Online Appendix

Contents

1 Additional context: plant scores

In this section we detail the exact algorithm used to compute the plant score. The score of plant j is given as follows:

$$s_j = \frac{\Big[(\text{Violation Points}) + (\text{Chronic Excessive Emission Points}) + (\text{Repeat Violator Points}) - (\text{Self Audit Points})\Big]}{(\text{No. of investigations} \times 0.1) + (\text{Complexity Points})} \tag{A-1}$$

The entire score s_j is also then multiplied by "Voluntary Program Points (if applicable)", but as we explain below these are extremely rare for plants in our dataset. The components in the above equation are as follows. The term "violation points" corresponds to the number of violations at a particular plant (these may be weighted by the magnitude of the violation). The terms "chronic excessive emission points" and "repeat violator points" correspond to adjustments if the plant has a history of the same violation. The term "self audit points" corresponds to minor score adjustments if plants have a formal environmental audit system. In the denominator, "no. of investigations" corresponds to the number of recent inspections.

1.1 Voluntary actions

Overall, we abstract away from self-audits and participation in the voluntary programs to ensure that the empirical model is tractable. We justify these modeling decisions in more detail below; overall, the justification is that these actions are extremely rare for the firms in our dataset.

We first obtained data on which entities actually "participated in voluntary environmental management systems" or "conducted a self-audit", through a Freedom of Information Request to the TCEQ. We were able to obtain data on which plants "conducted a self-audit" for the entire sample period 2012-2020. For plants that "participated in voluntary environmental management systems", we were able to obtain data on this element for the 2017-2020 period, which is not the full sample period but arguably enough to serve as a robustness check.

We find that the firms in our dataset rarely take these voluntary actions. Only 1.1% of plants in our dataset conducted a self audit at any point in 2012-2020. In terms of participation in voluntary environmental systems, the number is even lower at only 0.4%. Furthermore, we investigate the marginal effects of these voluntary actions on scores by doing an exercise where we adjust the raw scores of plants to undo the effects of these actions.[1] The R^2 of a regression on the raw scores on these adjusted scores is approximately 1.0, indicating these voluntary actions generate effectively none of the variation in observed scores.

In summary, due to the low frequency of these voluntary actions and the lack of overall effect on the scores, we do not model them explicitly.[2]

2 Data construction details

2.1 Matching scores, violations, inspections, and penalties

We discuss here the various datasets in the analysis and how they are merged together. There are two main data sources: the TCEQ and the federal EPA data on Texas facilities. Note that much of the data in the federal EPA database originates from the TCEQ, but is only readily available from the EPA. From the federal EPA's NPDES and RCRA databases, we combine data on violations, inspections, and penalties. We also use additional data from the federal EPA on pollution, which we describe further in Online Appendix Section 2.2. For the robustness tests on production reallocation, we also connect data on violations from co-owned plants in other states; we discuss this

[1] Specifically, if there was a voluntary program then we divide the total score by the maximum adjustment for "Voluntary Program Points" which is 0.75 (Texas Administrative Code [2002]. Note that using the maximum adjustment is a conservative assumption for this test because it would tend to change the scores more, thereby potentially reducing the R^2 in the regression. We then adjust for the total number of "Self-audit points" divided by complexity points.

[2] Moreover, since the counterfactuals involve a relatively small increase in the budget, and these voluntary actions are rare given the baseline budget, it is arguably unlikely that these voluntary actions would become substantially more common in counterfactuals.

in more detail in Online Appendix Section 3.1. From the TCEQ, we obtain annual score data from the website (for more recent files) and the Wayback Machine (for data on previous years back to 2012).

We match the annual score data by first connecting Texas plant IDs to their federal facility number (FRS ID) and then to the NPDES and RCRA databases using the EPA's crosswalk file. We aggregate the data to the yearly level. We take care to split out violations detected from "Discharge Monitoring Reports" in the NPDES because theseare typically discovered by means other than inspections. These violations are self-reported and are detected through firms' required self-monitoring equipment in bodies of water. For the descriptive analysis in Section 3 we focus on violations that are typically detected via inspections. However, data on violations detected by means other than inspections are still incorporated into the TCEQ scores and in our corresponding empirical model and overall results. We explain further how we handle these cases in Online Appendix Section 4.4 and when explaining the model in Section 4.1.

Finally, we match the above annual-level database (which is constructed using the EPA data) to the relevant compliance history scores in the TCEQ database. Scores are updated every year in November (not on January 1st), and so we take care to match the relevant TCEQ score the regulator is using to the EPA data.

2.2 Pollution index

Constructing the pollution index. We construct this index using two main sources. First, we use the percentage of dissolved oxygen in the "hydrological unit" (e.g. river or stream) that the facility expends pollution into, as identified from its Clean Water Act permit. Although there are many different chemicals and pollutants that facilities can emit, the percentage of dissolved oxygen is a commonly used "omnibus" measure of water quality. This measure was used in Kang and Silveira [2021] as well as Keiser and Shapiro [2019]. Higher percentages of dissolved oxygen typically correspond to higher quality water, since fish and other life require oxygen to survive. We follow Keiser and Shapiro [2019] (in their Appendix B.3.) to clean the data and compute the measure of dissolved oxygen. We match the pollution data to plants using the 8-unit hydrological code and use the annual average dissolved oxygen measurement.

The second source that we use is the EPA's "Risk-Screening Environmental Indicators" (RSEI) model. This takes the plant-level Toxic Release Inventory (TRI) data as an input. The TRI contains over 400 different chemicals, and the RSEI takes all these different chemicals, as well as varying toxicity of the releases, and aggregates them to a one-dimensional index of risk. The RSEI focuses on the risks to human health. This index is most useful in evaluating releases subject to RCRA (the law corresponding to hazardous wastes).

We then aggregate these two broad measures into a single pollution index, as follows. We normalize both indexes by their mean and standard deviation, so that they have a common scale and a mean of 0.0.[3] We then aggregate the measures with an equal weighting to our pollution index.[4] If a plant does not have a recorded value under one of the measures, then we construct the pollution index solely using the available index (typically, this will correspond to a plant that is regulated under one act but not the other). If the plant does not have a recorded value under either measure then we do not use their value in the analysis.

Using the pollution index in the empirical model. We then map the pollution data to violations in the model, allowing for heterogeneous social costs by industry. To do so, we look at the distribution of pollution across plants within an industry.[5] We compute the social costs of a marginal violation for each industry as moving from the median quantile of pollution to the 90th quantile of pollution among plants within that industry.[6] Finally, since only *relative* pollution matters for our percentage change results, we normalize the pollution per violation so that the utility sector has a value of 1.0 (as we similarly do for the regulator's "perceived" cost of violations).

We report the counterfactual results using the pollution index in Table A-1. For robustness, we construct the pollution index with different weights. We also construct the pollution index using

[3]Duflo et al. [2018] similarly aggregate measures of different types of pollution by standardizing each measure using the pollutant's standard deviation.

[4]Since weighting the indexes equally in the final pollution measure involves an assumption, we test robustness to alternative weights when we report the results in the next section.

[5]Our procedure shares a motivation with the procedure in Blundell et al. [2020] who write "We focus on the idea that the distribution of pollution across plants in an industry forms the basis of expectations about pollution quantities..."

[6]This is consistent with distribution of the number of violations per plant in the data: the 50th percentile of violations is 0 and the 90th percentile is a single violation. This procedure takes into account that there is an "average" amount of pollution that will occur even in the absence of a violation; for example, the RSEI score can be positive even in the absence of a regulatory violation.

Table A-1: Counterfactuals computed using pollution data

	↑ Inspections budget by 10%. Spent on:		
	Linked reg.	Unlinked reg.	50/50 mix
%Δ Social cost vs untargeted:			
Baseline (regulator preferences)	76.2%	55.3%	74.4%
Pollution index:			
Weight = 0.1	78.5%	55.7%	76.1%
Weight = 0.5	76.7%	55.3%	74.6%
Weight = 0.9	74.8%	54.6%	73.0%
Weight = 0.5, lower quantile = 0.1	78.9%	55.5%	76.2%
Weight = 0.5, upper quantile = 0.95	75.8%	55.2%	73.9%
Using environmental justice index	73.8%	54.3%	72.3%

Notes: This table illustrates the overall counterfactual results for different methods to quantify the social cost of violations. The first method uses the "perceived" social costs estimated from relative differences in the regulator's probability of inspection each industry sector. The methods using the pollution index measure the social cost of violations using actual pollution data. Here we illustrate robustness to different assumptions on how to weight the inputs into the pollution index. We also illustrate robustness to alternative assumptions about the lower and upper quantiles for the marginal violation used in the data construction.

alternative lower and upper quantiles of pollution. Finally, we also construct an alternative index using the relative mean EPA environmental justice index (which we also use in the main paper) at the industry level. Regardless of how we construct the index, or if we instead use the regulator's "perceived" social costs, the results are qualitatively robust. This suggests that the regulator's preferences are fairly well-aligned with the "true" social costs of violations.

3 Descriptive evidence: robustness

3.1 Robustness: production reallocation

We test whether firms respond to increased scrutiny at one plant by reallocating production to other plants. Overall, we do not find evidence that firms reallocate production and so do not include this channel as a margin for firms' responses to the regulation in the model.

Connecting Texas plants to commonly owned plants in other states While our TCEQ dataset covers firms in Texas, so measuring reallocation across plants location within Texas is straightforward, reallocation could potentially also be to plants outside Texas. Therefore, we first match plants in our dataset to co-owned plants in other states.

We use the "DUNS Number" from Dun and Bradstreet which is available in the EPA's Toxic Release Inventory (TRI) data. Most of the plants in our TCEQ dataset need to report to the TRI, and so this measure has good coverage for the plants in our data. One exception is that the definition of a "plant" is slightly more granular in the TCEQ dataset compared to the definition from Dun and Bradstreet. For example, if there are two contractors who regularly operate at the same gas station with different responsibilities, the TCEQ distinguishes their behavior, but Dun and Bradstreet would aggregate them since the DUNS number only varies by location. Therefore, we follow a conservative procedure of only linking plants from outside Texas that are connected by a unique DUNS number. In the end, this procedure allows us to link 7018 plant-year observations to firms in our Texas dataset. (Note that if a plant is unmatched then we therefore implicitly set the violations outside Texas to 0.)

Testing reallocation through plant exit rates Our production reallocation tests center on explaining the determinants of plant exits. We define an exit as any instance in which a plant either changed owner or no longer appears in the score dataset. Since 2020 is the final year of our data, we exclude the 2020 data from our reallocation regressions, since we cannot determine whether 2020 was the final year of a plant. Ideally we would also test for an intensive margin response, but we do not have production quantity data due to our focus on small and medium sized businesses (i.e. the firms are not large publicly traded firms where such data would be readily available).

We perform linear probability regressions in Table A-2 where the dependent variable is an indicator for whether the plant exited and the independent variables are firm and plant scores. We restrict the sample only to multi-plant firms, since these firms could potentially reallocate production across commonly-owned plants. To fix ideas, evidence consistent with production reallocation would involve a positive relationship between higher plant scores—which imply a worse environmental record—and the probability of exit, after controlling for the firm-wide score. However, the results across specifications (1), (2), and (3) in Table A-2 show exactly the opposite: in fact, a higher plant score slightly reduces the probability of exit. Explaining this negative relationship is outside of the scope of this paper. However, note that it is unlikely to dramatically affect plant decisions about whether to violate: scores change fast (returning to 0 after 5 years with no violations) but plants

Table A-2: Robustness: exit/reallocation regressions

	(1) 1[Exit]	(2) 1[Exit]	(3) 1[Exit]
Log(1+firm score)	0.005	0.003	0.003
	(0.002)	(0.002)	(0.002)
Log(1+plant score)	-0.01	-0.009	-0.009
	(0.002)	(0.002)	(0.002)
Year FEs	Yes	Yes	Yes
NAICS Category FEs	No	Yes	Yes
Region FEs	No	No	Yes
Only multi-plant firms	Yes	Yes	Yes
Control for # violations outside Texas	Yes	Yes	Yes
N	18622	18622	18622
Av. prob of exit per year	0.032	0.032	0.032

are long-lived (with an implied average lifespan of 25 years from the average probability of exit) and so short-run delays to exit are unlikely to be a first-order factor in decisions.

Testing reallocation through changes in violations in co-owned plants outside Texas We check for reallocation using a Poisson regression for the number of violations in plants located outside Texas on firm-level scores. In this test, evidence of production reallocation would be a positive and significant relationship between the Texas scores and violations from plants located outside Texas.

Across the three specifications in Table A-3, we find that the relationship is not significant and sometimes has a negative sign. Overall, we again do not find evidence here of production reallocation in our context, even with this improved test.

3.2 Robustness: strategic avoidance

We test whether firms respond to increased scrutiny at one plant by strategically avoiding the regulation at other plants. We are specifically interested in testing 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. Overall, our tests do not find evidence for this hypothesis.

Our strategic avoidance tests exploit that TCEQ inspectors are drawn from the local TCEQ regional

Table A-3: Testing production reallocation for plants outside Texas

	(1) Violations	(2) Violations	(3) Violations
Log(1+firm score)	-0.03 (0.0537)	-0.042 (0.0534)	0.014 (0.0582)
Firm-level FEs	Yes	Yes	Yes
Year FEs	No	Yes	Yes
Plant-level FEs	No	No	Yes
N	6698	6698	4527

office. Therefore, a violation that is discovered at a co-owned plant in the same region might provide more information for strategic avoidance than co-owned plants in other regions that are evaluated by different inspectors. Our results are presented in Table A-4. Here we restrict attention to plants that are in firms where there is at least one other plant in the same region, and at least one other plant in a different region.

In Part (1) of Table A-4 we provide a baseline correlation of within-firm violations (denoting by $\hat{v}$ the transformation $log(1 + \text{violations per inspection})$). Consistent with our findings for the full sample in the main paper, we find that violations are correlated within a firm in this subsample. Then, in Part (2) of Table A-4, we test whether this correlation depends on whether the within-firm violations are discovered by plants within the same TCEQ region or by plants in other regions. Evidence consistent with strategic avoidance would be that violations are more strongly correlated in plants that are situated in different regions than in the same region. However, we find that this is not the case: violations in one plant j are correlated with both plants situated in the same and also in different regions. Furthermore, we find that this correlation is approximately of the same magnitude. While it is difficult to provide a comprehensive test for strategic avoidance, these results do not suggest that it is a primary concern.

3.3 Robustness: role of federal EPA regulations

The vast majority (96.5 percent) of inspections in our dataset are conducted by the state regulator (the TCEQ). This is consistent with the federal EPA's description of the CWA and RCRA as

Table A-4: Robustness: strategic avoidance

(1)	
	$\hat{v}$ at j
$\hat{v}$ at $k \in \mathcal{J}_f/j$	0.17
	(0.04)
NAICS Category FEs	Yes
Region FEs	Yes
Only multi-region firms	Yes
N	795

(2)	
	$\hat{v}$ at j
$\hat{v}$ at $k \in \mathcal{J}_f/j$, same region	0.18
	(0.06)
$\hat{v}$ at $k \in \mathcal{J}_f/j$, other regions	0.13
	(0.04)
NAICS Category FEs	Yes
Region FEs	Yes
Only multi-region firms	Yes
N	795

primarily being enforced by states.[7] However, a small percentage of inspections (3.5 percent) are conducted by the federal EPA.

A decision that we needed to make is whether to retain this small number of federal inspections in the dataset or remove them. As we explain more below, and based on an exploration of the institutional details, we choose to retain them. Overall, accounting for the federal EPA's inspections has extremely minor effects on the quantitative findings, consistent with them accounting for a small fraction of total inspections.[8]

In Table A-5 we perform similar inspection probability regressions to Table 2 in the paper using only federal EPA inspections. Overall, we find that throughout specifications (1)-(3) federal inspections do not seem to respond significantly to the risk scores that the TCEQ uses. This is true both in a statistical sense and in an economic sense; the marginal predicted effects of an increase in scores on federal inspections is an order of magnitude smaller than for state inspections, as we report in Table A-6. As a result, we think that the best way to characterize the federal EPA's inspection decision is — from the point of view of each firm in Texas — an untargeted inspection that is not a function of the risk scores.

[7] For example, the EPA's website when describing RCRA states that: "EPA delegates the primary responsibility of implementing the RCRA hazardous waste program to individual states in lieu of EPA".

[8] Moreover, this suggests that if we made an alternative modeling decision and eliminated them entirely from the dataset, it would not make much difference to the results. However, based on the institutional details set out below, we think that the better modeling decision is to more carefully account for federal inspections.

Table A-5: Inspection probability regressions: federal EPA inspections only

Dependent Variable	(1) Inspection	(2) Inspection	(3) Inspection
Log(1+firm score)	0.073	0.072	0.065
	(0.073)	(0.073)	(0.074)
Log(1+plant score)	0.025	0.027	0.013
	(0.071)	(0.071)	(0.072)
Env. justice index	-	0.251	-
	(-)	(0.247)	(-)
Year FEs	Yes	Yes	Yes
NAICS category FEs	Yes	Yes	Yes
Region FEs	No	No	Yes
N	50864	50864	50864
Log-likelihood	-3015	-3014	-2972

Table A-6: Marginal effect of scores on inspection probabilities

Dependent Variable	State Inspection	Federal Inspection
Marginal effect of		
Log(1 + firm score)	0.012	0.0006
Log(1 + plant score)	0.023	0.0002

Notes: Marginal effects based on specification (1) of Table A-5 and specification (1) of Table 2. All marginal effects calculated at mean values of regressors.

Implications for the analysis in the paper Aside from data on violations discovered by the TCEQ's actions, the only other information that the TCEQ uses to compute risk scores are violations in the EPA's federal database (which we use), and this was formalized in 2011.[9] As a result, the small number of violations discovered by federal EPA inspections should be included in the updating rule for scores, and we do so in the model and descriptive results.

In terms of how federal EPA regulation affects the model, first recall that there are two types of agents in the model: a set of firms and the regulator. On the firm side, their problem can be characterized as responding to a "regulatory machine" (encoded by the empirical conditional choice probabilities (CCP) of inspection). Therefore, since violations enter into future scores regardless of whether they were discovered by the TCEQ or the federal EPA, firms can be interpreted as responding to a composite CCP that includes both state inspections as well as the 3.5 percent of federal inspections. Since, as mentioned, the federal EPA's inspection policy appears to not depend on the risk scores, federal EPA inspections will enter into the composite CCP as a slight increase in the intercept of the logit model for an inspection (i.e. as an untargeted inspection).

On the regulator's side, using a composite CCP for estimation will only be an issue if we need to consider alternative allocations where the total probability of an untargeted inspection falls below the federal EPA's probability of inspection (since they we would be cutting into the federal EPA's budget allocation). However, since estimation involves a marginal analysis (i.e. that the marginal benefit of an inspection is equalized across industries) that involves us simulating small deviations from the observed allocation, this issue does not occur.

A similar argument applies for why the counterfactuals are robust to using the composite CCP. That is, since our counterfactuals involve a budget *increase*, we are not considering counterfactuals where the number of inspections per industry is decreased so much that implicitly the allocation of federal inspections is changed.

[9] According to https://www.tceq.texas.gov/assets/public/legal/rules/hist_rules/Complete.11s/11032060/11032060_ado_clean.pdf in terms of the Clean Water Act: "Section 4.04 of HB 2694 amends TWC, §5.753(b) to remove enforcement actions from other states and the federal government, except actions by the United States Environmental Protection Agency (EPA), as mandatory components of compliance history and to clarify that enforcement actions from the EPA are mandatory components to the extent readily available to the commission."

4 Model and estimation details

4.1 Firm's problem

We rewrite our starting point, which is Equation 5:

$$\max_{a_j} \ \pi_j(a_j; s_j, s_f) + \beta \mathbb{E}_{\mathbf{z},\mathbf{v}} \left[V_j(\mathbf{s}') + \sum_{k \in \mathcal{J}/j} V_k(\mathbf{s}') \middle| \mathbf{s}, \mathbf{a}^*_{-j}, a_j \right] \tag{A-2}$$

Then, applying the continuation value sufficiency assumption:

$$\max_{a_j} \ \pi_j(a_j; s_j, s_f) + \beta \mathbb{E}_{\mathbf{z},\mathbf{v}} \left[V_j(\mathbf{s}') + W_j' \middle| \mathbf{s}, \mathbf{a}^*_{-j}, a_j \right] \tag{A-3}$$

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

$$V_j(\mathbf{s}) = \pi_j(a_j^*(\mathbf{s}); s_j, s_f) + \beta \mathbb{E}_{\mathbf{z},\mathbf{v}} \left[V_j(\mathbf{s}') \middle| \mathbf{s}, \mathbf{a}^*_{-j}, a_j \right] \tag{A-4}$$

Note that the state transitions, as well as Equation A-3 and Equation A-4, are only a function of v_j, W_j, s_j, s_f. Therefore, the remaining states are not payoff-relevant and so we can replace $\mathbf{s}$ with $\hat{\mathbf{s}}_j = (s_j, s_f, W_j)$. Furthermore, the expectations can be written as simply over $\mathbf{z}_j$ and $\mathbf{v}_j$, which completes the derivation.

4.2 Regulator's problem

Simplifying notation and rearranging the regulator's problem (defined in Equation 9), the Lagrangian can be written as

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

When we assume that $\mathcal{Z}$ is the family of logit functions of the form in Equation 11 and that the intercepts for each industry sector are chosen optimally, the first order conditions with respect to these intercepts are

$$0 = \sum_j \left[h_{g(j)} \frac{\partial \mathbb{E}_{a,s,\theta}[a_j|\bar{z}]}{\partial \rho^z_{0i}} - \lambda \frac{\partial \mathbb{E}_{a,s,\theta}[z_j|\bar{z}]}{\partial \rho^z_{0i}} \right]$$

for each industry sector i (so, there are 6 in total). Stacking these 6 first order conditions yields the equation $\mathbf{0} = \mathbf{Ah} - \lambda\mathbf{z}$. Here, $\mathbf{A}$ is a 6×6 matrix with entry $A_{ik} = \sum_{j:g(j)=k} \frac{\partial \mathbb{E}_{a,s,\theta}[a_j|\bar{z}]}{\partial \rho^z_{0i}}$, $\mathbf{h}$ is a 6×1 vector of social costs h_g, and $\mathbf{z}$ is a 6×1 vector with ith element $\bar{z}_i = \sum_j \frac{\mathbb{E}_{a,s,\theta}[z_j|\bar{z}]}{\partial \rho^z_{0i}}$. Provided $\mathbf{A}$ is nonsingular, inverting this equation yields $\mathbf{h} = \lambda \mathbf{A}^{-1}\mathbf{z}$. The scalar Lagrange multiplier λ is unknown and, in general, is not identified without more information, which implies that the entire vector $\mathbf{h}$ is not identified. However, the direction of vector $\mathbf{h}$ is pinned down by this system of equations, implying that each h_g is identified up to the normalization that $h_{utility} = 1$.

4.3 Computing complexity scores

Complexity scores are not directly observed in our dataset. However, we can still estimate them in two steps and the relative weights are in Table A-7.

First, we use the fact that the firm score is a weighted average of the plant scores (defined in the main paper in Equation 2), where the weights are the complexity scores. Therefore, the weights are identified up to a scale normalization. As a result, we estimate the weights via non-linear least squares using the regulator's aggregation rule, normalizing the score on the utility sector to 1.0.

Second, we turn to the scale factor. Note that the way the scale factor of the complexity points enters into the estimated model is through its effect on how a violation is mapped into the own-plant score updating rule $r_v^{s_j}/Q_j$ = $(r_v^{s_j}/[\text{scale factor}]) \times (1/[\text{normalized complexity score}])$. Since $(1/[\text{normalized complexity score}])$ is known from the first step, we use the Texas Administrative Code to jointly calibrate the components of the parameter $r_v^{s_j}/[\text{scale factor}]$ using the following exercise.

We consider a typical plant in the utility sector: a small municipal stormwater plant regulated until the Clean Water Act which discharges into a nearby river. Consider the specific details of this plant: it is a "municipal minor" (+ 3 complexity points), with an EPA designated "facility identification number" (+ 1 points), with an external outfall (+1 points), which is in total 5 points. We further validate this scaling number by considering some example "compliance history reports" from the TCEQ which are available on the website.[10] We use the Texas Administrative Code Texas Administrative Code [2002] to obtain a reasonable value of $r_v^{s_j} = 45$; this is the increase in score from a "moderate violation", in the case where there is agreed final enforcement orders containing a denial of liability. Overall, this leads to our approximation of $r_v^{s_j}/[\text{scale factor}] \approx 9.0$, as reported in Table A-7.

4.4 Estimating the probability a violation is discovered by means other than inspections: γ

In this section, we describe how we estimate γ, the probability that a violation is recorded even if no inspection is conducted. To do this, we make two assumptions. Our first assumption is that γ is constant across all firms and years. As in our model, v is the number of violations the firm commits (the "true" violations), and $\hat{v}$ is the number of violations that are recorded. Let $[\text{inspection}] \in \{0, 1\}$ record whether an inspection occurred, i.e., the realization of a Bernoulli random variable with probability $\bar{z}$, where $\bar{z} = \bar{z}_{g(j)}(s_j, s_f)$ is determined by the inspector's policy function. Our first assumption implies

$$\text{Pr}(\hat{v} > 0|\text{inspection} = 0, \bar{z}) = \gamma\text{Pr}(v > 0|\text{inspection} = 0, \bar{z}). \tag{A-5}$$

If an inspection does occur, violations are definitely recorded, so:

$$\text{Pr}(\hat{v} > 0|\text{inspection} = 1, \bar{z}) = \text{Pr}(v > 0|\text{inspection} = 1, \bar{z}).$$

Rewrite Equation A-5 as $\gamma = \frac{\text{Pr}(\hat{v}>0|\text{inspection}=0,\bar{z})}{\text{Pr}(v>0|\text{inspection}=0,\bar{z})}$. Our second assumption is that, conditional on $\bar{z}$, the true violations for any plant are independent of whether an inspection actually occurs. This assumption is reasonable because plants choose actions based on probabilities of inspection $\bar{z}$ but before inspections actually occur. This means we can write

$$\gamma = \frac{\text{Pr}(\hat{v} > 0|\text{inspection} = 0, \bar{z})}{\text{Pr}(\hat{v} > 0|\text{inspection} = 1, \bar{z})}.$$

[10] For a similar type of plant in the utility sector (in this example, a sewage treatment facility), a score of 5 is reasonable: https://www.tceq.texas.gov/assets/public/comm_exec/agendas/comm/backup/HR-RFR/2015-1792-MWD-Misc.pdf.

The odds ratio on the right hand side can be estimated from observables. We only need to estimate the probability a violation is observed conditional on (a) inspection probabilities and (b) whether an inspection actually occurred. We implement this by estimating a logit model of the form: $Pr(\text{violation} > 0|\text{inspection}, \bar{z}) = Logit\{\beta_1[\text{inspection} = 1] + \beta_2\hat{z} + \lambda_{plant}\}$ where $\hat{z}$ is the estimated inspection probability from specification 1 in Table 2 and λ_{plant} is a plant fixed effect. We include a plant fixed effect because plant level unobservables that correlate with scores and inspection outcomes could affect true violations, v. We are interested in the odds ratio from turning off and on $[\text{inspection} = 1]$. In our logit, this equals $e^{\hat{\beta}_1}$, so our estimate of γ is $\hat{\gamma} = e^{\hat{\beta}_1}$. Implementing the procedure results in $\hat{\gamma} = 0.110$.

4.5 Detailed algorithm: firm's problem

Overview In this section we describe the estimation routine for recovering the dynamic parameters that govern firms' decisions. Overall, we use the simulated method of moments which requires us to compute model-predicted moments for different values of the parameters σ_F^2, σ_J^2, $\bar{\theta}_g$, y. We use a diagonal matrix to weight our moments, setting a weight = 1 for all the moments except for the "responsiveness" moment which we weight by 0.1 so that the moments have a common scale.

In order to compute the simulated moments — for a given set of the parameters — we compute a separate problem for each firm. So, for each objective function evaluation we need to compute thousands of individual portfolio problems corresponding to each firm f we observe in the data. When implementing the algorithm we exploit that the problem can be parallelized across each firm given the regulators' conditional choice probabilities. Hence, for a given set of the parameters, we can compute the problem for each firm simultaneously using a supercomputer cluster, which lowers the overall computational time.

Computational algorithm for each firm f We now describe the algorithm that we use to compute the problem for an individual firm f in our data with portfolio $\mathcal{J}_f$ plants. The algorithm consists of an outer loop, where we solve for state transition beliefs over the reduced state space $\hat{\mathbf{s}}_j$ for each plant that are consistent with the firm's overall actions over its entire portfolio of plants, and an inner loop where we solve for the optimal action at the plant level given the state transition beliefs of each plant j.

1. *(Outer loop):* Guess the parameters governing the state transition beliefs for each plant

$j \in \mathcal{J}_f$ i.e. the parameters underlying the matrices $R_{0,j}$ and $R_{1,j}$.[11]

2. *(Inner loop):* For each plant $j \in \mathcal{J}_f$ run the following algorithm until convergence:

   (a) Denote the inner loop iteration number by n, starting from $n = 0$. We iterate over two objects: the (continuous) optimal action at each state $a_j^{*,n}(\hat{\mathbf{s}}_j)$ and the value function at each state $V_j^n(\hat{\mathbf{s}}_j)$ (where $\hat{\mathbf{s}}_j = (s_j, s_f, W_j)$). At each iteration n, we store each of these objects as a set of nodes and use linear interpolation to compute intermediate values between the nodes. To initialize the inner loop we provide a guess of the optimal action at each node, and a guess of the value function at each node.

   (b) Update the optimal actions at each node using the first order condition derived from the optimal action choice Equation 6. This results in $a_j^{*,n+1}(\hat{\mathbf{s}}_j)$.

   (c) Update the value function at each node using Equation 7. This results in $V_j^{n+1}(\hat{\mathbf{s}}_j)$.

   (d) Check for convergence in the optimal actions and the value functions. If the algorithm has not converged, return to Step 2(a).

3. *(Outer loop, simulation step):* The inner loop results in policy functions that generate the optimal actions for each plant j at each state, given the current parameters of the state transition beliefs for each plant: $\{a_j^*(\hat{\mathbf{s}}_j)\}_{j \in \mathcal{J}_f}$. The inner loop also results in value functions for each plant defined over the reduced state space: $\{V_j(\hat{\mathbf{s}}_j)\}_{j \in \mathcal{J}_f}$. We next simulate the stochastic process of inspections and detected violations, given these policy functions and value functions, as follows (note that we also "burn in" the simulation to avoid dependence on the initial state and we are interested in the properties of the stationary distribution). We can think of the simulation step as generating a simulated "dataset" of plant scores, firm scores, continuation values at the other plants, and violations, for each period t in the simulation and for each plant j: $\{s_{j,t}, s_{f,t}, W_{j,t}, v_{j,t}\}_{t \geq t_{burn-in}, j \in \mathcal{J}_f}$.

   (a) Given the current state at each plant $\hat{\mathbf{s}}_{j,t}$ compute the optimal actions using the policy functions from the inner loop. From this action, draw violations at each plant from a Poisson distribution with the mean as the negligent action.

[11] These consist of a guess for the parameters for the state transitions corresponding to the firm score s_f and the state W_j - the plant score updating rule is directly observed in the data.

(b) Given the scores at each plant (and the observed industry of each plant $g(j)$) take a Bernoulli draw from the regulator's conditional choice probability of inspection at each plant to determine whether a plant will be inspected. If the plant is inspected and violations are discovered, record the number of violations in the simulated "dataset" $\hat{v}_{j,t}$. If the plant is not inspected, violations are discovered with probability γ and not discovered ($\hat{v}_{j,t} = 0$) with probability $1 - \gamma$.

(c) Update the plant scores at each plant using the plant state evolution and the detected violations in the previous step. Record the current value of each plant j score in the simulated "dataset" $s_{j,t}$.

(d) Update the firm score using the updated individual plant scores and the complexity score of each plant using Equation 1. Record the current firm score in the simulated "dataset" $s_{f,t}$.

(e) For each plant $j \in \mathcal{J}_f$:

   i. Record the "true" state of the continuation values of the other plants in the simulated dataset $W_{j,t} = \sum_{k \in \mathcal{J}/j} V_k(\hat{\mathbf{s}}_{k,t})$ using the value functions computed in the inner loop in Step 2.

   ii. Update the continuation value state in $\hat{\mathbf{s}}_{j,t}$ using the current guess of the transition coefficients for the continuation value score.

   Comment: One way to intuitively think about Step 3(e) is that there are two versions of $W_{j,t}$ in the simulation. The first is the version in $\hat{\mathbf{s}}_{j,t}$ which we get from Step 3(e)(ii) (which we could denote with a hat: $\widehat{W_{j,t}}$). This could potentially differ from the "true" $W_{j,t}$ we use to estimate the transitions in the outer loop which are computed from $W_{j,t} = \sum_{k \in \mathcal{J}/j} V_k(\hat{\mathbf{s}}_{k,t})$ in Step 3(e)(i). Note, however, that the objective of the algorithm is to minimize this discrepancy: in the outer loop we continue to iterate until the series that correspond to the two versions of $W_{j,t}$ converge.[12]

[12]Therefore, these transitions are consistent with – and disciplined by – the "true" continuation values. As a result, and amongst other things, they capture the first-order cross-plant effect that a detected violation at plant j will reduce the continuation value at other plants in the same portfolio.

4. The above simulation step results in a simulated "dataset" of state transitions through time. We update the state transition parameters by performing an AR(1) regression over this "dataset" for each plant $j \in \mathcal{J}_f$.[13]

5. Return to step 2 until the state transitions converge.

4.6 Why is it optimal to target inspections on firms with a worse history of violations?

A key determinant of the optimal targeting of firms is the specification of the benefit of polluting. In our model this is $\theta_j b(a_j) = \theta_j a_j^y$, where y is a parameter. To see why, consider a single plant firm and switch off dynamics (i.e. set $\beta = 0$). Therefore, this firm's problem is: $a_j^* = \arg\max_{a_j} \{\theta_j a_j^y - z_j a_j\}$, where z_j is the probability of inspection. In this case—so long as $y < 1$ (which is the case empirically) and z_j is held constant—the marginal deterrence benefit of an extra inspection is also increasing in θ_j $(\partial^2 a_j^* / \partial\theta_j \partial z_j)$. As a result, across the population of plants with varying θ_j but constant z_j, plants with more past violations are likely to have a higher marginal deterrence benefit.

While much of the intuition in the above example carries through to the estimated dynamic model, it is still possible for higher types to exhibit lower marginal deterrence because z_j is not constant across plants. Specifically, there is diminishing marginal deterrence $(\partial^2 a_j^* / \partial z_j^2 < 0)$ and higher type plants are inspected with higher probability. Therefore, it is ultimately an empirical question about whether this is the case.[14]

4.7 Operationalizing the decomposition

We operationalize the decomposition of the total effect of regulation into the correlated targeting and firm-wide moral hazard effects in the following way.

Correlated targeting effect We take the following steps to compute the correlated targeting effect, for each of the counterfactuals:

1. Compute the policy functions $a_j^*(\hat{\mathbf{s}}_j)$ and value functions $V_j(\hat{\mathbf{s}}_j)$ *at baseline* for each plant in

[13] To clarify – using the language from the comment on Step (e)(ii.) – this dataset uses the updated "true" values of $W_{j,t}$, not $\widehat{W}_{j,t}$.

[14] For example, suppose plant 1 is a "bad" plant with a high θ_1 and plant 2 is a less bad plant: $\theta_2 < \theta_1$. If plant 1 is in an industry that the regulator inspects more frequently than plant 2 (which could occur due to differences in perceived social costs of violations), then z_1 may be substantially higher than z_2, driving down the responsiveness of plant 1 to extra inspections to a level below the responsiveness of plant 2.

the data.[15]

2. Then, using the *counterfactual* regulator's policy function $\bar{z}_{g(j)}(s_{jt}, s_{ft})$, simulate the stationary distribution of actions, scores, and violations, across all plants in the data, using Step 3 in the computational algorithm in AppendixSection 4.5.

3. Record the change in the number of violations (weighted by the regulator's perceived social costs).

4. Compare this change in the perceived social cost of violations relative to the change in the perceived social cost of violations if the increase in the inspection budget was spent on untargeted inspections. This is the 'correlated targeting effect'.

    - Note that it is important to compare the counterfactual increase relative to spending the budget on untargeted inspections. This is because even an untargeted increase in inspections could reduce violations. Therefore, the benefit to correlated targeting is additional to this reduction.

Consistent with our intuitive description of this correlated targeting effect, if linked regulation directs the regulator to inspect plants j that are more likely to respond to the regulation (i.e. with a higher magnitude $\frac{\partial a_j^*(\hat{\mathrm{s}}_j)}{\partial z_j}$) then the correlated targeting effect will have a higher magnitude.

Firm-wide moral hazard effect We then compute the firm-wide moral hazard as the residual of the total effect, after subtracting out the correlated targeting effect. This isolates the remaining channel by which the model allows for a response to regulation: by changing the mapping of equilibrium scores to actions $a_j^*(\hat{\mathrm{s}}_j)$ (note that there are also second-order effects on the equilibrium distribution of scores). For example, for a given state $\hat{\mathrm{s}}_j$, the firm will now internalize that the probability of an inspection has increased, and so take a less negligent action.

Table A-7: First-stage estimates

Variable	Estimate	Std. Err.	Variable	Estimate	Std. Err.
Inspection probability function			*Law of motion for plant scores*		
Plant Rating	0.109	0.017	Lagged plant rating	0.786	0.003
Firm Rating	0.056	0.018	Lagged # violations	9.0	(Calibrated)
Intercept			*Complexity Weights*		
Utility	-0.853	0.034	Utility	1.0	(Normalized)
Services	-1.165	0.041	Services	1.432	0.042
Manufacturing	-0.743	0.035	Manufacturing	1.476	0.034
Resources	-1.362	0.043	Resources	0.887	0.022
Transportation	-1.224	0.054	Transportation	0.219	0.027
Trade	-1.264	0.057	Trade	0.941	0.032

4.8 First stage estimates

5 Model robustness and additional counterfactuals

5.1 To what extent is correlation in compliance costs within the same industry driving the value added from linked regulation?

Suppose that the complexity weights that the regulator uses are outdated due to an industry-level cost shock and, for the period of our sample, one particular industry has relatively high compliance costs that are not captured through the "fixed" parts of the regulation (e.g. the complexity weight system and also industry-specific differences in untargeted inspections).[16] For plants in this industry, the corresponding firm score will then be relatively higher, nudging the regulator to inspect these firms more. Overall, this may help the regulator adjust to this shock.

With this as background, our robustness test involves switching off industry heterogeneity by setting the $\bar{\theta}_g$ to be equal to the average across industries. We then rerun all of our counterfactuals and investigate the results. Clearly, industry-level compliance costs may differ for a variety of reasons beyond the scenario described above. However, shutting off industry-specific heterogeneity is still informative because it provides an upper bound on how much industry-specific differences matter to the results.

[15] In the code, the estimation routine conveniently delivers equilibrium policy functions $a_j^*(\hat{s}_j)$ and value functions $V_j(\hat{s}_j)$ as objects that interpolate over a set of nodes given any $\hat{s}_j$. So this step is relatively simple.

[16] This example could apply more generally to technology or demand shocks. We have in mind the example where the time-varying shock happens after the "fixed" parts of the regulation are determined, but before the first year in our sample of 2012.

Table A-8: Value of linked regulation at mean industry type

	$\uparrow$ Inspections budget by 10%. Spent on:		
	Linked reg.	Unlinked reg.	50/50 mix
%Δ Social cost vs random:			
Using regulator's preferences:			
Mean industry type	74.7%	54.5%	74.4%
Baseline	76.2%	55.3%	74.4%
Using pollution index:			
Mean industry type	74.0%	54.1%	73.7%
Baseline	76.7%	55.3%	74.6%

Notes: These results shutdown industry-level heterogeneity in compliance types by setting the industry shifters $\bar{\theta}_g$ for each industry g, equal to the average. "Using regulator's preferences" corresponds to measuring the social cost of violations using the estimated "perceived" social cost of violations. "Using pollution index" corresponds to measuring the social cost of violations using the pollution index, constructed with weight 0.5.

We present these results in Table A-8. Across all the counterfactuals, and regardless of whether we look at the results weighted by the "perceived" costs of violations or the pollution index, relative differences in industry compliance costs do not seem to be driving the main results. Specifically, the correlation in compliance costs within firms' portfolios driven by common industry compliance costs accounts for a positive but small share of the benefits of linked regulation. However, when considering the two mechanisms, the one percent difference is a notable share of the correlated targeting mechanism.

5.2 Are random crackdowns optimal?

We simulate a counterfactual in which the regulator employs a "random crackdown" in the spirit of Eeckhout et al. [2010]. A random crackdown involves an increase in the intensity of regulation for an arbitrarily-selected subset of agents in a population. Eeckhout et al. [2010] show that a crackdown can be optimal if there are increasing marginal returns to deterrence—so focusing enforcement resources intensely on a small group of agents can more effectively deter them.

We operationalize this counterfactual by partitioning firms into two groups: a crackdown group and a holdout group. We order firms from largest portfolio (most plants) to smallest. Then, we

Table A-9: Evidence of diminishing marginal deterrence

Dependent Variable	(1) Violations
Pr(Inspection)	-10.035
	(2.198)
Pr(Inspection)^2	6.942
	(2.473)
Year FEs	Yes
Unit FEs	Plant
Only Inspected	Yes
N	10531

assign every other firm according to this ordering (i.e. the largest, the third-largest, the fifth-largest, and so on) to the crackdown group. We then increase the inspection budget by 10 percentage points by increasing the intercept in the inspection rule—i.e, increasing untargeted inspections—for only the plants assigned to the crackdown group, until equilibrium average inspection probabilities are 10 percentage points larger than baseline. This is a budget increase comparable to our main counterfactuals.

We compute the regulator's perceived social cost under this scheme. We find that perceived social costs are 19% higher under the random crackdown scheme than under a 10 percentage point increase in the budget allocated to untargeted inspections. Crackdowns perform worse than an increase in the budget devoted to untargeted inspections, and therefore also perform worse than our unlinked and linked regulation counterfactuals. We conclude random crackdowns would not be effective in our setting.

We provide intuition underlying this result in the stylized model in Appendix 6. The reason is that, unlike in Eeckhout et al. [2010], actions are convex in the inspection probability (i.e., there is diminishing marginal deterrence). We provide model-free evidence of this in Table A-9: the expected number of violations is decreasing, but convex, in inspection probability.

5.3 Counterfactual results by sector

We break out the counterfactual results by NAICS sector in Table A-10. A key rationale for providing this table is that it illustrates robustness of the results to the specification of, and estimates from, the regulator's problem. The regulator's problem is only used to estimate the perceived social costs of a violation by sector, with all other estimates identified and computed independently using the firm's problem. These perceived social cost estimates are then used to aggregate violations across sectors into a single statistic. Hence, disaggregating counterfactual change in violations by sector shows whether the results hold making only assumptions on the firm side of the model, and not the regulator.

The numbers in Table A-10 center on the headline results for multi-plant firms. We find that the conclusion that linked regulation does substantially better than unlinked regulation holds in every sector. Overall, these results show that the qualitative conclusion that linked regulation adds significant value is relatively robust across sectors and not driven by assumptions on the regulator's problem.

Table A-10: Counterfactuals: multi-plant firms only, split by sector

	$\uparrow$ Inspections budget by 10%. Spent on:		
	Linked reg.	Unlinked reg.	50/50 mix
%Δ Violations vs random: multi-plant firms			
Manufacturing	57.7%	49.5%	58.9%
Resources	31.2%	18.5%	33.0%
Services	38.1%	27.0%	40.2%
Trade	53.5%	40.2%	55.9%
Transportation	199.5%	122.4%	186.0%
Utility	109.1%	75.5%	102.6%

5.4 Should the regulator use longer histories?

We simulate counterfactuals in which the regulator uses longer histories to inform scores in the following sense: we change the autoregressive coefficient in the AR(1) process for scores ($r_{s_j}^{s_j}$ in Equation 8). Compared to the estimated autoregressive coefficient (0.786), less persistent scores perform worse. For instance, a scoring rule with persistence of 0.5—corresponding roughly to a rule where histories are "forgotten" 1.5 times as fast as in actual practice—leads social costs to be 25.5% higher. Somewhat higher persistence than actually practiced improves the efficiency of regulation, but at extreme levels of persistence, adding more persistence makes regulation less

Table A-11: Changing the persistence of scores

Persistence	Harms rel. to baseline
0	194.4%
0.5	125.5%
Baseline (0.786)	100%
0.9	90.1%
0.95	85.1%
0.99	96.5%

efficient. For instance, an AR(1) process with autoregressive coefficent 0.95 performs better than an AR(1) process with autoregressive coefficient 0.99.

These counterfactuals reflect the intuition in the theoretical analysis of Hörner and Lambert [2021]. Adding more persistence presents a trade-off: longer histories mean scores encode more information and make targeting more efficient, but if histories become too long, scores lose sensitivity to present-day violations, and they lose their ability to deter bad behavior. In the limit, a scoring rule that reflects an infinitely long history without any forgetting will not be affected at all by the present period and provides no deterrence effect.

5.5 Descriptive and model results

Table A-12: Firms with worse histories are more responsive to regulation

	(1)	(2)
Dependent Variable	Violations	Violations
Plant Types	Low-violation	High-violation
Pr(Inspection)	-0.539	-3.189
	(0.693)	(0.832)
Year FEs	Yes	Yes
Unit FEs	Plant	Plant
Only Inspected	Yes	Yes
N	4923	15553

Note: The dependent variable ‘Violations’ corresponds to violations conditional on inspection (since we only look at periods where plants were inspected). We show in these regressions that plants with higher overall propensities to violate (likely high-type plants) are more responsive to changes in the probability of inspection than plants with low propensities to violate. Low-violation plants are those below the median violations per inspection, high-violation plants are those above the median violations per inspection. There are more observations in the high-violation regression (column (2)) because plants with violations above the median are inspected more frequently.

Table A-13: Penalties scale approximately linearly in number of violations

Dependent Variable	(1) Penalty ($1000s)	(2) Penalty ($1000s)	(3) Penalty ($1000s)
# Violations	4.9	5.8	5.7
	(1.9)	(1.9)	(2.0)
# Violations Squared	-0.082	-0.098	-0.089
	(0.055)	(0.055)	(0.055)
Year FEs	No	Yes	Yes
NAICS Category FEs	No	No	Yes
N	15954	15954	15954

Note: These regressions show that while the linear term is significant, the quadratic term is not. For example, in our preferred specification in column (3), the quadratic term is not significant even at the 90% confidence level. These result support the modeling assumption that penalties scale approximately linearly in the number of violations.

Table A-14: Regression of violations per inspection at one plant vs. other co-owned plants

Dependent Variable	(1) Leave-out avg viols	(2) Leave-out avg viols
Log viols per inspection	.103	.074
	(.007)	(.007)
NAICS category FEs	No	Yes
N	2933	2933

Note: These regressions show that the within-firm correlation of violations across plants is significant and robust to controls for observables.

5.6 What if the regulator cares about compliance costs?

In our main specification, the regulator's objective is to minimize its perceived social harms of pollution, subject to a budget constraint on inspections. Our estimates in Table 6 are in terms of these perceived social harms. Here, we explore an alternative model in which the regulator also cares about the plants' costs of compliance with the regulation in addition to the social harms of pollution.

We begin by noting that plants' payoffs are quasi-linear in penalties, so the flow payoff of (non)compliance, $\theta_j b(a_j)$, is in dollars.[17] We can then re-estimate an alternative version of the regulator's problem

[17]Recall that in estimation we then take the additional step of normalizing the baseline penalty per violation to 1.0. This is without loss of generality because we did not normalize the scale of plant types θ in the firms' problem, and

with an alternative objective function that incorporates both plants' private abatement costs as well as the social costs of a violation. For each plant j, define "net social cost" as:

$$NSC_j(a_j) = \underbrace{h_{g(j)}\mathbb{E}_{v_j}(v_j|a_j)}_{\text{Social cost of violations}} - \underbrace{\theta_j b(a_j)}_{\substack{\text{Negligence benefits} \\ \text{/abatement costs}}} \tag{A-6}$$

$$= h_{g(j)}a_j - \theta_j b(a_j). \tag{A-7}$$

where the second equation follows since violations are drawn from a Poisson distribution with mean a_j. The regulator's problem is then identical to Equation 9 except with $NSC_j(a_j)$ in place of $h_{g(j)}a_j$ in the objective function.

Our main counterfactuals only depend on the *ratio* of perceived marginal social harms across industries; therefore we normalize $h_{g(j)}$ for one industry—utilities—to 1. In this case, we must also estimate the ratio of marginal social harms to private abatement costs for a single industry. We choose a ratio of $0.4533 = 1431.0/3157.1$ for plants in the industry "utility", which is derived from Kang and Silveira [2021] who estimate social costs and private abatement costs for the Clean Water Act in California for wastewater plants (a type of utility). We also test robustness by adding and subtracting 20% to this ratio and re-computing estimates and counterfactuals under these alternative assumptions.

With our new estimates, we re-compute all counterfactuals from Table 6. The re-computed counterfactuals are in Table A-15. Quantitatively, including these abatement costs in the results tends to slightly reduce the benefits of both linked regulation and unlinked regulation compared to an untargeted allocation of inspections. Intuitively, both linked and unlinked regulation disproportionately target plants who are the worst violators, and these plants also happen to have the highest abatement costs. Although the results illustrate that it is still optimal to target these plants, including abatement costs in the results reduces the implied benefits by a relatively small amount.

---

the firms' private costs did not directly enter into overall costs and benefits of the regulation in the counterfactuals. However, in this robustness test we need to modify our approach slightly (as we explain below), because we need to also account for the dollar value of private abatement costs in the counterfactuals.

Table A-15: Robustness: regulator internalizes firms' abatement costs

	↑ Inspections budget by 10%. Spent on:		
	Linked reg.	Unlinked reg.	50/50 mix
%Δ Social cost vs untargeted:			
Baseline	76.2%	55.3%	74.4%
Including abatement costs	69.1%	49.8%	67.8%
Including abatement costs (ratio+20%)	67.3%	48.3%	66.2%
Including abatement costs (ratio-20%)	70.8%	51.1%	69.4%

Notes: The"baseline" case corresponds to the case where the regulator's objective is to minimize only the "perceived" social cost of violations; this corresponds to the main results in the paper. The case "including abatement costs" corresponds to the case where the regulator minimizes the total net social costs of a violation, which incorporates both the total social costs of a violation as well as plants' private abatement costs.

6 Two-period stylized model

In this Appendix, we develop a stylized two-period model that incorporates key features of our empirical setting. A representative firm owns two plants that have private pollution types. They may either be high pollution (H) or low pollution (L); types may be correlated through a parameter ρ. In each period t, the firm chooses hidden polluting actions a_{it} (which generate violations v_{it} with probability a_{it}) at each plant i. The regulator inspects plants in each period; if a plant is inspected, violations are recorded.

Actions are influenced by a regulator's policy, denoted $\bar{z}$, to which the regulator commits at the outset of the game. $\bar{z}$ governs the probability that a plant is inspected in period 2 based on whether violations were observed in period 1. For example, $\bar{z}(v_{-i} = 1)$ is the probability of inspecting plant i in period 2 if plant $-i$ (the other plant the firm owns) violated in the previous period. $\bar{z}$ allows for linked and unlinked escalations, as well as untargeted inspections (i.e., inspections that are not contingent on histories).

6.1 Setup

Agents. There are many firms and one regulator. Each firm has two plants. We focus on the interaction between the regulator and a representative firm, whose plants are indexed by i. We generically use i to refer to the focal plant and $-i$ to refer to the other plant. Each plant is assigned a type $g_i \in \{L, H\}$ where L is a low-pollution type and H is a high-pollution type. The type profile for the firm is $g = (g_i, g_{-i})$.

Timeline. The following timeline gives the order of events.

$t = 0$	(i)	Nature draws plant types from a distribution we describe later. The regulator does not observe types but the firm does.
	(ii)	The regulator chooses inspection policy $\bar{z} : s_i \rightarrow [0, 1]$ where s_i is plant i's state.
$t = 1$	(i)	The firm chooses actions $a_{i1} \in [0, 1]$ and $a_{-i1} \in [0, 1]$. In equilibrium, this will depend on g and $\bar{z}$. Let the function $a_1(\bar{z}\vert g_i, g_{-i})$ be the firm's optimal action for plant i given the inspection policy, the plant's type, and the co-owned plant's type.
	(ii)	The regulator inspects exactly one plant owned by the firm. Each plant is inspected with probability $\frac{1}{2}$.
	(iii)	If i is inspected, it violates ($v_{i1} = 1$) with probability a_{i1} and is assessed a penalty of 1. With probability $1 - a_{i1}$, it does not violate ($v_{i1} = 0$) and no penalty is assessed. Same for plant $-i$. If a plant is not inspected, no penalty is assessed on that plant.
	(iv)	State s_i is updated to record which plant was inspected and whether a violation occurred. We next discuss states in more detail.
$t = 2$	(i)	The firm chooses actions $a_{i2} \in [0, 1]$ and $a_{-i2} \in [0, 1]$. Let function $a_2(\bar{z}\vert g_i, g_{-i}, s_i)$ be the firm's optimal action for plant i and let $a_2(\bar{z}\vert g_{-i}, g_i, s_{-i})$ be the optimal action for plant $-i$.
	(ii)	Plant i is inspected with probability $\bar{z}(s_i)$ and $-i$ is inspected with probability $\bar{z}(s_{-i})$. Conditional on the state, the Bernoulli draws of whether each plant is inspected are independent.
	(iii)	Conditional on being inspected, plant i violates ($v_{i2} = 1$) with probability a_{i2} and is assessed a penalty of 1, and it does not violate ($v_{i2} = 0$) with probability $1 - a_{i2}$ and assessed no penalty. Same for plant $-i$.

States. The state s_i encodes the outcome of period 1. Because the state will only depend on violations in the first period, we drop the subscript that indicates $t = 1$ from v_{i1} when describing the state. Since the regulator inspects exactly one plant, the state s_i can take one of four values: $s_i = (v_i = 1)$ indicates that plant i was inspected in period 1 and a violation was uncovered. $s_i = (v_{-i} = 1)$ indicates that plant $-i$ was inspected in period 1 and a violation was uncovered. $s_i = (v_i = 0)$ and $s_i = (v_{-i} = 0)$ indicate plant i was inspected and did not commit a violation, and that plant $-i$ was inspected and did not commit a violation, respectively.

Note that, for instance, $s_i = (v_i = 1)$ if and only if $s_{-i} = (v_{-i} = 1)$. $\bar{z}$ is therefore characterized by the vector in $[0, 1]^4$ that governs the probability of inspection for each state. For instance, $\bar{z}(v_i = 1)$ is the probability that plant i is inspected in period 2 after committing a violation in period 1. We adopt the convention $\bar{z} = [\bar{z}(v_i = 1), \bar{z}(v_i = 0), \bar{z}(v_{-i} = 1), \bar{z}(v_{-i} = 0)]$.

Types and plant payoffs. g is the realization of a random variable denoted by G. The distribution of G is governed by parameter $\rho \in [0, 1]$. The probability $\Pr(G = (H, H)) = \frac{\rho}{2}$, and the probability $\Pr(G = (L, L)) = \frac{\rho}{2}$. The probability $\Pr(G = (H, L)) = \frac{1-\rho}{2}$ and $\Pr(G = (L, H)) = \frac{1-\rho}{2}$. Therefore, if $\rho = .5$, plant types are independent. If $\rho = 1$, plant types are perfectly correlated.

The flow payoff for a plant that will be inspected with probability z is $b_g(a) - za$, where $b_g(a)$ is the flow benefit from polluting action a. We assume L-type plants get a flow benefit of 0, while H-type plants get a flow benefit $b_H(a)$ that is increasing, smooth, and strictly concave in a.[18] There is no discounting.

Regulator payoffs. The regulator chooses $\bar{z}$ to minimize expected actions across many ex ante identical firms over the course of the two periods, subject to a budget constraint B on inspections in period 2 and optimal behavior by the firms:

$$\begin{gathered} \bar{z}^* = \text{arg min}_{\bar{z}}\ \mathbb{E}_G\left[\sum_i a_{i1} + a_{i2}\right] \\ s.to \quad \mathbb{E}_G\left[\sum_{i,s_i} \Pr(s_i)\bar{z}(s_i)\right] \leq B, \qquad a_{i1} = a_1(\bar{z}|g_i, g_{-i}), \qquad a_{i2} = a_2(\bar{z}|s_i, g_i, g_{-i}) \end{gathered} \tag{A-8}$$

Examining Equation A-8, there are two key elements: equilibrium actions, $EA(\bar{z}) \equiv \mathbb{E}_G\left[\sum_i a_{i1} + a_{i2}|\bar{z}\right]$ and equilibrium inspections $EI(\bar{z}) \equiv \mathbb{E}_G\left[\sum_{i,s_i} \Pr(s_i)\bar{z}(s_i)\right]$. The Kuhn-Tucker conditions for an

[18]This ensures that optimal actions are unique and vary smoothly with the components of $\bar{z}$.

interior solution are

$$-\frac{\partial EA(\bar{z})}{\partial z} \bigg/ \frac{\partial EI(\bar{z})}{\partial z} = \lambda, \; \forall \, z \in \bar{z}.$$

Denote the *marginal equilibrium deterrence per inspection* of component z evaluated at policy $\bar{z}$ as $-\frac{\partial EA(\bar{z})}{\partial z} \big/ \frac{\partial EI(\bar{z})}{\partial z} = MEDI_z(\bar{z})$. At any interior optimum, $MEDI_z(\bar{z})$ must be the same for all components of $\bar{z}$.[19] The optimal action in period 2 for a plant of type L is zero, while the optimal action in period 2 for an H type plant is: $a_2(\bar{z}|H, g_{-i}, s_i) = \arg\max_a b_H(a) - \bar{z}(s_i)a$. This does not depend on g_{-i} and only depends on s_i through the inspection probability $\bar{z}(s_i)$, so we write $a_{i2} = a_2(\bar{z}(s_i))$ if the plant is type H and $a_{i2} = 0$ if it is type L.

6.2 Correlated targeting versus firm-wide moral hazard.

In Section 7.3 of the paper, we decompose the effects of regulation into two parts: a "firm-wide moral hazard" mechanism and a "correlated targeting" mechanism. The "correlated targeting" mechanism measures the change in actions stemming from a change in the regulator's policy function, holding firms' response functions (mapping scores and types to actions) fixed. The "moral hazard" mechanism measures the additional change in equilibrium actions from changing firm policy functions.

Our stylized model, which replaces scores with states s_i, lends support to this decomposition. The expression for the marginal effect of increasing $\bar{z}(v_{-i} = 1)$—that is, the effect of targeting the "linked" co-owned plant—on equilibrium expected actions can be decomposed into the two mechanisms:

$$\frac{\partial EA(\bar{z})}{\partial \bar{z}(v_{-i}=1)} = \rho\Bigg(\underbrace{\left(1 + 0.5\Delta^{H,H}_{a_2}\right) \frac{\partial a_1(\bar{z}|H,H)}{\partial \bar{z}(v_{-i}=1)}}_{\text{Firm-wide moral hazard}} + \underbrace{\overbrace{0.5a_1(\bar{z}|H,H)}^{\text{Prob. violation detected}} \frac{\partial a_2(\bar{z}(v_{-i}=1))}{\partial \bar{z}(v_{-i}=1)}}_{\text{Correlated targeting}} \Bigg) \quad \text{(A-9)}$$

where $\Delta^{H,H}_{a_2} \equiv a_2(\bar{z}(v_i = 1)) + a_2(\bar{z}(v_{-i} = 1)) - a_2(\bar{z}(v_i = 0)) - a_2(\bar{z}(v_{-i} = 0))$. Equation A-9 contains two terms. The first term captures moral hazard by modeling the consequences of increased deterrence on firms' (hidden) actions. This is analogous to the "firm-wide moral hazard" mechanism in our empirical decomposition. The second term captures the effect of increased

[19] Recall that there are a large number of firms, and so it is possible to transfer part of the inspections budget from firms in one state to firms in another state. In other words, this is why these conditions are with respect to the *expected* inspections and actions, where the expectation is over the distribution of firms.

targeting and is analogous to the "correlated targeting" mechanism in our empirical decomposition. We describe each term in more detail below.

The first term captures the effect of the threat of escalation in period 2 on actions in period 1. Period 1 actions change even though the probability of inspection in period 1 does not change, so this term is analogous to changing firms' equilibrium policy functions (mapping states to actions) in our empirical model. This term also captures the change in the equilibrium distribution of states (and, therefore period 2 inspection probabilities and actions) that arises from a change in policy functions. Because this captures the effect of changes in firms' (hidden) actions, we call these a "firm-wide moral hazard" mechanism.

The second term does not rely on the threat of escalation. Instead, it captures the effect of changes in inspection probabilities in period 2 on actions in that same period. It is driven by the assumption that plants with more observed violations tend to be more responsive, which arises from the assumption in this stylized model that high type plants are responsive to regulation whereas low type plants uniformly choose actions of zero.[20] Therefore, we call this a "correlated targeting" effect. Our decomposition isolates this effect by examining a change in the inspector's policy function without changing firm strategies.

Three additional notes on this expression: first, the "firm-wide moral hazard" effect relies on the regulator's ability to commit to the policy function $\bar{z}$ at the outset of the game. The "correlated targeting" effect, on the other hand, does not rely on commitment.[21] Second, whereas the correlated targeting effect relies on plants with more violations being more responsive to regulation, the firm-wide moral hazard effect does not. So long as firms are penalized for violating in period 1 with higher inspection probability in period 2, they will be deterred from violating in period 1. Third, this expression is nonzero only for (H, H) firms.[22] Therefore, the effects of linked regulation are increasing in the correlation of types ρ.

[20]This is a feature of our empirical setting as well. We justify our assumption that high type plants are more responsive to regulation with model-free evidence in Table A-12.

[21]We explore the consequences of commitment via an alternative no-commitment model in the next subsection.

[22]For (H, L) firms, L plants choose $a = 0$, so $s_i = (v_{-i} = 1)$ cannot happen for the H plant in this portfolio, and therefore $z(v_{-i} = 1)$ is not relevant to the firm's problem.

6.3 Example and analysis

For this example, we set $b_H(a) = a(1 - \frac{2\sqrt{a}}{3})$ which yields $a_2(z) = (1-z)^2$. This implies static deterrence ($a_2(z)$ decreasing in z) and decreasing static marginal deterrence ($a_2(z)$ convex in z). The continuation value is $U(z) = \frac{(1-z)^3}{3}$. In period 1, the optimal actions are:

$$
\begin{aligned}
a_1(\bar{z}|H,H) &= \frac{1}{4}\left[U(\bar{z}(v_i=1)) + U(\bar{z}(v_{-i}=1)) - U(\bar{z}(v_i=0)) - U(\bar{z}(v_{-i}=0))\right]^2 \\
a_1(\bar{z}|H,L) &= \frac{1}{4}\left[U(\bar{z}(v_i=1)) - U(\bar{z}(v_i=0))\right]^2 \\
a_1(\bar{z}|L,L) &= 0
\end{aligned}
$$

We consider the value of linked and unlinked regulation. Linked regulation corresponds to increasing $\bar{z}(v_i = 1)$ and $\bar{z}(v_{-i} = 1)$ simultaneously at the same rate, which is roughly equivalent to increasing the importance of a firm-wide score in our empirical setting. Unlinked regulation corresponds to increasing only $\bar{z}(v_i = 1)$ which roughly corresponds to increasing the importance of the plant-level score. Finally, we consider an increase in untargeted inspections, which implies all four components of $\bar{z}$ increase at the same rate. Marginal deterrence per inspection for each of these three policy levers corresponds to the following expressions:

$$
\begin{aligned}
MEDI_{\text{unlinked}}(\bar{z}) &= -\frac{\partial EA(\bar{z})}{\partial \bar{z}(v_i=1)}\Big/\frac{\partial EI(\bar{z})}{\partial \bar{z}(v_i=1)} \\
MEDI_{\text{linked}}(\bar{z}) &= -\frac{\frac{\partial EA(\bar{z})}{\partial \bar{z}(v_i=1)} + \frac{\partial EA(\bar{z})}{\partial \bar{z}(v_{-i}=1)}}{\frac{\partial EI(\bar{z})}{\partial \bar{z}(v_i=1)} + \frac{\partial EI(\bar{z})}{\partial \bar{z}(v_{-i}=1)}} \\
MEDI_{\text{untargeted}}(\bar{z}) &= -\frac{\frac{\partial EA(\bar{z})}{\partial \bar{z}(v_i=1)} + \frac{\partial EA(\bar{z})}{\partial \bar{z}(v_{-i}=1)} + \frac{\partial EA(\bar{z})}{\partial \bar{z}(v_i=0)} + \frac{\partial EA(\bar{z})}{\partial \bar{z}(v_{-i}=0)}}{\frac{\partial EI(\bar{z})}{\partial \bar{z}(v_i=1)} + \frac{\partial EI(\bar{z})}{\partial \bar{z}(v_{-i}=1)} + \frac{\partial EI(\bar{z})}{\partial \bar{z}(v_i=0)} + \frac{\partial EI(\bar{z})}{\partial \bar{z}(v_{-i}=0)}}
\end{aligned}
\tag{A-10}
$$

Optimally, all three of these measures should be the same. We illustrate using a baseline value of $\rho = .75$, which indicates a moderate level of correlation in plant types.

Optimal degree of targeting. In the example in Figure A-1, we show there are diminishing marginal benefits to unlinked escalations above a baseline of 10% untargeted inspections. Unlinked escalation is an increase in the baseline in the probability of inspecting a plant that violated in the first period. We compare the marginal equilibrium deterrence per inspection of unlinked

Figure A-1: Marginal deterrence per inspection of unlinked regulation is higher than untargeted inspections at first but declines as it is used more

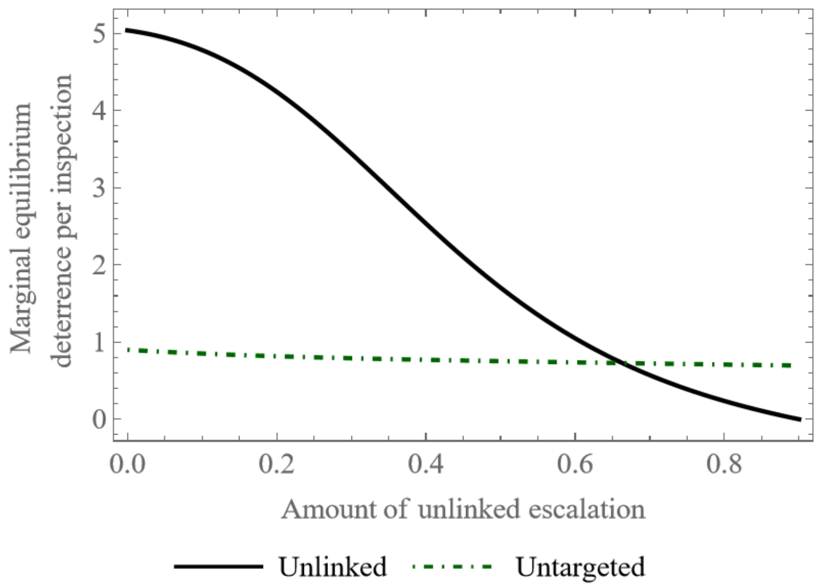

Notes: On the horizontal axis, the left-most point corresponds to the baseline regulator's policy. This baseline policy is a 10% probability of an untargeted inspection in the second period, and no unlinked escalations (i.e., plant i is inspected with 10% probability in the second period regardless of whether it committed a violation in the first period). Moving along the horizontal axis from this point, we increase the amount of unlinked regulation (i.e., plant i is inspected with $(10 + x)\%$ probability if it violated in period 1). We plot the corresponding marginal equilibrium deterrence per inspection for untargeted and unlinked regulation on the vertical axis.

escalation on this plant to the marginal equilibrium deterrence per inspection of an untargeted inspection. Unlinked escalations are more effective at the margin up to a point (here, roughly 65 percentage points). Beyond this point, additional inspections are better off allocated in an untargeted manner. Therefore, while some degree of unlinked regulation is optimal, it is possible to over-target at the plant level. Formally, we do this by setting

$$\bar{z} = [\bar{z}(v_i = 1), \bar{z}(v_{-i} = 1), \bar{z}(v_i = 0), \bar{z}(v_{-i} = 0)] = [.1 + z, .1, .1, .1] \quad \text{(A-11)}$$

and computing $MEDI_{\text{untargeted}}(\bar{z})$ and $MEDI_{\text{unlinked}}(\bar{z})$ as we vary z from 0 to 0.9.

Unlinked versus linked regulation. Next, we compare the value of unlinked and linked regulation. In the top panel of Figure A-2 we compare the marginal equilibrium deterrence per inspection of unlinked regulation (increasing the probability of inspecting a plant that violated in period 1) to that of linked regulation (increasing the probability of inspecting both plants owned by the firm after one violates in period 1). We then vary the degree of linking from only plant-level escalations to only firm-wide escalations. The optimal degree of linking—where linked and unlinked regulation have the same marginal deterrence per inspection—is interior.

In more detail, we compute and plot $MEDI_{\text{unlinked}}(\bar{z})$ and $MEDI_{\text{linked}}(\bar{z})$ evaluated at $\bar{z} = [.5, .1 + z, .1, .1]$ from $z = 0$ to $z = 0.4$. Here, z is an index of the degree of linking that is normalized to lie between 0 and 1. When $z = 0$, there is no linking. When $z = 0.4$, $\bar{z}(v_i = 1) = \bar{z}(v_{-i} = 1) = 0.5$—i.e., both plants are inspected with probability 0.5—there is only firm-wide linked regulation.

Role of correlation. Generally, as plant types become more correlated (i.e., ρ increases), linked regulation becomes more effective. We demonstrate this by taking the same analysis underlying the top panel of Figure A-2 and increasing the degree of correlation in plant types in the bottom panel. Specifically, the probability that one plant is a "high" type given the other plant is a "high" type increases from 0.75 to 0.95. The point at which the marginal deterrence per inspection of linked and unlinked regulation are equal—that is, the optimal degree of linking—is significantly higher when the correlation is higher.

Role of commitment. We can study commitment by allowing the regulator to freely change $\bar{z}$ after period 1. The regulator would want to set $\bar{z}(s_i)$ to optimally minimize actions in period 2 given its interim beliefs after period 1. An uncommitted regulator (who still faces a budget constraint) will set $\bar{z}(s_i)$ equal to

$$\begin{aligned} \bar{z}(s_i) = \quad & \arg\min_z \Pr(i = H | s_i, \bar{z})\, a_2(z) \\ \text{s. to} \quad & \mathbb{E}_G \left[\sum_{i, s_i} \bar{z}(s_i) \middle| \bar{z} \right] \leq B. \end{aligned}$$

The interim belief $\Pr(i = H | s_i, \bar{z})$ is based on Bayes' law and depends on the firm's strategy. For instance, because $a_1 = 0$ for all type-L plants, if $s_i = (v_i = 1)$, then $\Pr(i = H | s_i, \bar{z}) = 1$. That is, if the regulator observes a violation, it knows that the plant must be type H. If $s_i = (v_i = 0)$, beliefs are more complicated, generally between 0 and 1, and depend on the equilibrium values of $a_1(\bar{z} | H, H)$ and $a_1(\bar{z} | H, L)$.

Generally, violations in period 1 imply a higher belief that a plant is an H type, and higher beliefs imply that the regulator expects deterrence to have a greater effect in period 2. Therefore, the

Figure A-2: Marginal equilibrium deterrence per inspection of linked regulation is high at first but declining; more makes linking more effective

(a) Moderate correlation in types

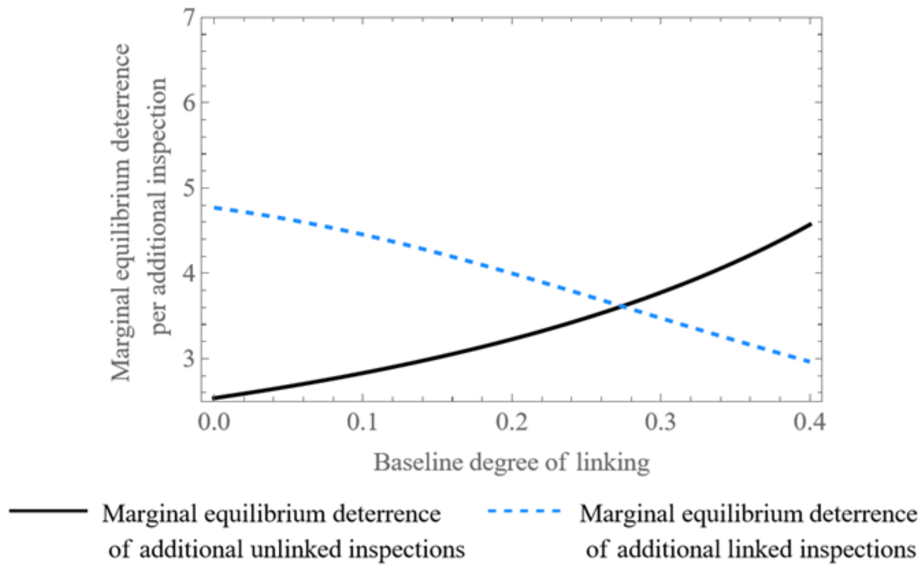

(b) High correlation in types

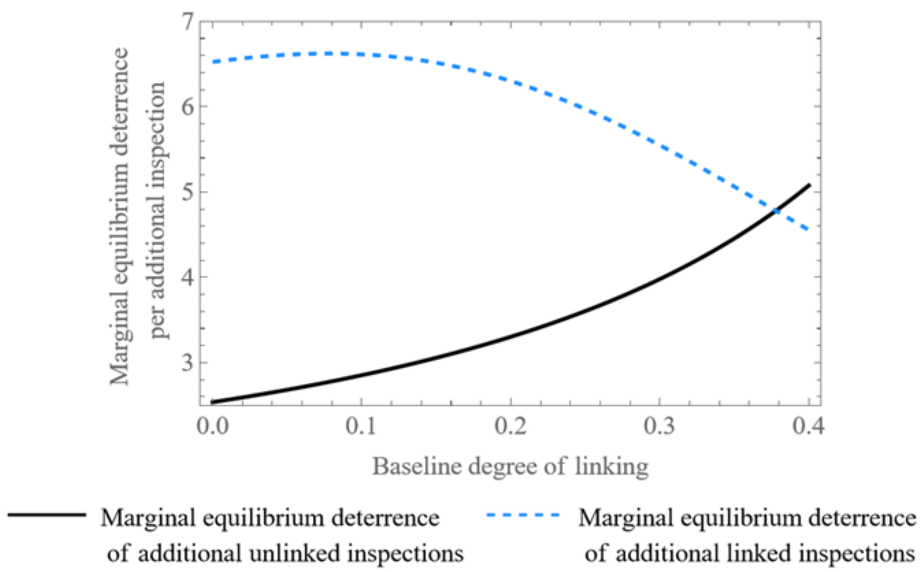

Notes: The solid black line reports the marginal equilibrium deterrence (across both plants) from increasing unlinked regulation (i.e. additional inspections that don't spill over to other sibling plants in a multi-plant firm). The dashed blue line reports the marginal equilibrium deterrence (across both plants) from increasing linked regulation. The x-axis varies the baseline policy at this marginal equilibrium deterrence is evaluated. The left-most point corresponds to a baseline policy whereby a plant that violated in period 1 is inspected with 50% probability and all other plants—including plants co-owned with a plant that violated in period 1—are inspected with 10% probability. The right-most point corresponds to a policy where both a plant that violated in period 1 and plants co-owned with such a plant are inspected with 50% probability in the baseline. Therefore, the x-axis represents a smooth transition from a baseline of plant-level (unlinked) escalation to a baseline of firm-level (linked) escalation. For instance, the solid black line, evaluated at $x = 0.1$, is the marginal equilibrium deterrence per inspection from an increase in unlinked regulation, when the policy is to inspect a plant that violated in period 1 with 50% probability and to inspect a plant that is co-owned with a plant that violated in period 1 with $(10 + 10)\% = 20\%$ probability.

regulator will allocate more of its inspection budget to states where beliefs are higher. This means that escalating—allocating more inspections to states where violations have occurred and beliefs are higher—is optimal for an uncommitted regulator, just as it is for a committed regulator.

Role of private types. We next study the role that private types play in this problem. Suppose the regulator could observe and contract on types. Inspection probabilities are thus a function of types i.e. $\bar{z}_H(s)$ for types H and L, respectively. Because L-type plants always choose actions of 0, the regulator should devote as much of its budget as possible toward deterring H-type plants. Only if there is extra room in the budget would a regulator set $\bar{z}_L(s) > 0$ for any s. Even in this case, the inspector would be indifferent between setting $\bar{z}_L(s) > 0$ and leaving the budget slack. Contrast this to the case where the regulator cannot contract on types, and the optimal level of linked and unlinked regulation is generally interior. Furthermore, there is no role for conditioning inspection probabilities on histories when the regulator can contract on types.

Crackdowns. We examine the role of optimal crackdowns as in Eeckhout et al. [2010], who show that "random crackdowns"—where a randomly-chosen subset of agents is inspected with higher probability than others—can be efficient depending on the curvature of polluting (or otherwise illegal) actions as a function of inspection probability. If pollution is concave (convex) in inspection probability, then there are increasing (decreasing) marginal returns to deterrence, and crackdowns are (not) optimal.

In this model, the optimality of crackdowns also relates to the curvature of equilibrium expected actions as a function of inspection probabilities. In the example of this subsection, the static optimal action is convex in z ($a(z) = (1 - z)^2$). We show in Figure A-3 that the expected *equilibrium* action is convex. This persists even in the presence of linked and unlinked regulation (as an example, we set the firm-wide escalation probability to 0.4 and the plant-level escalation probability to an additional 0.2 in Figure A-3). We also plot the expected equilibrium action for a specification where the static optimal action is concave, and show the expected equilibrium action is also concave. In summary, if the static optimal action is concave (convex) in inspection probability z, the equilibrium expected action is also likely to be concave (convex) in z, and crackdowns are likely (unlikely) to be efficient.

In our empirical model, plants' optimal static actions are concave in the probability a plant is inspected. We expect this concavity to generally flow through to the plant policy functions in a fully dynamic setting. Therefore, crackdowns are not optimal in our setting.

Furthermore, even if plants' actions were convex in inspection probabilities, a random crackdown would still not use any information in the histories of firms. Therefore, the firm-wide moral hazard and correlated targeting mechanisms — which underlie the significant value of linked regulation that we find — would not be present.

We also provide model-free empirical evidence that suggests that actions are convex in the inspection probability. We do so in Table A-9. We produce the table using a similar procedure to the deterrence regressions in Section 3.3, but also include a quadratic term for inspections. The results show that, in the empirical range of inspection probabilities, violations are convex in the probability of inspection when we fit this quadratic term.[23]

Figure A-3: Convexity of equilibrium expected action, different functional forms

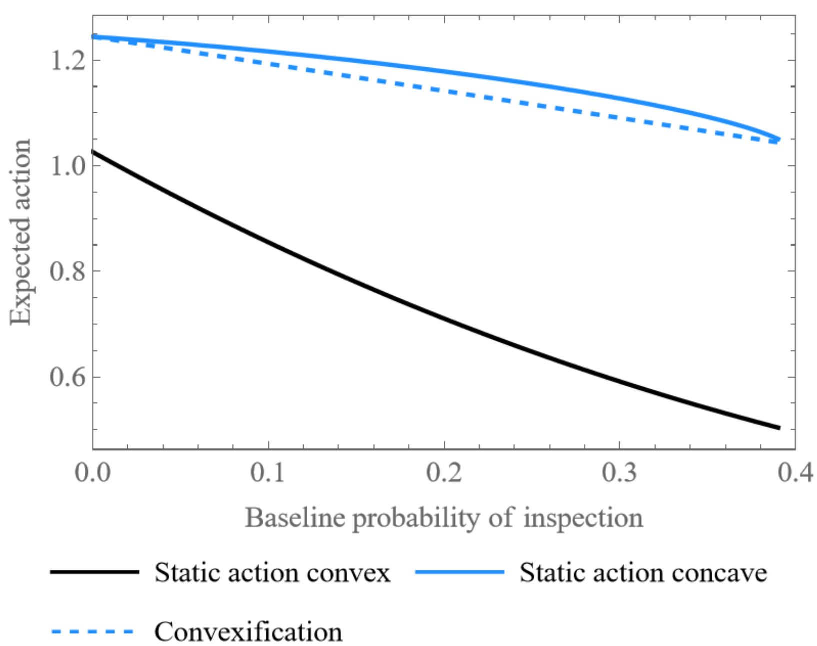

Notes: Expected action, $EA(\bar{z})$ versus z when $[\bar{z}(v_i = 1), \bar{z}(v_{-i} = 1), \bar{z}(v_i = 0), \bar{z}(v_{-i} = 0)]$ $= [z + 0.6, z + 0.4, z, z]$. Black line corresponds to functional form where static optimal action, $a_2(z)$, equals $(1 - z)^2$. Blue line corresponds to $a_2(z) = \sqrt{1 - z}$. Dotted line is the convexification of $EA(\bar{z})$.

[23]One caveat with this robustness check is that it only checks convexity "on average". Ideally we would check that there are no smaller regions of concavity in the optimal action. Data limitations would make such a test challenging.

References

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

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

Jan Eeckhout, Nicola Persico, and Petra E Todd. A theory of optimal random crackdowns. *American Economic Review*, 100(3):1104–1135, 2010.

Johannes Hörner and Nicolas S Lambert. Motivational ratings. *The Review of Economic Studies*, 88(4): 1892–1935, 2021.

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

David A Keiser and Joseph S Shapiro. Us water pollution regulation over the past half century: burning waters to crystal springs? *Journal of Economic Perspectives*, 33(4), 2019.

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