g. We estimate the model in three steps.

First step: We estimate/calibrate the following parameters "offline". We calibrate the discount parameter $\beta = 0.95$. We use the estimates for ρ^z from Column (1) of Table 2, including the NAICS sector fixed effects. We estimate the complexity weights using non-linear least squares, finding the Q_g that minimize the squared residual of the true firm score observed in the data compared to a predicted firm score built from the observed plant scores and Q_g using equation (1). We provide more details in Online Appendix 4.3. We compute the coefficient $r_{s_j}^{s_j}$ in the plant score updating rule directly from the data using a linear regression. We calibrate the parameter for how current violations affect future plant scores $r_v^{s_j}$ using the regulator's rule in Texas Administrative Code (2002). We explain how we calibrate the parameter γ (for violations discovered by means other than inspections) in Online Appendix Section 4.4.

Second step (firms' problem): We estimate the parameters governing the firms' problem $(\sigma_F, \sigma_J, \bar{\theta}_g, y)$ using simulated method of moments. We provide a short description of the algorithm here and leave a more detailed explanation of the algorithm—as well as the method of moments implementation—in Online Appendix 4.5. For each of the thousands of firm portfolios $\mathcal{J}_f$ we observe in the data we run an individual solution algorithm that contains two nested loops. In the inner loop we take the state updating process beliefs for each plant $j \in \mathcal{J}_f$ (*i.e.* $R_{0,j}$ and $R_{1,j}$) as given, and solve for a fixed point of the optimal actions and value functions at each state that satisfies equations (6) and (7) subject to the regulator's conditional choice probabilities. Then, we simulate sequences of "data" given the results from the inner loop, computing a time-series of optimal actions at each plant in the portfolio given the state, corresponding violations, inspections, and penalties. At the end of each period we update the scores and the state to the next period. We then run AR(1) regressions on this simulated "data", updating the state transition beliefs for each plant j over the reduced state space: $R_{0,j}$ and $R_{1,j}$. We continue the outer and inner loop procedures until these state transition beliefs converge.

Third step (regulator's problem): Finally, we estimate the social cost parameters h_g by leveraging the assumption that the regulator chooses the inspection policy function optimally. Our estimates of the firm parameters in Step 2 allow us to compute the stationary distribution of scores and actions $F(\mathbf{a}, \mathbf{s}; \bar{z}, \boldsymbol{\theta})$ for alternative inspection policies $\bar{z}$. Therefore, we can also compute the regulator's objective function V^R for alternative inspection policies. We show in Online Appendix 4.2 that h_g can be analytically computed up to a scale normalization (due to the budget constraint) by inverting the first-order-conditions of the regulator's optimal policy defined by equation (9). We approximate derivatives of V^R in the analytical formula via finite differences.

5.3 *Identification*

5.3.1. Second step (firms' problem) We informally discuss the moments selected and how the moments selected identify the parameters of the model that are recovered in step two 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, suppose that violations are deterministic: $v_j = a_j$, and denote z as the probability of inspection. 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

$$\log(a_j^*) = \frac{1}{1-y}\left(\log(\theta_j) - \log(z) + \log(y)\right). \tag{12}$$

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

Though our empirical model is more complicated, these comparative statics inform our choice of moments. To identify the average type for each NAICS sector, θ_g, we use the average number of (detected) violations per year, conditional on inspection occurring, for all plants in that sector. This yields six moments (one for each industry sector). 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 inspections occurring, we decompose the total variance of detected violations $\hat{v}$ as follows:

$$\mathrm{Var}[\hat{v}] = \underbrace{\mathbb{E}_j\left[\mathrm{Var}_t(\hat{v}_{jt} \mid j)\right]}_{\substack{\text{Within plant,}\\ \text{over time}}} + \underbrace{\mathbb{E}_f\left[\mathrm{Var}_j\left[\mathbb{E}_t(\hat{v}_{jt} \mid j) \mid j \in \mathcal{J}_f\right]\right]}_{\substack{\text{Across plants,}\\ \text{within firm}}} + \underbrace{\mathrm{Var}_f\left[\mathbb{E}_{j,t}(\hat{v}_{jt} \mid j \in \mathcal{J}_f)\right]}_{\text{Across firms}} \tag{13}$$

We divide the second term on the right hand side of this equation by the total variance $\mathrm{Var}[\hat{v}]$ to identify σ_J^2. With this identified, we also include the ratio of the second and third terms to identify σ_F^2. All else equal, higher types 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/across plants identifies σ_J^2, while the ratio of the within/across variance identifies σ_F^2.

As evident in equation (12), the curvature parameter y is identified by the responsiveness of violations to changes in inspection probabilities. We construct the relevant moment in two steps: first, we compute the predicted probability of inspection using our first-stage estimates. Next, we regress violations on this predicted inspection probability and plant fixed effects, conditional on inspection occurring. We compute the same moment analogously in our simulated data.

5.3.2. Third step (regulator's problem) The discussion in Online Appendix 4.2 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—which corresponds to a greater probability of an untargeted inspection that does not depend on plant or firm scores—the larger the relative implied harms of pollution from that sector. 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 inspection budget efficiently across sectors in equilibrium. This allows us to identify the *relative* marginal harms of pollution. Therefore, we normalize $h_{\text{utility}} = 1$ and estimate the marginal harms of pollution in other industries relative to this.

TABLE 4
Structural estimates

Parameter	Estimate	Std. Error	Parameter	Estimate	Std. Error
Mean type $\bar{\theta}_g$			*Perceived social cost h_g*		
Manufacturing	0.445	0.084	Manufacturing	0.984	0.037
Resources	0.019	0.04	Resources	1.076	0.043
Services	0.17	0.102	Services	0.945	0.043
Trade	0.023	0.035	Trade	1.248	0.086
Transportation	0.39	0.074	Transportation	1.561	0.275
Utility	0.45	0.084	Utility	1.0	–
Type variances					
Plant-level, σ_J^2	0.192	0.039			
Firm-level, σ_F^2	0.271	0.061			
Shape parameter, y	0.6	0.077			

Notes: We normalize the regulator preferences on the "utility" sector to 1.0, since regulator preferences are identified only up to a scale normalization. Standard errors are computed using 100 bootstrap iterations, sampling with replacement at the firm level.

TABLE 5
Model fit: moments

Moment	Simulated	Empirical
Mean violations		
Manufacturing	0.218	0.194
Resources	0.109	0.075
Services	0.140	0.134
Trade	0.118	0.090
Transportation	0.121	0.145
Utility	0.212	0.187
Viol. variance share		
Within-firm	0.195	0.214
Within-to-across-firm ratio	0.621	0.618
Responsiveness	−6.393	−6.512

6. RESULTS

The first-stage estimates are in Online Appendix Table A-7. The second and third-stage estimates 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 utility plants tend to have higher types, and trade and resources plants tend to have lower types.

6.1 *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 a binscatter representing the joint distribution of plant types with the leave-one-out average of other co-owned plant types. The results indicate that the underlying distribution of types is correlated within firms.

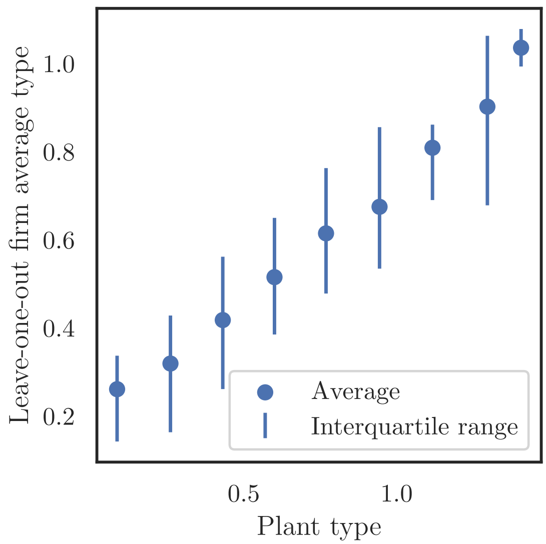

FIGURE 4
Estimated correlation of types

7. COUNTERFACTUALS

7.1 *Overview*

To evaluate the efficiency of linked regulation, we simulate a counterfactual budget increase that would increase the probability of inspecting a plant by an average of 10 percentage points. We measure effectiveness in terms of the decrease in total perceived social costs of violations: $\sum_j h_{g(j)} v_j$. Note that although we can measure changes to the absolute number of violations v_j, the social costs $h_{g(j)}$ are only identified up to a scale normalization. Therefore, we focus on percentage changes in total social costs, which are unaffected by this normalization. We also measure effectiveness using a pollution index; we describe the index construction in Online Appendix 2.2.

Overall, we investigate how much more effective this budget increase would be compared to an increase in untargeted 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 equally between linked and unlinked inspections. In Online Appendix 5.3 we show that these results are overall robust to assumptions on the regulator's problem by also presenting changes in violations disaggregated by industry sector.[20] We focus mainly on multi-plant firms because these are the targets of the regulation.

7.1.1. Theoretical properties of the optimal policies We provide an extended discussion of the theoretical properties of an increase in linked regulation, as well as the other

20. As previously mentioned, the regulator's problem is only used to compute estimates for the (perceived) social costs $h_{g(j)}$, which are then used in the counterfactuals to aggregate up violations across different sectors into a single number. All the other parameters are identified and estimated separately with no assumptions on the regulator's problem. Therefore, presenting the results by industry sector tests whether assumptions about the regulator are driving the results; ultimately, we find that this is not the case.

counterfactuals considered in this section, through the lens of a two-period stylized model in Online Appendix 6. We also use the model to describe theoretical properties of optimal regulation.

7.1.2. Computational details We model changes to untargeted inspections, unlinked inspections, and linked inspections, by modifying the parameters in the regulator's inspection function: $\rho^z_{0g(j)}$, ρ^z_1, ρ^z_2. Specifically, we interpret untargeted inspections as an increase in the constant of the inspection function $\rho^z_{0g(j)}$. Similarly, changing the level of unlinked inspections corresponds to changing the coefficient on plant scores (ρ^z_1), and changing the level of linked inspections corresponds to increasing the sensitivity to manager scores (ρ^z_2). A 50/50 mix of linked and unlinked escalation involves changing the level of ρ^z_1 and ρ^z_2 equally. We use a root-finding algorithm to search for the coefficient where the total probability of inspection equals the (increased) inspections budget. Note that to compute the counterfactuals we completely recompute the new equilibrium, which includes recomputing the equilibrium beliefs of agents.

7.2 *Decomposition: correlated targeting vs firm-wide moral hazard*

There are two main channels by which linked regulation may outperform untargeted inspections. We provide an intuitive discussion of these mechanisms here and leave a more complete theoretical justification to Online Appendix 6.[21] In addition, we provide more details about how we operationalize these decompositions numerically in Online Appendix 4.7.

The first channel is a "correlated targeting mechanism". The idea here is that, because types are correlated within a firm, the links between plants provide useful information about where next to investigate and allow the regulator to allocate inspections towards the most responsive plants.

We isolate the correlated targeting mechanism by first fixing each firms' mapping from scores to optimal actions ($a^*_j(\hat{\mathbf{s}}_j)$) at the baseline. Note that this also fixes the responsiveness of each plant to scores *e.g.* $\partial a^*_j/\partial s_j$. We then simulate a new stationary distribution of actions, violations, and scores, using the *counterfactual* regulator's inspection function. We record the decrease in the total social cost of violations. We also repeat the same procedure for the case where the budget is allocated to untargeted inspections. The difference between the two totals is what we call the correlated targeting mechanism. Intuitively, when types are correlated, linked regulation uncovers more violations at the most responsive plants. This then leads these plants to have relatively higher scores in the stationary equilibrium, lower actions, and fewer violations, compared to untargeted inspections.

The second channel is a "firm-wide moral hazard mechanism". Here, the links allow the regulator to punish a common owner for bad behaviour, so a violation at one plant increases deterrence at co-owned plants. Intuitively, owners internalize that a violation at one plant will escalate scrutiny at other commonly-owned plants, which causes them to change their optimal choice of actions conditional on scores. We isolate this mechanism by allowing the mapping from scores to actions $a^*_j(\hat{\mathbf{s}}_j)$ to change, which accounts for the remainder of the total effect of regulation.

We also show using a stylized model in Online Appendix 6 how both of these effects depend on the correlation structure in plant types. We show that, as the plant types become more correlated, the effect of both mechanisms becomes larger.

21. The stylized model in that section provides a theoretical justification for this decomposition and also illustrates that the total effect can be divided additively into these two effects.

TABLE 6
Counterfactuals

	↑ Inspections budget by 10%. Spent on:		
	Linked reg.	Unlinked reg.	50/50 mix
	%ΔPerceived social cost vs. untargeted:		
Multi-plant firms	−76.2%	−55.3%	−74.4%
Decomposition:			
= Correlated targeting mechanism	−8.7%	−9.6%	−9.7%
+ Firm-wide moral hazard mechanism	−67.5%	−45.7%	−64.7%
Single-plant firms	−34.0%	−35.0%	−36.4%
All firms	−48.0%	−41.7%	−49.0%
	%ΔSocial costs derived from pollution index vs. untargeted:		
Multi-plant firms	−76.7%	−55.3%	−74.6%

Notes: In the baseline model the average number of violations per plant (both discovered and not discovered by the inspector) is 0.74. The row "%Δ Social costs derived from pollution index vs. untargeted" uses the pollution index. Details about the pollution index construction are in Online Appendix 2.2. Robustness to alternative constructions of the pollution index are in the same section.

7.3 *Discussion*

We present the counterfactual results in Table 6. Recall that the objective of the regulator is to *minimize* the social cost of violations and so reductions in this metric correspond to "more effective" regulation. To produce the numbers in Table 6 we first consider an increase in the inspection budget devoted to untargeted inspections, and then compare how much more effective alternative types of regulation would be compared to this. In other words, the numbers in Table 6 are the change in the (perceived) social cost of violations compared to untargeted additional inspections.

7.3.1. Results: linked regulation We first consider a counterfactual policy where additional inspections are entirely devoted to linked regulation. We report the results in the first column of Table 6. Linked regulation performs dramatically better than an increase in untargeted inspections. Concretely, they reduce the social cost of violations by 76.2% overall compared to untargeted inspections.

The decomposition shows why: the correlated targeting effect of 8.7% indicates that linked regulation shifts inspections towards plants that are more responsive to regulation, while the firm-wide moral hazard effect of 67.5% shows that linked regulation is a powerful form of deterrence. Firms now face the threat of escalations not only on the offending plant, but on all co-owned plants, thereby amplifying this effect. Although the correlated targeting effect is smaller than the firm-wide moral hazard effect, there may be other settings in which the correlated targeting effect is comparatively larger. For instance, if firms can shift pollution from one plant to another to avoid detection (as in Gibson, 2019), correlated targeting would be an effective means of predicting the sites onto where the firm is shifting its pollution.

We also report the effect on single-plant firms. Here, the plant score and the firm score are identical, so we might expect the overall effects of unlinked and linked regulation to be similar. However, this intuition is less obvious when accounting for the equilibrium effects of extra linked inspections: one might be concerned that extra linked regulation may reveal more information about violations at multi-plant firms, and, therefore the large effects on multi-plant firms are due to a reallocation of existing enforcement resources from single-plant firms. The results for single-plant firms show that this is not the case: the decrease in the social cost of violations is qualitatively similar regardless of whether unlinked or linked regulation is used. Finally, the total effect is just a weighted average of the effects on single-plant and multi-plant firms.

7.3.2. Results: unlinked regulation We next consider the results for the counterfactual policy where additional inspections are entirely used for unlinked regulation (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 the perceived social cost of violations by 55.3% more than untargeted inspections in multi-plant firms.

The decomposition illustrates how unlinked regulation compares to linked inspections. The correlated targeting effect shows that unlinked regulation targets more responsive plants roughly as well as linked regulation. This is because types are fixed at individual plants across time. The firm-wide moral hazard effect of 45.7% shows that unlinked escalations have a deterrence effect. However, this effect is smaller than for linked regulation. Unlike linked regulation, unlinked regulation escalates at only one plant in the firm's portfolio, so the deterrence effect is more limited.

7.3.3. Results: 50/50 split of linked and unlinked escalations We also compute an equal mix of linked and unlinked escalations for the extra inspections budget in the third column. We find that the mix does slightly worse for multi-plant firms—but slightly better overall—compared to only linked escalations.

7.3.4. Results: using actual pollution data We also consider replacing the perceived social costs inferred from the regulator with an index constructed from actual pollution data linked to plants. This is important because the regulator may consider economic factors other than environment and health when deciding where to focus inspections (see, *e.g.* Deily and Gray, 1991 for an example). The details of how we incorporate these data are in Section 2.2 of the Online Appendix. The results are in the final row of Table 6 and are overall very similar to our main specification.

7.3.5. Discussion: heterogeneity and targeting Figure 5 breaks down the change in inspection probability by type quantile for a 10 percentage point increase linked regulation and in untargeted inspections. The linked regulation counterfactual allocates more of these inspections towards high type plants. These high type plants have, on average, higher baseline marginal deterrence: the correlation between baseline marginal deterrence and plant type is 0.54.

7.3.6. Discussion: role of commitment An assumption that we make throughout the paper is that the regulator can commit to a policy. While this is arguably reasonable in our empirical setting, one may ask how the results extend to a scenario where the regulator cannot fully commit. A complete characterization of the set of self-enforcing policies is outside the scope of this paper (and, indeed, is known to be challenging more generally).[22] However, our decomposition into the two mechanisms suggests that linked regulation can be implemented and add some value even in the absence of commitment. The firm-wide moral hazard mechanism certainly requires a credible commitment from the regulator to go through with their firm-level punishment to ensure deterrence. However, the correlated targeting mechanism merely directs the regulator to inspect the plants with the highest marginal abatement benefits, which may be self-enforcing

22. A useful summary of common difficulties is in Skreta (2006) in Section "Technical Difficulties and the Procedure".

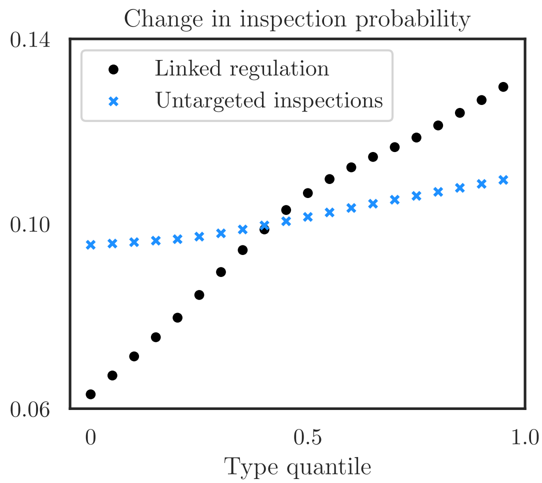

FIGURE 5
Targeting under linked regulation

Notes: Horizontal axis reports bins of plant type quantiles. The figure reports the LOWESS-smoothed change in inspection probability between baseline and coutnerfactual inspection policies, by type quantile

behaviour even in the absence of commitment. As an additional exercise in the Online Appendix 6.3, using a stylized version of our empirical setting, we show that targeting plants with higher responsiveness to regulation is optimal both with and without commitment.

7.4 *Robustness and additional results*

We run several tests to check how robust our conclusions are to alternative modelling assumptions, and also to provide more detail about what is driving the value-added of linked regulation in the model. In Online Appendix 5.1, we examine the extent to which correlation of compliance costs within plants of the same industry is driving the value added from linked regulation. Although the correlation does not play a major role in the overall results, it does comprise a notable share of the correlated targeting mechanism. This suggests that linked regulation could act as a way to "nudge" the regulator in a socially beneficial direction if the fixed parts of the regulatory scheme (here, the complexity points) are not optimally specified.

In Online Appendix 5.6, we test robustness to the assumption that the regulator does not internalize the compliance costs of the firms; the estimates are qualitatively similar to our baseline estimates.

We next examine counterfactual policies that use longer violation histories to inform scores. We do this by changing the persistence of the AR(1) process that governs scores, and re-simulate the model to equilibrium. Details of our implementation and results are in Online Appendix 5.4. We find that using slightly longer histories is beneficial, but if scores grow too persistent, they become insensitive to present-day violations and regulation becomes less effective. This resembles the intuition in Hörner and Lambert (2021).

Finally, we simulate a "random crackdown" as in Eeckhout *et al.* (2010). We provide more information about our empirical implementation in Online Appendix 5.2. We find that increasing inspections for only a random subset of plants is less efficient than our main policy counterfactuals, with perceived social harms 19% higher under a 10 percentage point increase in the budget devoted to crackdowns rather than untargeted inspections. The stylized model in the Online

Appendix provides intuition for why crackdowns are less useful in our setting. One reason is due to the curvature of equilibrium actions with respect to inspections. A second reason is that random crackdowns do not take advantage of histories and the resulting benefits of dynamic targeting through the firm-wide moral hazard and correlated targeting mechanisms. This may explain that, while random crackdowns are used for traffic infractions where histories are not directly observed (as in Eeckhout *et al*., 2010), they are not used—to our knowledge—by the TCEQ who regulate firms where more prior information is available.

8. 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 study ‘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 the effectiveness of regulation.

We develop a new empirical framework to study dynamic linked regulation. Our framework allows us to go beyond a simple quantification of whether linked regulation adds value in practice. In addition, we are able to test two alternative theoretical mechanisms about *why* it works: a firm-wide moral hazard mechanism and a correlated targeting mechanism. Disentangling these two mechanisms is important to understanding if linked regulation would work in other settings (for example, when the regulator cannot commit to a policy). Overall, we show that linked regulation can add significant value compared to other forms of regulation, and that the firm-wide moral hazard mechanism is primarily responsible for its benefit.

Acknowledgments 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 editor Bård Harstad and to two anonymous referees. Also thank you to Wesley Blundell, Hector Chade, Scott Holladay, Panle Jia Barwick, Karam Kang, Nick Kuminoff, Ashley Langer, Alvin Murphy, Devesh Raval, Mathias Reynaert, Conor Ryan, Jonathan Williams, and Gabriel Ziegler, as well as audiences at AERE, ASU, FTC, IIOC, NBER Decentralization Conference, NBER EEE, SEA, U Wisconsin Madison AAE, DSE Summer Conference for useful comments.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability

The data and code underlying this research are available on Zenodo at https://doi.org/10.5281/zenodo.16987699.

REFERENCES

ABRAMO, A. and LEHMAN, J. (2015), “How N.Y.’s Biggest For-Profit Nursing Home Group Flourishes Despite a Record Of Patient Harm” propublica.com website. https://www.propublica.org/article/new-york-for-profit-nursing-home-group-flourishes-despite-patient-harm.

BHATTACHARYYA, S. and LAFONTAINE, F. (1995), “Double-sided Moral Hazard and the Nature of Share Contracts”, *The RAND Journal of Economics*, **26**, 761–781.

BLOOM, N., BRYNJOLFSSON, E., FOSTER, L., *et al.* (2019), “What Drives Differences in Management Practices?”, *American Economic Review*, **109**, 1648–1683.

BLUNDELL, W., EVANS, M. F. and STAFFORD, S. L. (2021), “Regulating Hazardous Wastes under U.S. Environmental Federalism: The Role of State Resources”, *Journal of Environmental Economics and Management*, **108**, 102464.

BLUNDELL, W., GOWRISANKARAN, G. and LANGER, A. (2020), “Escalation of Scrutiny: The Gains from Dynamic Enforcement of Environmental Regulations”, *American Economic Review*, **110**, 2558–2585.

CHEN, Q., CHEN, Z., LIU, Z., *et al.* (2025), "Regulating Conglomerates: Evidence from an Energy Conservation Program in China", *American Economic Review*, **115**, 408–447.

COLMER, J., EVANS, M. F. and SHIMSHACK, J. (2025), "Environmental Citizen Complaints", *Journal of Political Economy: Microeconomics*. Accepted.

CRÉMER, J. and MCLEAN, R. P. (1988), "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions", *Econometrica: Journal of the Econometric Society*, **56**, 1247–1257.

DEILY, M. E. and GRAY, W. B. (1991), "Enforcement of Pollution Regulations in a Declining Industry", *Journal of Environmental Economics and Management*, **21**, 260–274.

EECKHOUT, J., PERSICO, N. and TODD, P. E. (2010), "A Theory of Optimal Random Crackdowns", *American Economic Review*, **100**, 1104–1135.

EPA. (2005), "Cosmed Agrees to Pay $1.5 Million for Clean Air Act Violations-Agreement Will Result in Improved Air Quality in Six States" epa.gov website. https://www.justice.gov/archive/opa/pr/2005/August/05_enrd_423.htm.

—— (2015), "Rcra's Critical Mission & the Path Forward" (Report).

GIARDILI, S., RAMDAS, K. and WILLIAMS, J. (2023), "Leadership and Productivity: A Study of us Automobile Assembly Plants", *Management Science*, **69**, 1500–1517.

GIBSON, M. (2019), "Regulation-induced Pollution Substitution", *The Review of Economics and Statistics*, **101**, 827–840.

GOLDFARB, A. and XIAO, M. (2011), "Who Thinks about the Competition? Managerial Ability and Strategic Entry in US Local Telephone Markets", *American Economic Review*, **101**, 3130–3161.

Government Accountability Office (2018), "SEC Inspections of Financial Industry Regulatory Authority's Governance Were Consistent with Internal Guidance", Accessed: 2025-04-04. https://www.gao.gov/assets/gao-18-522.pdf.

GOWRISANKARAN, G. and RYSMAN, M. (2012), "Dynamics of Consumer Demand for new Durable Goods", *Journal of Political Economy*, **120**, 1173–1219.

HÖRNER, J. and LAMBERT, N. S. (2021), "Motivational Ratings", *The Review of Economic Studies*, **88**, 1892–1935.

JIN, G. Z. and LESLIE, P. (2009), "Reputational Incentives for Restaurant Hygiene", *American Economic Journal: Microeconomics*, **1**, 237–267.

KANG, K. and SILVEIRA, B. S. (2021), "Understanding Disparities in Punishment: Regulator Preferences and Expertise", *Journal of Political Economy*, **129**, 2947–2992.

MCAFEE, R. P. and RENY, P. J. (1992), "Correlated Information and Mechanism Design", *Econometrica: Journal of the Econometric Society*, **60**, 395–421.

MILLIMET, D. L. (2014), "Environmental Federalism: A Survey of the Empirical Literature", *Case Western Reserve Law Review*, **64**, 1669.

MOOKHERJEE, D. and PNG, I. P. L. (1994), "Marginal Deterrence in Enforcement of law", *Journal of Political Economy*, **102**, 1039–1066.

OSHA (2024), "29 cfr §1903.3—Authority for Inspection". Accessed: 2025-04-01. https://www.osha.gov/laws-regs/regulations/standardnumber/1903/1903.3.

—— (n.d.), "Having Multiple Facilities Can Lead to Repeat Violations" ilt.safetynow.com website. https://ilt.safetynow.com/having-multiple-facilities-can-lead-to-repeat-violations.

PCAOB (n.d.), "Public Company Accounting Oversight Board Inspections Policy". Accessed: 2025-04-04. https://pcaobus.org/oversight/inspections.

POLINSKY, A. M. and SHAVELL, S. (1998), "On Offense History and the Theory of Deterrence", *International Review of Law and Economics*, **18**, 305–324.

RIJAL, B. and KHANNA, N. (2020), "High Priority Violations and Intra-Firm Pollution Substitution", *Journal of Environmental Economics and Management*, **103**, 102359.

SCUR, D., SADUN, R., VAN REENEN, J., *et al.* (2021), "The World Management Survey at 18: Lessons and the way Forward", *Oxford Review of Economic Policy*, **37**, 231–258.

SEGERSON, K. (2022), "Group Incentives for Environmental Protection and Natural Resource Management", *Annual Review of Resource Economics*, **14**, 597–619.

SHIMSHACK, J. P. and WARD, M. B. (2022), "Costly Sanctions and the Treatment of Frequent Violators in Regulatory Settings", *Journal of Environmental Economics and Management*, **116**, 102745.

SKRETA, V. (2006), "Sequentially Optimal Mechanisms", *The Review of Economic Studies*, **73**, 1085–1111.

TCEQ (2022), "Annual Enforcement Report" (Report).

Texas Administrative Code (2002), "Environmental Quality", 30.1(60.2).

Texas Environmental Public Participation Guide (2017), "The University of Texas Law School Environmental Clinic" https://law.utexas.edu/wp-content/uploads/sites/11/2019/01/2017-EC-EnviroPublicParticipationGuide.pdf.

VARAS, F., MARINOVIC, I. and SKRZYPACZ, A. (2020), "Random Inspections and Periodic Reviews: Optimal Dynamic Monitoring", *The Review of Economic Studies*, **87**, 2893–2937.