The Value of Regulatory Discretion: Estimates from Environmental Inspections in India*

Esther Duflo[†], Michael Greenstone[‡], Rohini Pande[§] and Nicholas Ryan[¶]

March 10, 2018

Abstract

High pollution persists in many developing countries despite strict environmental rules. We use a field experiment and a structural model to study how plant emission standards are enforced. In collaboration with an Indian environmental regulator, we experimentally doubled the rate of inspection for treatment plants and required that the extra inspections be assigned randomly. We find that treatment plants only slightly increased compliance. We hypothesize that this weak effect is due to poor targeting, since the random inspections in the treatment found fewer extreme violators than the regulator's own discretionary inspections. To unbundle the roles of extra inspections and the removal of discretion over what plants to target, we set out a model of environmental regulation where the regulator targets inspections, based on a signal of pollution, to maximize plant abatement. Using the experiment to identify key parameters of the model, we find that the regulator aggressively targets its discretionary inspections, to the degree that half of plants receive fewer than one inspection per year, while plants expected to be the dirtiest may receive ten. Counterfactual simulations show that discretion in targeting helps enforcement: inspections that the regulator assigns cause three times more abatement than would the same number of randomly-assigned inspections. Nonetheless, we find that the regulator's information on plant pollution is poor, and improvements in monitoring would reduce emissions.

I Introduction

Recent pollution levels in emerging economies like China and India exceed the highest levels ever recorded in rich countries. Such pollution reduces lifespans (Chen et al., 2013; Greenstone et al., 2014; Ebenstein et al., 2017) and labor productivity (Graff-Zivin and Neidell, 2012; Chang et al., 2016; Adhvaryu et al., 2016; He et al., 2016). High pollution persists despite strict emission

*We thank Sanjiv Tyagi, R. G. Shah, and Hardik Shah for advice and support over the course of this project. We thank Pankaj Verma, Eric Dodge, Vipin Awatramani, Logan Clark, Yuanjian Li, Sam Norris, Nick Hagerty, Susanna Berkouwer, Kevin Rowe and Jonathan Hawkins for excellent research assistance and numerous seminar participants for comments. We thank the Sustainability Science Program (SSP), the Harvard Environmental Economics Program, the Center for Energy and Environmental Policy Research (CEEPR), the International Initiative for Impact Evaluation (3ie), the International Growth Centre (IGC) and the National Science Foundation (NSF Award #1066006) for financial support. All views and errors are solely ours.

†MIT, eduflo@mit.edu

‡University of Chicago, mgreenst@uchicago.edu

§Harvard, rohini_pande@harvard.edu

¶Corresponding author. Yale, nicholas.ryan@yale.edu.

standards on the books. Regulatory enforcement is thus the crucible for environmental quality, but we know little about why enforcement fails. Regulators blame a lack of resources to carry out regulations. Other observers offer less charitable explanations, for example that regulators with wide discretion choose not to enforce standards, out of corruption, laziness or incompetence (Stigler, 1971; Leaver, 2009).[1] Whether regulatory discretion helps or hinders enforcement is generally uncertain. While discretion can be abused, it also allows regulators to use local information to strengthen enforcement.

Gujarat, India is an ideal setting in which to study regulatory enforcement. Gujarat is one of India's most industrialized states and contends with major pollution challenges.[2] Pollution is not for a lack of standards, as there are strict maximum limits on air and water emissions from industrial plants. Nor is it from an inability to punish violators; in 2008, before our study, the Gujarat Pollution Control Board (GPCB) ordered 9% of the plants in our sample to close, at least temporarily, sometimes cutting off their utilities.

While punishments are severe when meted out, the chance of being caught is low. The GPCB has a limited inspection budget and chooses which plants to inspect. Half of plants are inspected less often than the prescribed rate, while other, similar plants are inspected many times more. This discretion in inspection targeting may hurt or help regulatory enforcement. It would hurt enforcement if plants bribe the regulator to avoid inspections, or if regulators shirk and avoid the dirtiest plants to minimize conflict and monitoring costs. Discretion may also help, if it allows the regulator to use local information to target more polluting plants. Understanding the effects of limited resources and regulatory discretion on enforcement is generally hard due both to poor data on regulation and outcomes and the endogeneity of inspection targeting.

This paper uses a field experiment and structural estimation to unbundle the roles of resources and discretion in regulatory enforcement. The experiment covered 960 industrial plants and ran for two years. All sample plants came from the highest, 'red' category of pollution potential. The inspection treatment, assigned to half of plants, was cross-randomized with an audit reform

[1]The view that discretion leads to regulatory abuse of power was clearly expressed by the current Indian Prime Minister when he unveiled a scheme of randomly assigned inspections for compliance with labor rules: "Now computer draw will decide which inspector (labor) will go for inspection to which factory and he will have to upload his report online in 72 hours. These facilities are what I call minimum government, maximum governance. I have been hearing about 'inspector raj' since childhood." (The Economic Times, 2014)

[2]For example, all seven of the cities in Gujarat that are monitored for air pollution exceed the national standards for fine particulate matter (Central Pollution Control Board, 2012).

experiment in the subset of plants that were eligible for environmental audits (Duflo et al., 2013).[3] The inspection treatment met the *de jure* inspection rate by providing the resources needed to bring all treatment plants up to at least the required minimum number of inspections. The treatment also removed the regulator's discretion over these extra inspections, by allocating them randomly across all treatment plants. It did not alter pollution standards or the regulatory penalties for violations. The regulator continued to exercise discretion in allocating its existing budget of inspections across both the treatment and control groups.

The analysis is conducted with unusually rich, perhaps unique, data on the regulatory process, plant abatement and pollution. On the regulatory process, we code nearly 10,000 pieces of correspondence between the regulator and sample plants, which record their interactions over five years (from two years before the experiment through one year after). These documents include plant inspections, pollution readings, regulatory notices and penalties, on the regulator's side, as well as written responses from plants, such as documentation of abatement equipment. We also ran an independent endline survey of plant pollution and abatement costs.

Our experimental results pull on each link in the chain from inspections to emissions—which, in the end, the treatment did not meaningfully reduce. First, the experiment was implemented, in that inspection rates in the treatment group were twice those in the control, and treatment plants report higher perceived inspection rates, showing the scrutiny was felt. Second, treatment plants were more often found in violation of pollution standards and received more citations for those violations. Third, GPCB followed-up on inspections in both treatment arms the same way: all inspections were entered in the same database and were judged by the same officials. We empirically verify that, conditional on an inspection's findings, treatment status did not affect the regulator's follow-up. Yet, fourth, despite more citations, treatment plants were no more likely to be penalized. Fifth, we cannot reject the null of a zero treatment effect on average plant pollution emissions, although we find a small increase in the share of firms in compliance.[4]

Why did the bundled treatment, including both additional resources and reduced discretion, prove so weak at lowering emissions? A pattern of evidence suggests that a main reason is the

[3]In 1996, the High Court of Gujarat ordered GPCB to instate a third-party audit system wherein plants from polluting sectors must provide an annual audit report to GPCB. Duflo et al. (2013) evaluated a reform of this audit system and found that making third-party auditors more accountable to the regulator and less beholden to the plants they audit improves truth-telling and lowers pollution.

[4]The effects of both the inspection and audit treatments on pollution are negative, but their interaction is positive and significant, consistent with the information obtained from these channels being substitutable.

removal of regulatory discretion over what plants to inspect. Data on status quo inspections and the process of regulatory sanctions, after an inspection, shows that the regulator reserves the most costly penalties for extreme violations of regulatory standards. The treatment did identify many more plants that violated emissions standards, but did not find any more *extreme* violators, which would have been candidates for the most costly penalties. This gap suggests that the regulator's discretionary inspections, while done at a low rate on average in the status quo, nonetheless found many of the dirtiest plants. Adding random inspections mostly picked up smaller violators that the regulator would not have penalized in any case.

Motivated by this evidence, the second part of the paper uses a structural model to separate the roles of resources and regulatory discretion. We model environmental regulation under imperfect information and use the experimental variation in inspections to identify key parameters of the model. In the model, the regulator is benevolent and seeks to reduce pollution but is constrained by resources and information. The model includes two stages. In the *targeting* stage, the regulator choses whom to inspect, subject to its inspection budget, based on plant observables and noisy signals of plants' pollution, unobserved by the econometrician but known to the regulator. Plants decide whether to abate pollution given the threat of inspections. More polluting plants are both more likely to abate, because an inspection offers them a greater threat of penalties, and abate more conditional on taking action, because the abatement technology is proportional to pollution levels. The second stage is a *penalty stage* where, after an inspection, penalties may be levied on non-compliant plants. The regulator must follow an exogenous process for applying penalties to polluting plants. We estimate this process as a policy function, using our rich data on the regulatory process, and hold the estimated policy fixed in counterfactuals.

The plant's objective is to minimize the total cost of environmental regulation by trading off costly abatement actions against the risk of future inspections and the penalties they may beget. We use the penalty stage of the model to recover these unobserved penalties, using the plant's choice, in a dynamic problem, of when to install costly abatement equipment. This revealed preference approach has the benefit of capturing all the costs of regulation, including formal penalties, like a mandated plant closure, and informal costs like disruption and bribes.

The model estimates yield three broad sets of results. First, the regulator aggressively targets plants with high pollution signals, even though its signals are weakly related to true plant pollution. The estimates imply that half of plants receive fewer than one inspection per year, while plants expected to be the dirtiest may receive ten. This finding suggests that targeting, rather than neglect, is a reason why many plants are left largely alone. A sensitivity analysis, using the method of Andrews et al. (2017), confirms that the estimates of key model parameters are especially sensitive to the variation created by the experiment. If the experimental treatment effects on pollution had been greater, for example, we would have estimated the parameter that governs the efficacy of abatement in our model to be higher.

Second, regulatory penalties are costly when applied, but the risk of penalties is low. We estimate that a realized plant closure costs about $50,000, inclusive of formal and informal costs, or roughly two months of mean plant profits in our sample. Using these costs and the probability of penalties, the expected discounted value of an initial inspection to a plant is negative $2,000. Even for plants where an initial inspection reveals a pollution reading of at least five times the regulatory standard, the expected value of the inspection is negative $6,000—greater than for a plant with average pollution, but far smaller than the cost of certain punishment.

Third, counterfactual exercises reveal that regulatory discretion in choosing which plants to inspect is valuable, especially for tight inspection budgets. At the GPCB's current inspection rate, the inspections chosen by the regulator induce three times more abatement than would the same number of randomly-assigned inspections. The value of discretion declines as the number of inspections available to the regulator increases. The experiment doubled inspections, from the status quo, and assigned them randomly. We simulate that the effect on abatement of the same number of added inspections would have been fifteen percent greater than the estimated treatment effect, if the added inspections were assigned according to the regulator's discretion.

In a regime with discretion, improving the regulator's information can boost the efficacy of inspections. A technology that gave the regulator perfect information on plant emissions would increase abatement by thirty percent, at the status quo number of inspections. This abatement is the same as would be achieved by a one-third increase in the inspection budget, if the added inspections were allocated with discretion. Such technology is not science fiction: continuous emissions monitoring systems (CEMS) are used widely in the United States, and India has announced plans to roll out these devices in heavily polluting sectors (Central Pollution Control Board, 2013, 2014).

The paper makes several contributions to the literature. First, to the best of our knowledge, it provides the first experimental evidence on how inspections change plant emissions.[5] Second, we also believe it the first study with such rich data on the process of environmental regulation, including data on regulatory actions, penalties and independently-measured plant pollution emissions. Third, our model demonstrates how structural analysis can be used to unbundle the channels through which an experimental treatment works. Our model captures the context of the experiment, including unobserved plant heterogeneity, and we use the model estimates to simulate counterfactuals that either were not part of the experiment (e.g., expanding inspections with discretion) or could not plausibly have been part of any experiment (e.g., removing discretion from status quo inspections). We find that regulatory discretion, which is seldom measured despite great theoretical interest, can be valuable.[6] We also highlight that poor information is a major constraint on regulatory efficacy.

II Context and Experimental Design

II.A. Context: Regulation of Industrial Pollution in India

In India, national laws set pollution standards, and practically all enforcement of environmental regulations occurs at the state level. States may make their standards more strict than the national standards, but cannot relax them (Ministry of Environment and Forests, 1986). State Pollution Control Boards, such as the Gujarat Pollution Control Board (GPCB), are re-

[5]Studies of regulatory inspections in the United States show that inspections reduce pollution significantly (Hanna and Oliva, 2010; Magat and Viscusi, 1990). The studies rely on observational data wherein dirtier plants are more likely to get inspections, and this endogeneity is a strong concern (See Shimshack (2014) for a recent survey). Studies of regulatory efficacy in emerging economies are more mixed, with Tanaka (2013), for example, finding large reductions in pollution from a control policy in China and Greenstone and Hanna (2014) finding cuts in pollution in India from policies targeting air, but not water, pollution. On the costs of regulation, Greenstone et al. (2012) find U.S. regulations lower manufacturing productivity, and Ryan (2012), using a dynamic model, finds that the U.S. Clean Air Act Amendments raised entry costs in the cement industry. There are few studies of environmental regulation in developing countries, but Aghion et al. (2008) and Besley and Burgess (2004) document negative productivity effects of rigid industrial regulation for India.

[6]Environmental regulation is a classic setting for incentive regulation (Laffont and Tirole, 1993; Laffont, 1994; Boyer and Laffont, 1999). Limited regulatory capacity and commitment typically change the optimal regulatory policy in emerging economies with incomplete markets (Laffont, 2005; Estache and Wren-Lewis, 2009). Papers in organizational economics, such as Aghion and Tirole (1997), suggest that formal and real authority may then optimally diverge leading to substantial value for discretion. More broadly, our findings resonate with the literature on effective policy design when state capacity is limited (Besley and Persson, 2010). For instance, consistent with this paper's findings, Rasul and Rogger (2013) report significant gains from providing Nigerian bureaucrats autonomy in decision-making. Other papers note, to the contrary, bureaucrats and politicians may misuse discretion in environmental regulation (Burgess et al., 2012; Jia, 2014).

sponsible for enforcing the provisions of the Water Act (1974), Air Act (1981) and Environmental Protection (1986) Act, and their attendant command-and-control pollution regulations.

Turning to our study partner, the GPCB is responsible for monitoring and regulating approximately 20,000 industrial plants in the Indian state of Gujarat. The practices of the GPCB are largely common with other Indian states. Each GPCB regional office has several inspection teams and a Regional Officer who assigns inspections to plants. During an inspection, the team observes plant conditions and its environmental management and often, but not always, collects pollution samples for laboratory analysis. Officers at the regional and head office review inspection reports, which describe the plant's condition, and analysis reports, which list pollution concentrations for air and water pollutants.

Regulations mandate routine inspection of plants in sectors with the highest pollution potential ("red" category plants) every 90 days if they are large- or medium-scale and once per year if they are small-scale.[7] In the year before the experiment, 42 percent of control plants were inspected at less than the prescribed rate. These routine inspections, which the experiment manipulated, make up 35% of total inspections. The remainder are due to plant applications to operate (30%), public complaints (11%), and follow-ups on prior inspections or penalties (24%).

Plants found in violation of pollution standards can be harshly penalized. The regulator can mandate that a plant install abatement equipment, post a bond against future performance or even shut down, by ordering that a plant's water and electricity be cut. Utility disconnections remain in force until the plant has shown progress towards meeting environmental standards; the median duration of closure in our data is 24 days. Because abatement equipment is observable, plants may install equipment to show compliance, even when operational changes could fix an initial violation.[8] In addition to formal penalties like closure, plants may incur other costs of regulation, such as disruptions to plant operations during inspections or bribes.

[7]The GPCB follows a government classification for plants based on their reported scale of capital investment, with small-scale being investment less than INR 50m ($1 million), medium INR 50m to 100m ($1 million to 2 million) and large above INR 100m ($2 million) (throughout, we use an exchange rate of US$ 1 = INR 50). Prescribed inspection rates for these plants are comparable to those applied to large plants, by air pollution potential, in the United States (Hanna and Oliva, 2010).

[8]The regulator, in principle, can also take a violator to court for criminal sanction, but this is rare and does not occur in our data, because documenting violations is burdensome and there are long prosecutory delays.

II.B. Experimental Design

The goal of our experiment was to estimate the impact of moving from the status quo, infrequent inspections allocated with discretion, to regular inspections of all plants at prescribed inspection rates. Such a reform would bring the GPCB into compliance with its own prescribed inspection rates and the Central Pollution Control Board's (CPCB) inspection rules.

To this end, between August 2009 and May 2011 we worked with GPCB to increase inspection frequency for a random subset of highly polluting plants. We identified the population of 3,455 red-category (i.e., high pollution potential) small- and medium-scale plants in three regions of Gujarat (Ahmedabad, Surat and Valsad), which constitutes roughly 15% of the more than 20,000 regulated plants in Gujarat. By CPCB rules, these plants are supposed to be inspected either once per year if they are small-scale or once in three months if they are medium scale (Ministry of Environment and Forests, 1999). From this population, the sample of 960 plants was drawn in two batches. First, we selected all 473 audit-eligible (i.e., "super red") plants in Ahmedabad and Surat. Second, we randomly selected 488 plants from the remaining audit-ineligible population.

Inspection treatment assignment was randomized within region by audit-treatment-status strata (treatment: 233 plants, control: 240 plants and non-audit-eligible: 487 plants). The treatment was thus cross-randomized and implemented concurrently with the audit reform treatment studied by Duflo et al. (2013). The 481 plants assigned to the inspection treatment were assigned at least one annual initial (routine) inspection and up to four per year. In the first quarter the plant was assigned one initial inspection, after which it was randomly assigned on a quarterly basis to be inspected again with probability 0.66. After four quarters this cycle started over.[9]

Regional GPCB teams consisting of an environmental engineer and scientist conducted treatment inspections. To not overburden current staff, we worked with GPCB to rehire and integrate three recently retired GPCB scientists into the overall team. Rehired staff were sometimes allocated to regular inspections, and regular staff were often allocated to inspections assigned under the treatment, so that teams were well-mixed in practice.[10] Each morning in each region, the

[9]Towards the end of the inspection treatment, in the month prior to the endline survey, we also assigned, randomly and independently of the other treatments, some plants to receive a letter from GPCB reminding them of their obligations to meet emissions limits. This letter reiterated the terms of plants' environmental consent, which in principle they already knew, but it may have also served to increase the salience of regulatory compliance. The letter had no effect on emissions or compliance (Appendix Table A8).

[10]Administrative data on staff assignments across all three regions, shows that only one staff member, who was newly rehired, participated in treatment inspections only, whereas 32 staff members, mostly current employees,

designated inspection team was randomly assigned a list of plants from the treatment group at which to conduct initial "routine" inspections that day. This mimicked GPCB's practice of assigning teams to plants, except that the plant assignment was random, rather than being based on an official's discretion.

In all respects but targeting, control and treatment inspections were the same. Treatment and control inspection reports entered the same database without any distinguishing flag, had samples analyzed by the same GPCB labs and had the same GPCB officials deciding on follow-up inspections and punishment.

Two of our experimental design choices are worth discussing. First, our treatment simultaneously modifies the number of inspections and the method of assignment. Separate experiments on these two components would have been interesting: increasing the budget but with regulatory discretion, and, separately, asking the regulator to randomly inspect plants while holding the existing budget constant. We could not get the regulator's buy-in for the second option. Regarding the first option, we lacked the budget to do two different treatment arms, one with and one without discretion, and we concluded there was more to be learnt by testing the *de jure* policy, with a prescribed rate of inspections for all plants. Our joint treatment implies that we need the structural model to separate the impacts of resources and discretion.

Second, we explicitly asked the regulator to follow-up on the treatment and control initial inspections identically. If randomly assigning inspections at a prescribed rate was adopted permanently, then the regulator might change its follow-up behavior in response, which our experimental estimates will not capture. Identical follow-up allows us to focus on the one dimension, of inspection targeting, that did change. Moreover, had the regulator been free to vary how inspections were handled in the two groups, pure experimental or Hawthorne effects would have been a concern (e.g., the regulator trying to look tough, or ignoring the treatment inspections).

II.C. Data

The paper uses two sources of data, an endline plant survey and GPCB administrative records. The endline survey was conducted between April and July, 2011 (the experiment ended in May 2011), by independent agencies, mainly engineering departments of local universities, supervised

participated in both treatment and control inspections.

by the research organization J-PAL South Asia. The survey collected pollution readings, expenditures for abatement equipment investment and maintenance, and data on other aspects of plant operations. The GPCB issued letters that required plants to cooperate with the surveyors and stated truthfully that the results would not be used in regulation. Attrition was low and did not differ by treatment status (12.9% of plants closed during the study and only 4.7% attrited for other reasons, see Appendix Tables A5 and A6).

The second source of data is 9,624 GPCB documents on its interactions with plants. We categorize these documents by (a) whether they record an action of the regulator or a plant and (b) the type of action they record. Figure 3 shows the actions of the regulator and plant and Appendix Table A1 maps these actions to their documentation.

The regulator can choose to *Inspect*, *Warn*, *Punish* or *Accept*. To *Inspect* is to revisit the plant and gather another pollution reading. *Inspect* is documented by an inspection report and an analysis report giving lab results on pollution. To *Warn* is to threaten the plant that it is at risk of regulatory action and is documented by regulatory letters, citations for violations of pollution standards and closure warnings, that the plant will be closed, absent some remedial action. *Punish* records only costly punishments, mainly plant closure, documented by a closure direction sent to the plant or a notice to the utility to cut off the plant's water or electricity. *Accept* is to accept the plant is in compliance and is documented by the regulator revoking in writing a prior action, or simply taking no further action against a plant.

The plant has only two actions, *Comply* and *Ignore*. *Comply* is documented by the installation of abatement equipment, typically with a certification or invoice from the vendor that did the work. *Ignore* is documented by any letter the plant writes to the regulator that does not give evidence of compliance.[11] The action *Ignore* is also inferred from the absence of any plant response between regulatory actions.

Many regulator and plant actions occur in response to a prior action. We use two main rules to create chains of related actions. First, documents that explicitly cite one another are linked. Second, documents concerning the same plant that follow within a short time (usually 30 days)

[11] For example, a plant where GPCB found high air pollution readings claimed in correspondence: "At the time of visit our chilling plant accidentally failed to proper working, so chilling system of scrubber was not effective by simple water. Same time batch was under reaction and we were unable to stop our reaction at that time. Now it is working properly." That is, a piece of air pollution control equipment failed, causing pollution to be higher than normal during the visit. These types of explanations are common when plants are found well out of compliance.

are linked. We also impute additional *Ignore* actions, when none are documented, to enforce an alternating-move structure. Appendix A describes the linkage rules in more detail.

II.D. Randomization Balance Check

Plant characteristics and past regulatory interactions such as inspections, pollution readings and citations are balanced by treatment assignment. Appendix Table A4 presents a randomization check. Of 18 baseline measures reported, there is a significant difference between the treatment and control groups at the ten percent level on only one measure.

Many plants face costly penalties or take remedial actions despite the poor coverage of inspections. In the control group, 40 percent of plants had any pollution reading collected in the year prior to the experiment, and 34 percent of plants had a pollution reading above the limit (fully 85 percent of those with a reading taken). Many plants (22 percent in the control group) were cited for violations. More forcefully, 24 percent of control plants were mandated to install abatement equipment,[12] 7.5 percent were ordered temporarily closed, 2 percent had to post a bank guarantee (performance bond) and 1 percent had utilities cut off.

III Results: Experimental Estimates

This section examines how the experimental inspection treatment affects regulatory actions, plant abatement costs and pollution emissions. To motivate our structural analysis, we then document how the regulator targets discretionary inspections.

III.A. Regulatory Action

Table 1 presents differences in regulatory outcomes by treatment status during the experiment. Each row considers a different outcome. As we move down the table rows, we move along four links in the chain from inspections to penalties for violating plants.

First, the treatment was implemented faithfully (Panel A). Within a row columns 1 and 2 report the means for control and treatment plants, while column 3 reports the coefficient on the inspection treatment dummy from a regression of each outcome on treatment and dummies for

[12]This rate of equipment mandates is unusually high; an Air Action Plan issued a blanket mandate for all firms in some cities and sectors to upgrade their air pollution control devices (Gujarat Pollution Control Board, 2008).

strata used in randomization. Control plants were inspected an average of 1.40 times per year over the course of the experiment . Treatment plants were assigned to be inspected 2.12 more times per year and actually inspected an additional 1.71 times per year, more than doubling the annual rate of inspection, to 3.11 times. The treatment increased initial inspections, that start a new chain of interactions with the regulator, by 1.50 times per year.

Treatment-assigned inspections could, in principle, either crowd-out or crowd-in the regulator's discretionary inspections. Crowd-out would arise if the regulator diverts inspections away from treatment plants that are now being inspected at the prescribed rate. Crowd-in would occur if initial random inspections trigger follow-up inspections when a violation is found. On net, discretionary inspections were neither crowded-out nor in, perhaps because both effects cancel out.

Second, plants were aware of the increase in inspection frequency. Panel B reports perceived inspection frequency. Both control and treatment plants overstate how many inspections they receive in a given year. Though not officially told they would be inspected more, treatment plants recognized the change, and recalled being inspected a significant 0.71 times more than control plants in 2010. While correctly signed, the perceived difference understates by 58% the actual difference in inspection rates. A placebo check shows that there was no difference in perceived inspections in 2008, prior to the experiment.

Third, the additional treatment inspections led to more detected pollution violations and regulatory citations, which threaten action against plants. Panel C examines the number of regulatory actions against sample plants: the regulatory actions are ordered by increasing severity, from pollution readings, citations and warnings through to actions like mandated closures and utility disconnections that have a large cost to plants. Treatment plants are a significant 0.21 share more likely to have a pollution reading collected over the nearly two-year treatment, on a meager 0.38 base in the control. These readings lead directly to more treatment plants being found in violation of a standard (0.22 increase) and a greater number of citations (0.21 share per year) for these violations, more than doubling the citation rate in the control. Treatment plants see a statistically significant annual increase of 0.07 closure warnings, which formally threaten to close the plant unless remedial action is taken.

Fourth, despite the extra violations, there is no significant evidence of greater regulatory penalties for treatment plants. For example, closure directions, the mandated installation of equipment and utility disconnections are higher in the treatment but by small and statistically insignificant amounts (last two rows of Panel C). This fact will be central to our interpretation of the effect on compliance and abatement: despite increased inspections and violations, costly punishments did not increase.

Our experimental design was meant to rule out one potential explanation for the lack of additional punishments in the treatment—namely, that the regulator, despite regularly following-up on status quo inspections, just ignored the treatment inspections. We check the assumption of equal follow-up directly in Appendix Table A7, where we regress the probability that the regulator lets a plant go after an inspection on treatment status and the contents of that inspection. This probability does not differ by treatment status, without controls (column 1) or conditional on pollution (column 3). The other columns add interactions between the treatment and controls for pollution and other characteristics, and test for the joint significance of the interactions of treatment with these observables. We fail to reject that the treatment interactions are zero in all specifications. Thus the follow-up to an inspection is the same for treatment and control plants, conditional on the regulator's own information.

III.B. Plant Abatement Costs, Pollution and Compliance

Table 2, Panel A presents estimates of treatment effects on plant abatement costs. We use endline survey descriptions of abatement expenditures to separate abatement costs for capital and maintenance.[13] Abatement capital expenditures are observable by the regulator, and sometimes mandated in response to a high pollution reading. Maintenance expenditures are not directly observable by the regulator, but proxy for greater use of abatement equipment, and may thus be associated with lower pollution. More than half of control plants (0.57) install capital equipment, at an amortized cost of about \$2000 per year (column 1), while average maintenance costs are about \$264 per year with only 11 percent of plants reporting positive maintenance expenditures. As a basis of comparison, sample plants spend about \$145,000 on electricity annually. We find no meaningful effect on either capital abatement expenditures or whether any capital expenditure

[13] We distinguish maintenance from capital costs by searching descriptions of expenditures for strings associated with maintenance, like "maintain" or "change." See Appendix A for details. Capital costs are amortized into an annual flow of expenditures for comparison to maintenance costs.

was incurred (column 1 and 2). The column 3 estimate suggests that treatment plants did increase maintenance expenditure by \$838 (standard error \$499, p-value < 0.10). This effect is large relative to the control level of maintenance expenditures, but small in economic terms for plants of this size. Additionally, there is an insignificant treatment effect on the probability of reporting any maintenance expenditure (coefficient 0.01 share, standard error 0.02) (column 4).

Table 2, Panel B reports the results from regressions of pollution levels (column 1) and compliance (column 2) on treatment assignments for the inspection treatment, audit treatment and their interaction. Pollution is measured in standard deviations for each pollutant at the plant-by-pollutant level and standard errors are clustered at the plant level.[14] Compliance is an indicator for whether a plant-by-pollutant reading is below the pollutant-specific standard.

The treatment had modest impacts on pollution emissions. In column 1, the treatment reduced plant pollution emissions by 0.10 standard deviations (standard error 0.084). This effect is about half the size of the statistically significant -0.187 standard deviation reduction in pollution due to the audit treatment.[15] The treatment increased inspections by 1.71 per year; the implied local average treatment effect is therefore a reduction of 0.06 standard deviations of pollution per inspection. The audit-by-inspection interaction is large and positive. The sign of the interaction is as expected. Audits provide three reports of pollution in a year. If audit quality improves, then the informational value of extra inspections falls. Conversely, if plants are inspected regularly, then the audit adds less. The magnitude of the interaction is large enough to offset the sum of the main effects. However, we can reject neither that there is no effect of both interventions combined (p-value 0.97) nor that the joint effect for plants in both the inspection and audit treatment groups is equal to the inspection treatment main effect (p-value 0.19).

The column (2) entries indicate that the inspection treatment marginally increased compliance with pollution standards: treatment plants are 3.7 percentage points (standard error 2.1 percentage points, p-value $= 0.087$) more likely to comply, on a base of 61% compliant pollution readings in the control. (Multiple pollutants are observed in the survey and only 10% of plants are compliant on *all* pollution readings measured).[16]

[14] All specifications include region fixed effects and an audit-eligibility indicator. Since only Ahmedabad includes both audit-eligible and ineligible plants, this specification is equivalent to using region-by-eligibility fixed effects.

[15] The audit treatment effect on pollution reported in Duflo et al. (2013) was estimated in the inspection control group only and was slightly larger.

[16] To test the robustness of this compliance effect, Appendix Table A9 reports placebo checks where compliance is coded to occur at various multiples of the real standard. The effect of inspection treatment on compliance is

Compliance can increase without a large reduction in average pollution if plants near the standard were the most responsive to the inspection treatment. Figure 1, Panel A plots the coefficients on inspection treatment from regressions of indicators for a pollutant reading being in a given bin, relative to the regulatory standard, on treatment assignments (as in Table 2, Panel B, column 3, but with finer bins rather than a single dummy for compliance). Treatment reduces pollution readings just above the standard more than in any other bin, though this decrease is not statistically significant (p-value 0.17), and it significantly increases the number of readings just below the standard, in [-0.2,0.0]. The treatment thus shifted some plants that were modestly out of compliance with the *de jure* standard into compliance.

III.C. Status Quo Targeting of Inspections

The experimental results are puzzlingly weak. A doubling of inspections and citations failed to increase penalties or reduce average emissions and led only to small changes in abatement costs and compliance, despite that the regulator does punish plants found in violation and does so similarly in the treatment and the control groups. Did the marginal treatment inspections not generate much abatement because they are random, while discretionary inspections are targeted? We provide several pieces of evidence on targeting in the status quo.

First, Figure 7, Panels A and B demonstrate that the treatment did not appreciably increase the number of plants subject to five or more inspections in a year, who are typically severe violators. Instead, it increased the frequency of inspections for plants that would not have been inspected regularly, reducing the share of plants inspected at less than the prescribed rate from 50 percent in the control to 13 percent in the treatment group.

Second, despite the additional inspections, Figure 2 reveals that the treatment did not increase the number of extreme violators, plants with pollution readings 5 or 10 times the standard.[17] The treatment *did* find many plants that exceed the standard by smaller amounts. This lower intensity of marginally discovered violations suggests that the regulator is already inspecting the dirtiest plants, using the few inspections available in the control group.

Third, our endline survey pollution readings, in the control group, predict future regulatory

statistically significant only at the true standard.

[17] Thirty five (10) plants in the treatment group have a pollution reading greater than $5\overline{p}$ ($10\overline{p}$), compared to 33 (12) in the control group. These rates are practically identical, and the null hypotheses that detection probabilities for plants with readings $> 5\overline{p}$ and $> 10\overline{p}$ do not differ by treatment status cannot be rejected.

inspections conditional on plant observables and the regulator's own past readings (Appendix Table A10). Since the regulator did not see the endline survey readings, this prediction must mean the regulator has its own signals of plant pollution, and uses these signals to target inspections.

Thus, it appears that the regulator is selectively inspecting and punishing the most polluting plants. The failure of the treatment to use the regulator's private information may, in turn, explain the surprising finding that the treatment had only weak impacts on penalties and emissions. The following sections build on this insight to specify and estimate a model of the regulator's targeting problem.

IV A Model of Inspection Targeting and Enforcement

To understand why the treatment did not meaningfully reduce plant emissions, we set out a structural model of regulation and plant behavior to unbundle the roles of resources and regulatory discretion. We consider a benevolent regulator, who seeks to maximize abatement, given available information, resource constraints and the process of applying penalties. Thus, we abstract away from the possibility that the regulator's choice of inspections is corrupted and ask whether high plant pollution can be explained in terms of the constraints on the regulator's actions and information. (Our specification will allow corruption in the conduct of inspections, just not their assignment.)

We model regulator-plant interactions as a game in two stages.

1. *Targeting stage.*
    1. The regulator chooses an inspection targeting rule to minimize plant pollution subject to a budget of inspections;
    2. Plants choose whether to run their abatement equipment, given their abatement cost, known level of pollution, and the regulator's targeting and penalty rules;
    3. The regulator observes a part of plant pollution and inspects plants by applying the targeting rule (1.) to this signal, yielding a pollution reading from the inspection.
2. *Penalty stage.*
    1. The regulator acts as a *regulatory machine*, following exogenous rules for follow-up and punishment based on pollution measured in inspections and plant actions;
    2. Plants face a single-agent dynamic problem: they play against the regulatory machine, and

decide when to comply versus risk future penalties.

The model thus encompasses both the targeting of inspections, which our experiment changed, and the penalties from high pollution readings, which it did not. To accord with the experiment, we simplify the regulator's behavior in the penalty stage by estimating a *regulatory machine* policy function that maps states to action probabilities. The estimates and targeting counterfactuals therefore take the penalty stage policy as given.

IV.A. Targeting Stage

IV.A.1. Targeting stage actions

Plant j has a latent level of pollution in period m of

$$\log \tilde{P}_{jm} = \phi_0 + \phi_1 X_j + u_{1j} + u_{2jm}, \tag{1}$$

where X_j are observable plant characteristics, u_{1j} is a pollution shock known to both the plant and the regulator, and u_{2jm} is a pollution shock known only to the plant, which varies over time. We assume both pollution shocks are normal with $u_{1j} \sim \mathcal{N}(0, \sigma_1^2)$ and $u_{2jm} \sim \mathcal{N}(0, \sigma_2^2)$. The higher the share of the residual variance in pollution that is due to σ_1, the better the information of the regulator. At the extreme, if $\sigma_2 = 0$, the regulator has perfect information and observes pollution at each plant; this would be the case if the regulator had access to a perfectly functioning monitoring technology.

The regulator sets a targeting rule $\mathcal{I}(u_{1j}|X_j, T_j, \theta_T)$ that assigns an annual number of initial, routine inspections as a function of pollution shock u_{1j}, given plant characteristics, treatment status and targeting parameters θ_T. The regulator sets the rule first and then observes u_{1j} to assign inspections.

Plants know $\mathcal{I}(\cdot|\cdot)$, their characteristics, treatment status and pollution shocks and can therefore calculate how often they will be inspected. Plants also know their cost of abatement operations and maintenance c_j, where $\log c_j \sim \mathcal{N}(\mu_c, \sigma_c^2)$. The cost and pollution shocks are mutually independent, $c_j \perp u_{1j} \perp u_{2jm} \perp u_{2j,m+1}$. Plants use this information to decide whether to *Run* their existing abatement equipment, which action is not observed by the regulator.[18] Running abatement equipment reduces pollution proportionally to its latent level, such that

[18] Plants must install pollution control devices, depending on their sector and emissions potential, as a condition of opening.

$\log P_{jm} = \log \tilde{P}_{jm} + \phi_2 Run$, where $\phi_2 < 0$. The functional form assumption that abatement is proportional to pollution provides the regulator one incentive to target highly polluting plants. We cannot directly test this assumption, although it seems to be realistic for many production processes: for example, air pollution control equipment removes a fraction of pollution emissions that are sent up a plant's chimney.[19]

IV.A.2. Targeting stage payoffs and equilibrium

An equilibrium in the targeting stage consists of an abatement rule for the plant that minimizes the cost of regulation and a targeting function for the regulator that minimizes pollution, given the signal of pollution it observes.

The cost of regulation for plants in the targeting stage is summarized by a penalty value function $V_0(P_{jm})$, which gives the money value to the plant of an initial inspection (hence subscript 0) that finds pollution reading P_{jm}. We derive this function in Section V.A., below, as the expected discounted value to the plant of all regulatory actions in the penalty stage, including follow-up inspections, penalties and possibly bribes.

A plant anticipating I_j initial inspections will run its equipment if the reduction in expected penalties, from lower pollution at each initial inspection, exceeds its cost of maintenance

$$Run^* = \mathbf{1}\left\{ I_j(V_0(P_{jm}) - V_0(\tilde{P}_{jm})) > c_j \right\}. \tag{2}$$

We expect that the value $V_0(\cdot)$ will be decreasing in pollution, becoming more negative, so that, for a plant that *Runs* its equipment, $\phi_2 < 0 \Rightarrow P_{jm} < \tilde{P}_{jm} \Rightarrow V_0(P_{jm}) - V_0(\tilde{P}_{jm}) > 0$. That is, for a sufficiently small cost of maintenance, running abatement equipment will be worthwhile, since it will reduce expected penalties in the penalty stage that follows an initial inspection.

The objective of the regulator is to set an inspection rule that maximizes total abatement (i.e., minimizes total pollution). Targeting depends on endogenous parameters $\lambda \in \theta_T$ and additional

[19] Air pollution control devices like filters, electrostatic precipitators, cyclones and scrubbers are commonly installed in industrial plants in both India and developed countries. The US EPA rates such equipment by the fraction of a pollutant it removes and reports efficacies of 90% for cyclones, 95-99% for bag filters and 99% for scrubbers under their intended operating conditions (Environmental Protection Agency, 2012). As part of another project we physically measured the efficacy of air pollution control devices for a small number of plants in Surat, Gujarat, one of the areas in this paper's sample, by comparing pollution concentrations before and after control devices within the same plant's exhaust system. We found efficacies of 76% for cyclones and bag filters, somewhat worse than the EPA ideal.

exogenous parameters $\beta, \rho \in \theta_T$. The optimal targeting parameter vector λ^* solves

$$\lambda^* \in \arg\max_{\lambda} \sum_{j=1,\ldots,N} \int\int \mathcal{F}\Big(\mathcal{I}(u_{1j}|X_j, T_j, \lambda, \beta, \rho)(V_0(P_{jm}) - V_0(\tilde{P}_{jm}))\Big) \times \tilde{P}_{jm}(1 - e^{\phi_2}) dF(U_2) dF(U_1) \quad (3)$$

such that

$$\sum_{j=1,\ldots,N} \int \mathcal{I}(u_{1j}|X_j, T_j, \lambda, \beta, \rho) dF(U_1) = N \cdot \bar{I}. \quad (4)$$

The integrand of the objective (3) is the product of the probability of plant abatement and the quantity of abatement a plant achieves by choosing to *Run* its equipment. This plant-level expected abatement is integrated over the distributions of the two parts of pollution, which the regulator observes after setting the targeting rule (u_{1j}) or does not observe (u_{2jm}), and summed over plants $j = 1, \ldots, N$ to yield total abatement.[20]

The regulator's budget constraint (4) is that total inspections under the chosen targeting rule must be equal to the total inspection budget in expectation (i.e., the product of the number of plants and the average inspection rate, $\bar{I}$, per plant). While the targeting rule depends on a stochastic shock, we treat the budget constraint as exactly binding, since the regulator sets the rule before observing the u_{1j}, the observed pollution shocks are independent and there are a large number of plants N. The regulator can therefore work out how many inspections a given rule will yield in expectation and this expectation will be very nearly right.

For our estimation and counterfactuals we impose a probit link form for the targeting rule:

$$\mathcal{I}(u_{1j}|X_j, T_j, \lambda, \beta, \rho) = \lambda_2 \Phi\left(\frac{\lambda_1 + X_j'\beta_1 + T_j'\beta_2 + u_1}{\rho}\right), \quad (5)$$

where Φ is the normal cumulative distribution function. The parameters λ_1, λ_2, β and ρ determine the shape of the targeting rule: λ_2 sets the maximum number of inspections, λ_1 shifts the share of plants that will have inspections near, or far below, the maximum, β is the coefficient vector on plant observables, and ρ scales the argument of the targeting function. Loosely, a high λ_2 and a very negative λ_1 will concentrate inspections aggressively in the plants observed to be dirtiest. This functional form is restrictive: unconstrained, the regulator may not have chosen

[20] This objective function does not ascribe value to the penalty phase for the regulator. In particular, it does not account for the fact that, by targeting a more polluting plant, the regulator, in the penalty stage, could better compel them to install abatement equipment, providing a direct benefit of lower future pollution. We neglect this outcome in the targeting stage because (a) most plants have unused abatement equipment, so the installation of more equipment, on its own, is unlikely to reduce pollution (b) the cost of maintenance is far below the cost of new equipment, so the maintenance margin is a more likely channel for plant deterrence. We believe that mandated equipment installation is mainly a way to punish plants and of low marginal environmental value.

from the probit family. We specify this form for two reasons. First, because it parsimoniously fits a range of interesting targeting rules (See Appendix B for Monte Carlo simulations). Second, it greatly reduces the dimensionality of estimation, relative to a non-parametric targeting rule, and thereby makes it possible to constrain estimates by imposing the optimality of targeting.

IV.B. Penalty Stage

The targeting stage takes as given the value function $V_0(P_{jm})$ of an initial inspection conditional on pollution. To estimate this value function in the penalty stage, we model a plant's optimal compliance behavior *after* an initial inspection as a dynamic discrete choice problem, assuming that the plant's objective remains to minimize the overall cost of regulation.

The penalty stage starts with round 1, when an initial inspection takes place, and in subsequent rounds $t = 2, 3, \ldots$ the plant j and the regulatory machine R alternate moves. In all even rounds, the plant may *Comply* or $Ignore$ the regulatory machine, where *Comply* requires a plant to pay a constant amount to install abatement equipment. In any odd round after the first, the regulatory machine has four actions a_{Rt}: *Inspect*, $Warn$, *Punish* or *Accept*, which correspond to categorized regulatory data (Section II.C.).

Figure 3 shows these actions and their within-round payoffs for the plant. The plant's payoff for inspection includes any disruptions and bribes paid during the inspection. The payoff for punishment is the cost associated with temporary closure and any remediation. Thus, the plant seeks to minimize regulatory costs by choosing between a known abatement cost and the value of continuing the stage, possibly facing greater costs if the regulatory machine chooses to $Inspect$ or $Punish$. Each chain of interactions between the plant and the regulatory machine is treated as independent.[21] We assume that the plant knows the regulatory machine's action probabilities in each possible future state.

[21]Specifically, we assume that $u_{2j,m+1}$ is independent of u_{1j} and u_{2jm}. The regulator observes u_{1j}, but conditional on this, does not, for example, use past penalty stage readings to determine targeting. The data broadly supports this assumption: the average time between chains, about five months, is much larger than the average time between actions within a chain, two weeks. Further, recent pollution readings do not change regulatory targeting of inspections (Appendix Table A10, column 4). Lastly, the regulator has a short memory; 93% of the time when an action cites a prior inspection it is the most recent prior inspection.

IV.C. Simulations of Optimal Inspection Targeting

How much should the regulator concentrate inspections among the plants with high observed pollution shocks? Since the plant's value of the penalty stage decreases (i.e., becomes more negative) with pollution, the regulator can induce more abatement by allocating inspections to plants with high pollution and therefore high expected penalties. This argument favors a steep targeting function that concentrates inspections on heavily polluting plants. Moreover, plant reductions in pollution are proportional to the pollution level, so allocating inspections to higher-polluting plants yields higher abatement when those plants do abate. In favor of a flatter targeting function, however, abatement also depends on the cost of running the equipment. If the regulator targets all inspections to a few plants that it expects are highly polluting, it may miss some easy targets with low running costs.

Appendix B reports Monte Carlo simulations that illustrate how changes in the shape of the penalty function and regulatory information affect the choice of targeting rule, for one parameterization of the model. The simulations show that the results of this trade-off vary with the regulator's information (See Appendix Figure B2 for greater detail). If the regulator is poorly informed, it is better to concentrate inspections in the plants with the highest observable pollution shocks. If the pollution signal is imprecise, then the regulator targets large outliers, that it is confident will be polluted enough to abate when fearing inspection. If the regulator observes a larger fraction of the variance in pollution, the optimal targeting function is flatter. In this case the regulator is confident that plants observed to be moderately polluting may also abate if inspected and so spreads inspections around, to catch plants with not only high pollution but also low abatement costs.

V Estimation

The model is estimated moving backwards. First, we use backward induction within the penalty stage. Our penalty estimation pools treatment and control plants since, as we discussed, the regulator applies the same penalty rules for all plants. Next, we use the estimated value function from the penalty stage to obtain targeting stage parameters.

V.A. Maximum Likelihood for Penalty Stage

Building the likelihood for plant actions requires several preliminary steps. (1) We specify the common state of the game as comprised of the pollution reading, the last two actions of the regulator and plant and the game round. (2) We estimate state transition probabilities using a count estimator. (3) We estimate a multinomial logit model of action probabilities for the regulatory machine, conditional on the state. These steps are described in detail in Appendix IV.B. Here we focus on the specification of penalties and the value of regulation, taking the states, state transitions and regulatory policy as given.

The plant payoff if it complies by installing abatement equipment is $-k$. We assume all plants have a cost for installing abatement capital equal to the average value of abatement capital costs observed in our sample, $k =$ \$17,000.

The plant payoff, if the regulator chooses $Punish$, takes one of two specifications: a constant, $h(p_{jt}) = -\tau_0$, or a function of pollution

$$h(p_{jt}) = -\tau_1 \mathbf{1}\{\overline{p} < p_{jt} < 2\overline{p}\} - \tau_2 \mathbf{1}\{2\overline{p} \leq p_{jt} < 5\overline{p}\} - \tau_3 \mathbf{1}\{5\overline{p} \leq p_{jt}\},$$

where $\overline{p}$ is the legally mandated pollution threshold. This functional form allows the regulatory machine to punish high polluters with a higher probability and possibly different penalties.

In some specifications plants also have direct costs of inspections $b(p_{jt}, a_{j-}) = (1 - \mathbf{1}\{a_{j-} = Comply\}) \times (\nu_1 \mathbf{1}\{\overline{p} < p_{jt} < 2\overline{p}\} + \nu_2 \mathbf{1}\{2\overline{p} \leq p_{jt} < 5\overline{p}\} + \nu_3 \mathbf{1}\{5\overline{p} \leq p_{jt}\})$. This function specifies that inspections are costless for plants that have recently complied but, for plants that have not, inspections have a cost that depends on pollution emissions. The form is meant to capture the idea that recent compliance may excuse the plant from offering bribes or other disruptions.

Using these preliminaries we build the plant's action probabilities. The choice-specific utility of taking action a_{jt} for within-round payoff $\pi_j(a_{jt}|s_t)$ is

$$\begin{aligned} v_j(a_{jt}|s_t) &= \pi_j(a_{jt}|s_t) + e_j(a_{jt}|s_t) + \delta \sum_{s_{t+1}} f(s_{t+1}|a_{jt}, s_t) \sum_{a_{R,t+1}} Pr(a_{R,t+1}|s_{t+1}) \times \\ & \left\{ \pi_j(a_{R,t+1}|s_{t+1}) + \delta \sum_{s_{t+2}} f(s_{t+2}|a_{R,t+1}, s_{t+1}) V(s_{t+2}) \right\} \end{aligned} \tag{6}$$

We specify shocks $e_j(a_{jt}|s_t)$ to the utility of each action that are distributed identically and independently across actions with a type-I extreme value distribution of unknown variance, generating closed-form solutions for action probabilities (Rust, 1987). The plant discounts the value of future rounds by δ. The transition $f(s_{t+2}|s_{t+1})$, from the plant's point of view, contains both the machine's action and any other change in the state before the plant moves again. We assume that the machine's action probabilities $Pr(a_{R,t+1}|s_{t+1})$ and state transition probabilities $f(\cdot|\cdot)$ are stable and known to the plant.

The plant's optimal action in the penalty stage maximizes its expected discounted value at each state. The value of the state is the value of this best action $V_j(s_t) = \max_{a \in A_P} v_j(a_{jt}|s_t)$. In determining its move now, the plant takes into account current payoffs and the value of future states that are likely to follow. We use backward induction to solve for the values of each state for the plant, conditional on a given set of penalty- and inspection-cost vectors $\theta_P = \{\tau, \nu\}$.

Identifying the model parameters requires two known payoffs and a discount factor (Rust, 1994; Magnac and Thesmar, 2002). For the first payoff, we assume a zero payoff from *Ignore* for the plant. For the second, we assume the penalty function equals zero for states when plants' pollution reading is absent or below the standard. Given these two assumptions, the variance of the plant action shock σ_a is then a free, estimable parameter. We use a discount factor of $\delta = 0.991$ between rounds, which has been calibrated, given the average round duration, to match the annual returns on capital for Indian firms found by Banerjee and Duflo (2014).

The likelihood over chains n and rounds t is

$$\mathcal{L}(\theta_P) = \prod_n \prod_{t=1}^{t=T_{jn}} Pr(a_{jnt}|s_{jnt}, \theta_P).$$

We use a gradient-based search with numerical derivatives to find parameters that maximize the probability of plant actions that are observed in the data. Given the estimated parameters $\hat{\theta}_P$, we use backwards induction to calculate the value of the penalty stage $V_0(\cdot)$ for each level of pollution at the time of an initial inspection.

V.B. Generalized Method of Moments for Targeting Stage

In the targeting stage, the regulator sets a rule for how to inspect plants. Plants, anticipating the value of pollution that each inspection will yield and associated penalties, decide whether to *Run* their abatement equipment. The *Run* decision is endogenous to plant pollution shocks u_{1j} and u_{2jm}, both not observed by the econometrician. Taken together, the targeting stage is characterized by a system of equations for inspections, pollution and the *Run* decision. We use the generalized method of moments for estimation with both analytic and simulated moments.

V.B.1. Targeting Stage Estimation Moments

The parameters to be estimated are $\theta_T = \{\phi, \beta, \lambda_1, \lambda_2, \mu_c, \sigma_1, \sigma_2\}$, where ϕ are the parameters of the pollution equation (1), β and λ govern inspection targeting (5), μ_c is the mean of the log abatement maintenance cost, and σ_1, σ_2 give the standard deviations of pollution shocks, which are known to both the plant and regulator (u_1) or the plant only (u_2), respectively.

We additionally fix the values of two model parameters outside of the estimation: the variance of the maintenance cost shock σ_c and an inspection targeting parameter ρ. While in principle identified, we found that estimating these parameters along with θ_T in our sample yielded estimates too imprecise to be usable. Below, we discuss why this is the case and how our estimates vary over a range of assumed values for these two parameters.

We observe N_j, X_j, T_j, P_j and $c_j \times Run$ in the data and estimate $\widehat{V}_0(p_j)$ from the penalty stage, as described above. The estimation moments are chosen to match features of the pollution and inspection distributions, in particular the interactions of treatment with inspections and residual pollution. Appendix C.2 derives the moment conditions and a sensitivity analysis in Section VI.B.2. discusses the contribution of different moments to identification.

A first set of moments is based on the error in the pollution equation, which is orthogonal to treatment assignment in the model. Letting $Z_j = [\mathbf{1}\ X_j\ T_j]$ where T_j is the treatment assignment,

$$g_1(\phi) = Z_j'(\log P_j - \phi_0 - \phi_1 X_j - \phi_2 Run),$$

A second set of moments is based on expected inspections and inspections squared

$$\begin{aligned} g_2(\lambda, \beta) &= \mathbf{1}'(\mathbb{E}[\mathcal{I}(u_{1j}|X_j, T_j, \lambda, \beta, \rho)] - I_j) \\ g_3(\lambda, \beta) &= \mathbf{1}'(\mathbb{E}[\mathcal{I}^2(u_{1j}|X_j, T_j, \lambda, \beta, \rho)] - I_j^2). \end{aligned}$$

where the expectation is calculated analytically, in the model, based on the targeting function (5) and the distribution of u_1 shocks. Expected squared inspections are meant to capture regulatory information because dispersion in inspections, conditional on observables, reflects targeting on unobserved (to the econometrician) pollution shocks.

A third set of moments is based on the probability of running abatement equipment and the mean cost conditional on running. These moments are intended to target μ_c, the mean of the unconditional maintenance cost distribution, and ϕ_2, the efficacy of abatement.

Fourth and last, we form moments based on the variance of pollution shocks and their covariance with inspections. In the model, if the regulator observes a higher fraction of pollution variance, then inspections will have a higher covariance with residual pollution.

We fix σ_c and ρ outside the estimation. For σ_c, higher-order moments of the truncated cost distribution could in principle provide identification. In practice, these estimates are imprecise and sensitive to the choice of higher-order moments. With only around ten percent of plants choosing *Run*, it is difficult to use the observed, truncated costs to infer the shape of the unconditional maintenance cost distribution. Therefore, we set $\sigma_c = 0.5$, as this is in roughly the mid-point of the estimates we obtained by using different higher-order moments (albeit with large standard errors). We set the parameter $\rho = 0.25$ standard deviations of observed pollution. This parameter operates almost like a scaling factor in the targeting function argument.[22] Changes in the freely estimated targeting parameters, β and λ, can therefore closely replicate the effects of varying ρ on the inspection distribution in the model (See Appendix D). Section VI.B.2. considers the robustness of the targeting stage estimates to these assumptions.

V.B.2. Imposing the Constraint of Optimal Targeting

Optimal targeting is defined by maximizing abatement (3) subject to the inspection budget constraint (4). To impose optimality in the regulator's choice of targeting parameters λ, we require that the first-order conditions of the Lagrangian of the regulator's problem hold (see Appendix C.3 for the derivation). Under the assumed targeting functional form, the three first-order conditions impose two independent non-linear constraints. These conditions state that the marginal reduction in pollution from increasing either targeting parameter must equal the contribution of that parameter to the inspection budget. In addition, the parameters of an optimal rule must satisfy the inspection budget.

V.B.3. Targeting Stage Objective Function

We stack the moments to form $g(\theta_T) = [g_1' \; g_2' \ldots g_7']$ and minimize gWg' as a function of θ_T to estimate the parameter vector $\hat{\theta}_T$. In constrained estimation, we conduct this minimization subject to the optimal targeting constraints. We update the weighting matrix W to form a two-step

[22] The parameter ρ is not purely a scaling factor because u_1 appears outside the targeting argument, with known units of pollution, though it is not observed by the econometrician. However, estimation runs with free ρ did not reliably converge.

optimal estimator. Standard errors are calculated to account for the non-linear constraints and their correlation with the moments (Newey and McFadden, 1994), and adjusting for simulation bias, which is negligible with $S = 5,000$.

VI Structural Estimates of Regulatory Costs and Targeting

VI.A. Penalty stage

VI.A.1. Plant and regulatory choice probabilities in the penalty stage

Table 4 presents estimates of multinomial logit coefficients for the conditional action probabilities, for both the regulatory machine and plants. The machine is much more likely to punish when pollution is high (Columns 1 to 3). Past actions matter for current actions. The machine is less likely to *Warn* or *Punish* if it has warned before. It is also less likely to *Punish* if the plant has *Complied* before. Plant compliance drops the likelihood of any inside action, which would continue the penalty stage, and raises the probability that the machine *Accepts* to end the stage and the imminent threat of punishment. These estimates support the trade-off between compliance and future penalties for plants with high pollution that underlies the model.

VI.A.2. Revealed Preference Penalty Estimates

Table 5 presents the dynamic estimates for the penalty function. These estimates put a monetary value on mandated plant closings, utility disconnections, and other penalties that would not be estimable without the structural model.

In columns (1) and (2), we assume that *Inspect* does not entail any costs for the plant so that *Punish* is the only regulatory action that is costly. In column (1) we see that when punishment cost is constant, it costs a plant \$54,000 (standard error \$25,000; we round off penalty estimates). In column (2) we allow the cost to vary with pollution. We cannot reject that the penalty function is flat with respect to pollution (above the threshold), with estimates ranging from \$40,000, when observed pollution is slightly above (1-2$\overline{p}$) the standard, to \$54,000 (2-5$\overline{p}$) for higher levels.

Columns (3) and (4) consider the case where both *Punish* and *Inspect* are assumed to be costly to plants. We call this the case "with bribes" for short, though inspections may impose other costs like disruptions to plant operations. Relative to column (1), the estimated cost of punishment in column (3) declines to \$28,000 (standard error \$21,000) with a per inspection cost of \$10,000 (standard error \$3,000). Inspections are less costly than punishments, but more frequent; reflecting this, the estimated cost of inspection is lower than that of punishment and a lower cost of punishment is needed to rationalize plant compliance behavior when inspections are also costly. In column (4), we allow the cost of inspections to vary by pollution reading. We find, again, that inspections cost plants perhaps one-third or less of the value of punishment, and that the cost of inspections does not significantly vary with pollution.

Does the scale of estimated penalties and inspection costs make sense? We compare the estimates to abatement costs and plant profits as benchmarks. The highest penalties estimated for punishment are over three times the average equipment capital cost. This ratio of penalties to costs is reasonable given that penalties must meet, or exceed, costs required to induce abatement, and that penalties occur infrequently, even for violating plants.[23] On profits, mean plant annual sales in our endline survey are \$2.9 million. The typical penalty for severe pollution is plant closure, with a median duration of 24.5 days. For each 10 percentage point profit margin for plants, and assuming profit is proportional to closure (i.e., no substitution across periods), this duration of closure implies a loss in profits of \$20,000. Variable profit margins would then have to be in the range of 20-30% to match the column (1) model estimates, which is arguably on the higher side, but not unreasonable.

VI.A.3. The Value of Environmental Regulation to Plants

The value of the penalty stage to a plant summarizes all costs of environmental regulation. Figure 4 shows this value, at a variety of states, as calculated through backward induction given the estimated costs of regulatory penalties from Table 5, column (2). At each state, values are divided between expected discounted future abatement costs (light grey) and expected discounted future regulatory penalties (dark grey). The figure shows three different dimensions of the state: the time dimension is shown across panels, the pollution dimension is shown across clusters of bars within a panel, and the dimension of regulatory action across bars within a cluster.

States late in the penalty stage, when the machine is more likely to punish the plant, have sharply lower valuations for plants. Panel A shows the plant value at $t = 6$. The value to the

[23] A plant with an extremely high pollution reading has a 1/3 chance of punishment implying an expected value of penalties ($= 1/3 \times \$54,000$) that is about equal to the average abatement capital cost (\$17,000).

plant is sharply decreasing in pollution, for readings above the standard, reflecting the higher risk of punishment, costly plant compliance and continuation of the penalty stage associated with high pollution. The share of value due to penalties is also increasing in pollution.

Panel B shows the expected discounted value for the plant when it can first act ($t = 2$). The value is shown for only the machine's lagged action *Inspect*, which, by construction, is the only action that the machine can take in $t = 1$. The values are much less negative than in $t = 6$, since the probabilities of punishment, compliance and continuation are all sharply lower in the early going. There remains a steep gradient of penalties in pollution: the value ranges from negative $1,160, if the machine did not take a pollution reading, down to negative $6,240 if the inspection found a pollution reading more than five times the standard, a more than five-fold difference. The share of the expected value due to penalties is also increasing.

Overall, using the distribution of pollution on first inspection, the expected discounted value of regulation on first inspection is -$2,131, of which 40% is expected future penalties and 60% expected future abatement capital expenditures. Thus, a measure of regulatory costs that does not account for the monetary value of penalties would be greatly understated, and differentially understated for more polluting plants.

Using these expected discounted values, at discrete levels of pollution, we form the value of an initial inspection to the plant as a function of any level of pollution, $V_0(p)$. To approximate values for all pollution levels, we fit a piecewise-cubic Hermite interpolating polynomial function to the discrete Figure 4, Panel B bars, to obtain a smooth $\hat{V}_0(p)$ (Appendix Figure B1 plots the resulting function). The resulting value function determines plant incentives for preemptory abatement in the targeting stage.

VI.B. Targeting stage

We now turn to the targeting stage, which includes the pollution equation (1), the targeting function (5) and the distributions of pollution and cost shocks. Plants' decisions to *Run* abatement equipment (2) link pollution to inspection policy.

VI.B.1. Estimates

Columns 1 and 2 of Table 6 presents coefficient estimates where regulatory targeting is constrained to be optimal, conditional on the other parameter estimates. Columns 3 and 4 present unconstrained estimates. For each column pair, Panel A gives estimates of select parameters β, the effects of observables on targeting, in the inspection equation, and ϕ, the efficacy of abatement, in the pollution equation. Panel B gives estimates of targeting parameters λ and distributional parameters.

Panel A, column 1, replicates the reduced-form finding that treatment plants receive significantly more inspections. To put coefficient estimates in terms of inspections, we need to calculate marginal effects, which equate to about two inspections per year depending on the values of other plant covariates. In column 2, plants that run their abatement equipment are estimated to reduce pollution by -1.90 (standard error 0.16) logged standardized pollution points. A coefficient in logs of -1.90 is equivalent to an 85% reduction in pollution, which is similar to estimates of the efficacy of air pollution control equipment.

The estimates of the pollution shock distributions, in Panel B, suggest that the regulator observes only a small part of plant pollution, but uses this information to target plants with higher pollution signals. The standard deviation of the unobserved pollution shock is 1.03 (standard error 0.047) logged pollution points, as compared to 0.069 (0.003) for the observed component of pollution. Thus, while less than 1% of the variance of pollution ($= \sigma_1^2/(\sigma_1^2 + \sigma_2^2)$) is observable, the regulator targets aggressively. The maximum inspection coefficient $\hat{\lambda}_2 = 33$ and the shift parameter is $\hat{\lambda}_1 = -0.395$. These parameters imply large differences in inspections across plants for which the regulator observes different shocks, despite the fact that this observation is only a small part of overall pollution. For example, an audit-eligible, inspection control plant in Ahmedabad would expect to receive 0.56 inspections, with a shock at the 5th percentile of the observable part of pollution, and 3.48 inspections with a 95th-percentile shock.

In columns 3 and 4, we remove the optimal targeting constraints, and a few changes in the model estimates are notable. First, the inspection targeting parameters decline: $\hat{\lambda}_2$ falls from 33 to 10 and $\hat{\lambda}_1$ becomes less negative (-0.22, standard error 0.066, versus -0.395, standard error 0.003). Thus, targeting is slightly less aggressive in concentrating inspections in the dirtiest plants. At the same time, the share of the variance in pollution that is observable to the regulator increases substantially and consequently more variation in inspections is attributed to observable variance in pollution. Finally, the estimated effect of running equipment on abatement is smaller, -0.71 (0.308) log points instead of -1.90 (0.16). Thus optimal targeting, which is very aggressive, is justified, in the constrained model, if the amount of abatement achieved by running equipment is large. Otherwise, plants that are high-inspection but also high-cost will not abate, and the regulator does better spreading inspections around.

Given that optimal targeting is so aggressive, is it reasonable to assume that the regulator is targeting optimally? In Table 6, Panel C we report a distance metric test for the two independent constraints in the system. Under the null that the constraints are valid, the product of the sample size and the difference in the minimized value of the criterion for the constrained and unconstrained estimators is distributed χ^2_2 (Newey and McFadden, 1994). We reject the constrained estimates (p-value $= 0.0003$), implying that the regulator's targeting parameters (λ_1, λ_2) are not set optimally given the other parameters. However, Section VI.B.3. argues that the fit of the constrained model is still quite good.

VI.B.2. Sensitivity Analysis and Robustness Checks

To develop intuition for how features of our data affect parameter estimates, we use the measure introduced by Andrews et al. (2017) to provide a sensitivity analysis. This *sensitivity* measures how a parameter estimate would change, at the margin, given a change in one moment, holding fixed all other moments underlying the estimation. Since our structural estimates rely on experimental variation, this method is appealing to build intuition for how the estimates would change if the experimental results had been different.

Figure 5 reports the sensitivity of select parameter estimates (across panels) to select moments (row headers within panel). The length of each bar is the increase (solid bar) or decrease (hollow bar) in each parameter that would result from increasing the row moment by one standard deviation, *ceteris paribus*. Appendix D reports the full sensitivity matrix of parameters with respect to moments; here we highlight the sensitivity of targeting parameters λ, regulatory information $\sigma_1^2/(\sigma_1^2+\sigma_2^2)$ and abatement efficacy ϕ_2 to the experimental variation. All inspection moments refer to initial inspections and not follow-ups.

The differences in inspection moments between the treatment and control groups play a key role in estimating the targeting equation parameters. The targeting parameter λ_2 is most sensitive, amongst all moments, to the mean products of the treatment with inspections and inspections squared (Figure 5, Panels λ_1 and λ_2; see also Appendix Table D15). These sensitivities imply that if the observed mean inspections of treated plants increased by one inspection per year, from 2.93 to 3.93, with a corresponding increase in squared inspections, the estimated maximal inspections parameter would rise from $\hat{\lambda}_2 = 10.06$ to $\hat{\lambda}_2 = 26.23$ and λ_1 would decline from $\hat{\lambda}_1 = -0.22$ to $\hat{\lambda}_1 = -0.58$, showing more aggressive targeting in the status quo. Plainly, if the treatment increased initial inspections a lot, beyond the control level, it must have been because the targeting function was steep.

The share of pollution variance observed by the regulator (hereafter, information) is sensitive to the dispersion of inspections (mean inspections squared, conditional on mean inspections) and the interaction of inspections with treatment (Figure 5, Panel $\sigma_1^2/(\sigma_1^2 + \sigma_2^2)$). For example, a one standard deviation increase in mean inspections squared, from 7.53 to 17.12, conditional on mean inspections, would increase the estimated observable share of the pollution shock by approximately 0.31 on a small base of 0.02. As information improves, the regulator inspects plants more on the basis of pollution shocks observable to the regulator, but not the econometrician, increasing dispersion in inspections. Information is also increasing in the covariance between pollution and inspections because, in the model, if the regulator is more informed, it will assign more inspections to plants with high pollution.

The efficacy of pollution abatement $\hat{\phi}_2$ is sensitive to both pollution and cost moments. We find that $\hat{\phi}_2$ is increasing in the pollution residual in the treatment. If treatment pollution were higher by one standard deviation, then $\hat{\phi}_2$ would increase from -0.71 to -0.51 (Figure 5, Panel ϕ_2, moment Pollution Resid. $\times\, T$), indicating a decline in abatement efficacy. The interpretation is that if the treatment had reduced pollution less than observed, conditional on abatement decisions, the model would infer that abatement was less effective. The efficacy of abatement depends not just on the pollution equation, as it would in a single-equation model, but also, with a high sensitivity, on the cost of maintenance. If the mean maintenance cost conditional on running increased by \$100, the $\hat{\phi}_2$ coefficient would increase from -0.71 to -0.53 (Panel ϕ_2, moment $\mathbb{E}[c_j|Run]$).

Finally, we have fixed values for two parameters, σ_c and ρ, and here we check the robustness of our estimates to these assumptions. Appendix D shows estimates of the unconstrained targeting model for different values of σ_c and ρ. Changing the value of σ_c has very little impact on estimation of the targeting parameters. It has a larger effect on the estimated effect of abatement: moving from $\sigma_c = 0.25$ to $\sigma_c = 1.00$ increases the estimated $\hat{\phi}_2$ from -0.60 to -1.07. Both the lowest and highest estimate are within one standard deviation of the estimate we report for $\sigma_c = 0.50$. Changing the value of ρ changes the estimated β and λ vectors in the targeting function. The changes in β are roughly, but not exactly, proportional to changes in ρ, since ρ scales the argument of the targeting function. Because the parameters β and λ that are estimated adjust to offset changes in ρ, varying ρ does not have a large effect on the fit of the model to the inspection moments (Appendix Table D13).

VI.B.3. Model fit

There are several ways to assess the model's fit. Figure 6 compares the optimal targeting function with the unconstrained targeting function. The x-axis reports the observed component of pollution in standard deviations and the tick marks above the axis show the distribution of the observed component of pollution.[24] The y-axis reports the number of inspections for a control plant with the average value of $X_j'\hat{\beta}$, thus the figure shows the effect of the observable component of pollution on the number of inspections. The unconstrained targeting function is depicted with the solid black line and the dashed line represents the constrained optimal targeting function with current information. It is striking that the unconstrained and constrained targeting functions lie nearly atop one another throughout most of the distribution of the observed component of pollution; the exception is that the optimal function allocates more inspections to the plants with the highest observed pollution shocks. So although we statistically reject the constraints that require targeting be optimal, the constrained estimates fit the data well, as they describe a targeting rule that is very similar to the one derived from the unconstrained estimates. As a basis of comparison, the full information optimal targeting function is shown as the dotted line. For the estimated parameters, full information would lead the regulator to target somewhat more aggressively at plants with observed pollution shocks above the median.

To get a sense of the model fit, it helps to compare the distributions of inspections and

[24] Because the estimated $\hat{\sigma}_1$ differs between the constrained and unconstrained estimates, each is plotted in different units, corresponding to the standard deviation of observed pollution in that set of estimates.

pollution in the data to those generated by the model. We generate data using the estimated model parameters and a single simulation draw of the three model shocks. Figure 7, Panels C and D show the modeled distributions of inspections in the control and treatment groups, respectively, and can be compared to the true distribution in Panels A and B. The model and data distributions of inspections in the control are nearly identical, and show a similar truncation of the distribution at low levels. The treatment distribution in the model shows an upward shift in the mass of inspections and fits well, though the distribution of treatment inspections in the model is somewhat more skewed than in the data. Figure 7 shows the distribution of pollution in the data (Panel E) and the model (Panel F). The model was fit only to moments based on the mean and variance of pollution, but the fit throughout the distribution is good, with a similar modal value, in $[0.5\overline{p}, \overline{p}]$, and right-skewness. Overall, we conclude that our distributional assumptions provide parsimonious fits to the inspection and pollution data.

A more stringent test of the constrained model fit is the extent to which it matches the observed treatment effect on pollution compliance. Recall, from Figure 1, Panel A, that the experimental estimates show a narrow response to the treatment, with larger increases in pollution readings below the regulatory standard and decreases in readings spread out above the standard. Figure 1, Panel B reports coefficients from the same set of regressions for pollution bin dummies on treatment in the model-generated data, where the input data are plant characteristics, treatment assignments and draws from the distributions of pollution and cost shocks. It is striking that the model and the experimental results produce a similar pattern of abatement: the largest estimated increase in mass is beneath the regulatory standard, and the largest decrease in mass is just above, with only modest decreases in mass in the higher parts of the pollution distribution. At the same time, the fit is not perfect, as the model predicts a large increase in mass in the treatment that lies further beneath the regulatory standard than is seen in the data. This suggests that treated plants are able to control pollution down to the standard, but no more, whereas in the model abatement is assumed to be a discrete action and so may undershoot the standard. Nonetheless, the fit of the predicted and actual responses to treatment appears good.

VII Counterfactual Inspection Targeting

The value of the full structural estimation of the model is that it allows us to predict the effect of alternative policies that we did not experimentally evaluate. We use the model estimates to evaluate counterfactuals on optimal inspection targeting and pollution abatement that vary in regulatory budgets, discretion and information.

The basic framework is the regulator's problem of maximizing abatement (3) subject to the budget constraint (4). We take as given regulatory penalties and thus the outcome of the penalty stage, and examine the effect of cross-sectional changes in targeting on plants' abatement decisions. We consider these medium-run counterfactuals, matched to the horizon of the experiment, since they change inspection targeting but neither the penalty function nor the abatement capital available to plants. If the targeting policies were changed permanently, the penalty function might be adjusted in response.

Within this framework, we consider several targeting regimes. Within each regime, we vary the budget constraint $\overline{I}$ and measure the reduction in pollution achieved. The first regime is a uniform targeting rule that gives every plant the same number of inspections, regardless of the observed pollution shock. This regime requires no information on pollution. Second, a targeting regime with discretion, where the regulator solves (3), but where the regulator has only the sparse information in observed pollution. Third, a targeting regime with discretion where the regulator has full information. This regime is not feasible, with currently installed technology, but would be feasible with the installation of continuous emissions monitoring systems like those used in other countries for some pollutants (e.g., sulfur dioxide and nitrogen oxides in the United States). Fourth, a hybrid regime, like in the experimental treatment, where an initial 1.47 inspections (the control mean) are allocated with discretion, and additional inspections beyond that are allocated uniformly across plants.

Figure 8 traces out the abatement achieved by the alternative targeting regimes. Each line shows the total pollution abatement in units of the regulatory standard, relative to the latent pollution level $\tilde{P}$, (vertical axis) as a function of the total inspection budget per plant per year (horizontal axis) under a different regime. The dashed-and-dotted (blue) line shows abatement under the uniform rule that requires all plants to be inspected the same number of times. The solid (red) line shows abatement under the optimal targeting rule where targeting is based on observed pollution only. The dashed (black) line shows abatement under the optimal targeting rule where the regulator is assumed to observe all variation in pollution. The horizontal range of the figure extends from zero inspections per plant per year to four, which is the prescribed regulatory minimum for large-scale, high-polluting plants.

There are several notable findings. First, a uniform inspection rule does very poorly at low levels of the budget constraint. At one inspection per plant, the average abatement is negligible, and at the observed control inspection level of about 1.5 annual inspections per plant (i.e., the first vertical line), mean abatement is 0.06 standards. The reason for this poor performance is that few plants have trivial maintenance costs, and so spreading inspections out over all plants causes the regulator to be spread quite thin; the result is that plants, which know the targeting rule, generally find the cost of pre-emptive abatement to exceed its benefit.

Second, at any given budget constraint, regulatory discretion increases the abatement the regulator achieves. At the observed level of 1.5 annual inspections per plant in the control, shown by the first vertical line, the total abatement is about three times greater (0.17 standards vs. the prior 0.06) when the regulator allocates inspections with its information (red solid line), as compared to a constant rule (blue dashed-and-dotted line). Put another way, to achieve the same level of abatement with uniformly allocated inspections would require a budget of about 2.2 inspections per plant. The value of discretion using partial information is especially strong, in relative terms, at low budget constraints, but tapers off as the uniform rule eventually allocates enough inspections for expected penalties to cover the maintenance costs of many plants.

Third, our simulation of the treatment regime, combining discretionary inspections with additional random inspections, shows a weakened marginal response to treatment inspections. The treatment, adding roughly 1.5 additional uniform inspections per year, moves along the dotted red line from the left-hand to the right-hand dotted vertical line, increasing average abatement by 0.14 standards. The marginal effect on abatement of adding random inspections to discretionary inspections allocated by the regulator, as was done in the experiment, is small. This is because inframarginal plants, which were not targeted by the regulator, do not receive enough random inspections to induce abatement. Had the inspections in the treatment been allocated according to the regulator's optimal rule, along the solid red line, we estimate abatement would have been about 15% greater. In other words, doubling inspection while keeping the same level of discretion would have decrease pollution by 0.16 standard deviations. This increase in abatement is equivalent to an arc elasticity of pollution with respect to additional discretionary inspections of -0.27, relative to the control mean level of inspections and pollution.

Fourth, there is a substantial benefit to better information in a discretionary regime. The dashed black line gives the share of plants abating under an optimal targeting regime where the regulator observes all variation in pollution ($\sigma_2 = 0$) as would be the case where the regulatory has access to a perfectly functioning monitoring technology. The difference in abatement under an optimal targeting regime with full (dashed black) versus estimated (solid red) information is 30% of baseline abatement (on average across budgets), and the amount of abatement from better information is increasing in the inspection budget (the gap between the black solid and blue dashed lines increases). Full information in a discretionary regime, at the control budget, is as valuable as doubling inspections in a non-discretionary regime. It is apparent that full information allows the regulator to more precisely target its inspections and this substantially increases abatement.

Full information can in principle be achieved by the use of continuous emissions monitoring systems (CEMS), devices which report real-time pollution readings. The Indian Central Pollution Control Board has developed standards for Continuous Emissions Monitoring Systems for particulate matter, the most severe air pollution problem in India, and has recently mandated their use for hundreds of large factories around the country (Central Pollution Control Board, 2013, 2014). CEMS have much higher fixed costs but fairly low variable costs relative to inspections. We estimate CEMS monitoring of particulate matter in Gujarat today costs about USD 1800 per plant-year, on an amortized basis, whereas a single in-person inspection with air pollution sampling costs USD 70, including the costs of staff, travel and lab analysis. The efficacy of CEMS as a substitute for in-person inspections will depend on the evolution of costs, as devices are installed and used more widely, and whether a monitoring regime with CEMS provides incentives for accurate data reporting, rather than only the installation of monitoring devices.

These results help reconcile a number of facts about the effect of the inspection treatment and the constraints on regulation. The inspection budget, given the present penalty structure, would induce practically no abatement if inspections were allocated evenly across plants. Discretion has value because it allows the regulator to concentrate inspections in the plants with high observable signals of pollution and this greatly increases abatement, even though half of plants are left alone. This targeting would grow more valuable if information were improved.

VIII Conclusion

We conducted an experiment on environmental regulation of industrial plants in Gujarat, India. The treatment bundled increased inspection resources with a removal of discretion over which plants to inspect. The striking finding is that the treatment had little effect on plants' pollution emissions.

We unbundle the experimental results with exhaustive data and a structural model. Our data set on the regulatory process, pollution readings and penalties opens the black box of interactions between the GPCB and regulated plants; we are not aware of a comparable data set from any country on regulatory process and outcomes. We set out a structural model of the interactions between the regulator and plants to separate the roles of resources and discretion in regulatory enforcement. At the GPCB's current level of inspections, we find that removing discretion would be damaging: the inspections chosen by the regulator induce three times more abatement than would the same number of randomly-assigned inspections. With respect to the experiment, the abatement achieved by the intervention's increase in inspection resources was somewhat undercut by the removal in discretion in targeting the extra inspections.

The structural analysis also uncovered that poor information on plant emissions hinders enforcement when the regulator has discretion. A technology that provided the regulator with perfect information on plant emissions would increase total abatement by roughly thirty percent at the status quo inspection rate, the same reduction as would be achieved by a one-third increase in the inspection budget, if the added inspections were allocated with discretion. The importance of reliable information in discretionary regimes has not been widely appreciated in the literature or policy debates.

Our analysis underscores that strict environmental standards and high levels of pollution co-exist as long as regulators have weak enforcement tools. The study contrasts three prominent policy levers. First, the impact of adding resources alone is likely positive but, at least in settings similar to Gujarat, modest. Second, reducing regulatory discretion can undercut enforcement, even in settings with weak institutions. This result on discretion stands in contrast to the conventional wisdom in policy circles and the academic literature that removing discretion is the best safeguard against corruption. Third, improved monitoring of plant emissions can strengthen enforcement, as regulators have poor information on which plants are deserving of sanction. We saw a similar result in our parallel study of environmental audits, which found that plants reduce emissions when third-party auditors report their emissions to the regulator more truthfully (Duflo et al., 2013). While achieving perfect information, with continuous emissions monitoring systems, may be possible in principle, further research is needed to address whether these devices are in fact a reliable and cost-effective substitute for in-person monitoring, in a setting with weak institutional capacity.

Regulators and governments generally consider reforms on many different dimensions. A randomized experiment along each dimension will often be infeasible and so many policy interventions that are tested experimentally are bundled, in the manner of resources and reduced discretion in our experiment. This paper has combined an experiment and structural analysis to unbundle several aspects of regulatory enforcement in a critical policy domain.

References

Adhvaryu, A., Kala, N., and Nyshadham, A. (2016). Management and shocks to worker productivity. *Mimeo, University of Michigan.*

Aghion, P., Burgess, R., and Redding, S. (2008). The unequal effects of liberalization: Evidence from dismantling the license Raj in India. *American Economic Review*, 98(4):1397–1412.

Aghion, P. and Tirole, J. (1997). Real and formal authority in organizations. *Journal of Political Economy*, 105:1–29.

Andrews, I., Gentzkow, M., and Shapiro, J. M. (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics*, page qjx023.

Banerjee, A. V. and Duflo, E. (2014). Do firms want to borrow more? Testing credit constraints using a directed lending program. *The Review of Economic Studies*, 81(2):572–607.

Besley, T. and Burgess, R. (2004). Can labor regulation hinder economic performance? Evidence from India. *Quarterly Journal of Economics*, 119(1):91–134.

Besley, T. and Persson, T. (2010). State capacity, conflict and development. *Econometrica*, 78:1–34.

Boyer, M. and Laffont, J.-J. (1999). Toward a political theory of the emergence of environmental incentive regulation. *The RAND Journal of Economics*, pages 137–157.

Burgess, R., Hansen, M., Olken, B., Potapov, P., and Sieber., S. (2012). The political economy of deforestation in the tropics. *Quarterly Journal of Economics*, 127(4):1707–1754.

Central Pollution Control Board (2012). National ambient air quality status and trends in india: 2010. *Available at* `http://www.cpcb.nic.in/upload/NewItems/NewItem_192_NAAQSTI.pdf` *(last accessed September 18, 2017).*

Central Pollution Control Board (2013). Specifications and guidelines for continuous emissions monitoring systems (cems) for pm measurement. *CPCB/e-Publication/2013-14*.

Central Pollution Control Board (2014). Directions . . . regarding self-monitoring of compliance. *CPCB/B-29016/04/06/PCI-I/5401.*

Chang, T., Zivin, J. G., Gross, T., and Neidell, M. (2016). The effect of pollution on worker productivity: evidence from call-center workers in china. Technical report, National Bureau of Economic Research.

Chen, Y., Ebenstein, A., Greenstone, M., and Li, H. (2013). Evidence on the impact of sustained exposure to air pollution on life expectancy from China's Huai River policy. *Proceedings of the National Academy of Sciences*, 110(32):12936–12941.

Duflo, E., Greenstone, M., Pande, R., and Ryan, N. (2013). Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India. *The Quarterly Journal of Economics*, 128(4):1499–1545.

Ebenstein, A., Fan, M., Greenstone, M., He, G., and Zhou, M. (2017). New evidence on the impact of sustained exposure to air pollution on life expectancy from china's huai river policy. *Proceedings of the National Academy of Sciences*, 114(39):10384–10389.

Environmental Protection Agency (2012). Air pollution technology fact sheets. Technical report, Technology Transfer Network: U.S.-Mexico Border Information Center on Air Pollution.

Estache, A. and Wren-Lewis, L. (2009). Toward a theory of regulation for developing countries: Following Jean-Jaques Laffont's lead. *Journal of Economic Literature*, 47(3):729–770.

Graff-Zivin, J. and Neidell, M. (2012). The impact of pollution on worker productivity. *American Economic Review*, 102(7).

Greenstone, M. and Hanna, R. (2014). Environmental regulations, air and water pollution and infant mortality in India. *NBER Working Paper 17210.*

Greenstone, M., List, J. A., and Syverson, C. (2012). The effects of environmental regulation on the competitiveness of us manufacturing. Technical report, National Bureau of Economic Research.

Greenstone, M., Nilekani, J., Pande, R., Ryan, N., Sudarshan, A., and Sugathan, A. (2014). Lower pollution, longer lives: Life expectancy gains if India reduced particulate matter to air-quality standards. *Mimeo, Harvard Kennedy School.*

Gujarat Pollution Control Board (2008). Annual action plan 2008-09. Technical report, Govt. of Gujarat, http://gpcb.gov.in/pdf/aplan.pdf.

Hanna, R. N. and Oliva, P. (2010). The impact of inspections on plant-level air emissions. *The BE Journal of Economic Analysis & Policy*, 10(1).

He, J., Liu, H., and Salvo, A. (2016). Severe air pollution and labor productivity: Evidence from industrial towns in china. *IZA Discussion Paper No. 8916.*

Jia, R. (2014). Pollution for promotion. *Mimeo, UCSD.*

Laffont, J.-J. (1994). *Regulation of pollution with asymmetric information.* Springer.

Laffont, J.-J. (2005). *Regulation and Development.* Cambridge University Press.

Laffont, J.-J. and Tirole, J. (1993). *A Theory of Incentives in Procurement and Regulation.* MIT Press.

Leaver, C. (2009). Bureaucratic minimal squawk behavior: Theory and evidence from regulatory agencies. *American Economic Review*, 99(3):572–607.

Magat, W. A. and Viscusi, W. K. (1990). Effectiveness of the epa's regulatory enforcement: The case of industrial effluent standards. *JL & Econ.*, 33:331.

Magnac, T. and Thesmar, D. (2002). Identifying dynamic discrete decision processes. *Econometrica*, 70(2):801–816.

Ministry of Environment and Forests (1986). The environment (protection) rules, 1986. *Gazette of India*, 844(E).

Ministry of Environment and Forests (1999). Prevention and control of pollution (uniform consent procedure) rules, 1999. Technical Report Part II, Section 3(ii), Schedule IV, The Gazette of India, http://icpe.in/pdf/2.pdf.

Newey, W. K. and McFadden, D. (1994). Large sample estimation and hypothesis testing. *Handbook of econometrics*, 4:2111–2245.

Rasul, I. and Rogger, D. (2013). Management of bureaucrats and public service delivery: Evidence from the Nigerian civil service. *Mimeo, UCL.*

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033.

Rust, J. (1994). *Advances in Econometrics, Sixth World Congress*, chapter Estimation of dynamic structural models, problems and prospects: Discrete decision processses. Cambridge University Press.

Ryan, S. P. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–1061.

Shimshack, J. P. (2014). The economics of environmental monitoring and enforcement. *Annu. Rev. Resour. Econ.*, 6(1):339–360.

Stigler, G. (1971). The theory of economic regulation. *Bell Journal of Economics and Management*, 2(1):3–21.

Tanaka, S. (2013). Environmental regulations on air pollution in China and their impact on infant mortality. *Mimeo, Tufts University.*

The Economic Times (October 16, 2014). Pm narendra modi unveils labour reforms; launches universal account number for employees. Technical report, The Economic Times, https://economictimes.indiatimes.com/news/economy/policy/pm-narendra-modi-unveils-labour-reforms-launches-universal-account-number-for-employees/articleshow/44833365.cms.

IX Figures

Figure 1: Effect of Treatment on Pollution Distribution

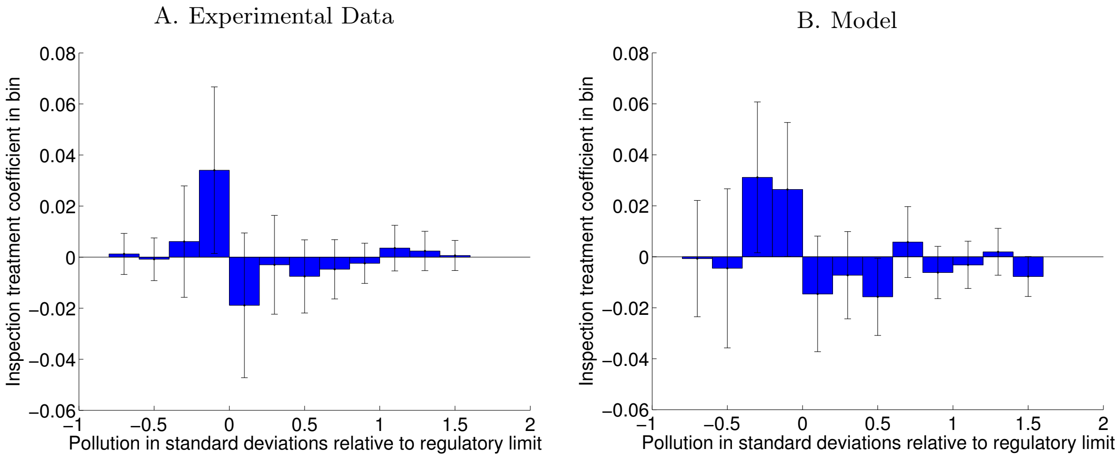

The figure panels report coefficients on the inspection treatment assignment from regressions of dummies for a pollution reading being in a given bin, relative to the regulatory standard, on inspection treatment, audit treatment, inspection $\times$ audit treatment, a dummy for being audit-eligible and region fixed effects. Panel A reports coefficients from such regressions on the experimental data and Panel B reports coefficients from the same regressions run on model-generated data using the constrained model estimates of Table 6. Pollution readings are standardized by subtracting the regulatory standard for each pollutant and dividing by the pollutant's standard deviation; bins are 0.2 standard deviations wide and centered at the regulatory standard shown by the vertical line. Each plant has multiple pollutant observations and regressions are run pooled for all pollutants together. The whiskers show 95% confidence intervals for the inspection treatment coefficient.

Figure 2: Regulatory Targeting of Extreme Polluters

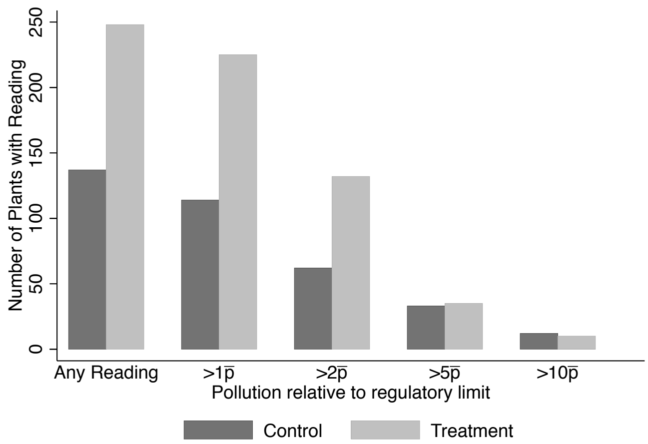

The figure shows the number of plants with pollution readings either taken or that fall in various bins, relative to the regulatory standard, during the first year of the intervention for the control and treatment groups, respectively. The first pair of bars shows the number of plants that had at least one pollution reading taken. The remaining four pairs show the number of plants with at least one reading above the standard ($>1\overline{p}$), more than 2 times the standard ($>2\overline{p}$), more than 5 times the standard ($>5\overline{p}$) and more than 10 times the standard ($>10\overline{p}$).

Figure 3: Actions of the Regulatory Machine and Plant

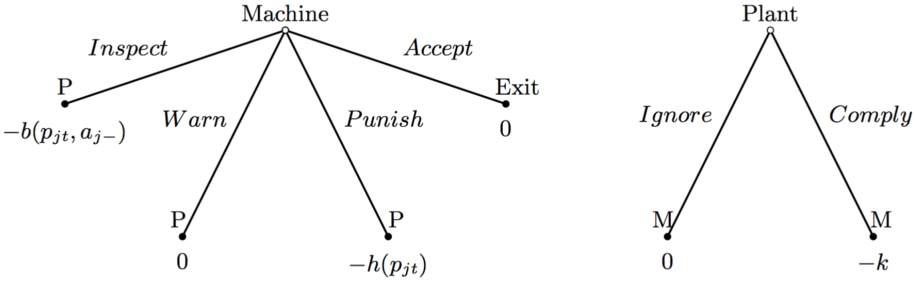

A. Actions of the Regulatory Machine B. Actions of the Plant

The figure gives the actions of the regulatory machine and plant at each node and the terminal nodes give the payoffs in each round for the plant. The penalty stage begins with an inspection where the Regulatory Machine (M) observes p_{j1}. The machine can take four actions. If M *Inspects*, M gets a new signal of pollution and the plant may have to offer a bribe with payoff $-b(p_{jt}, a_{j-})$. If M *Warns*, there is no cost to the plant. If *M Punishes*, the plant faces a cost $-h(p_{jt})$. After each of these moves the plant *Ignores* or *Complies* and *M* moves again. If M *Accepts*, the stage ends.

Figure 4: Value of Environmental Regulation for Plants

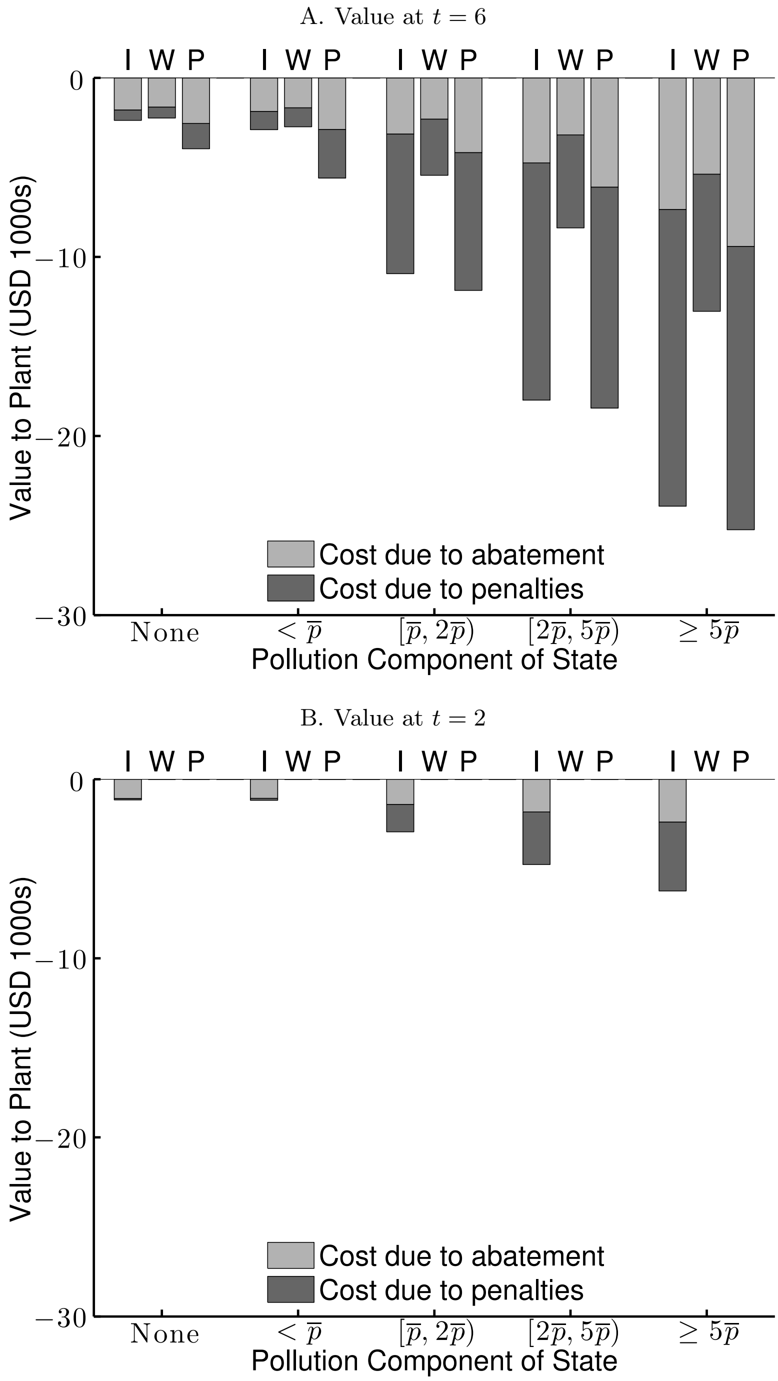

The figure shows the cost of regulation to plants in thousands of US dollars as measured by the expected discounted value of different states in the penalty stage. Values are divided between expected discounted future abatement costs (light grey) and expected discounted future regulatory penalties (dark grey), both of which, as costs to the plant, have negative value. The figure shows three different dimensions of the state along which plant value varies. First, the panels show the time dimension, with panel A evaluated when it is the plant's turn to move at t=6, and panel B at t=2. Second, the five clusters of bars on the horizontal axis show different maximum lagged pollutant readings observed during the prior inspection. Third, within each group, the letters I, W and P show how the value to the plant changes if the regulatory machine's lagged action was *Inspect*, *Warn* or *Punish*, respectively.

Figure 5: Sensitivity of Targeting Parameters to Moments

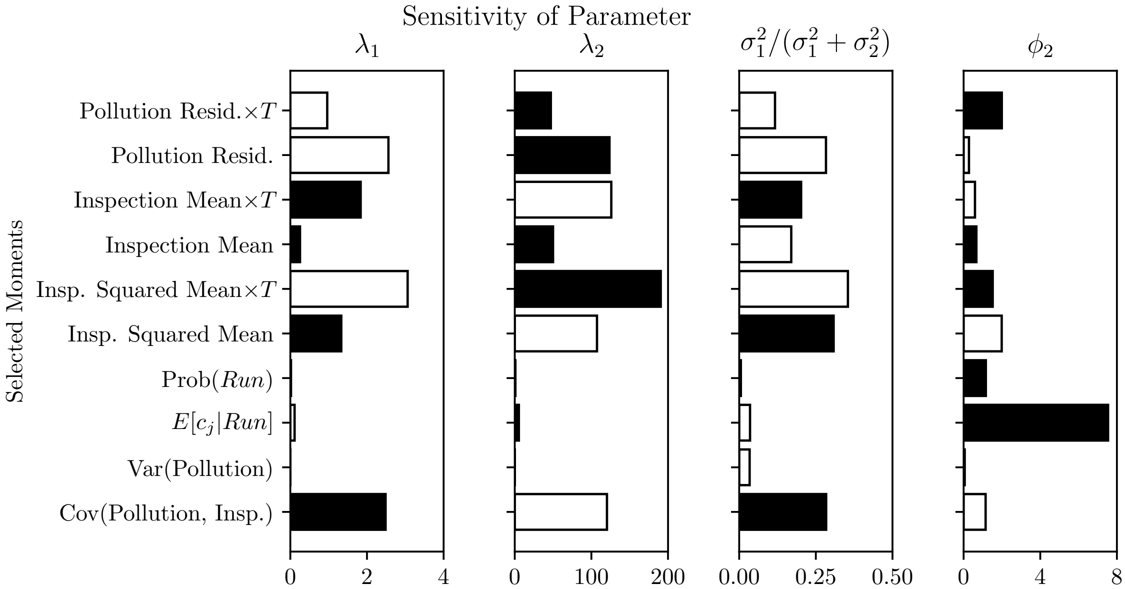

The figure shows the Andrews et al. (2017) sensitivity measure of selected targeting model parameter estimates with respect to selected moments used to estimate the model. The panels, left to right, show the sensitivity of the parameters or functions of parameters λ_1, λ_2, $\sigma_1^2/(\sigma_1^2+\sigma_2^2)$, with respect to moments indicated by the row headers. The length of each bar is the local sensitivity of the parameter with respect to the row moment, measured in units of the parameter value per one standard deviation change in the moment. We omit from the rows the products of pollution residuals and inspection means with observable plant characteristics other than treatment status. Black filled bars indicate positive sensitivity and hollow bars negative sensitivity.

Figure 6: Observed and Optimal Targeting Rules

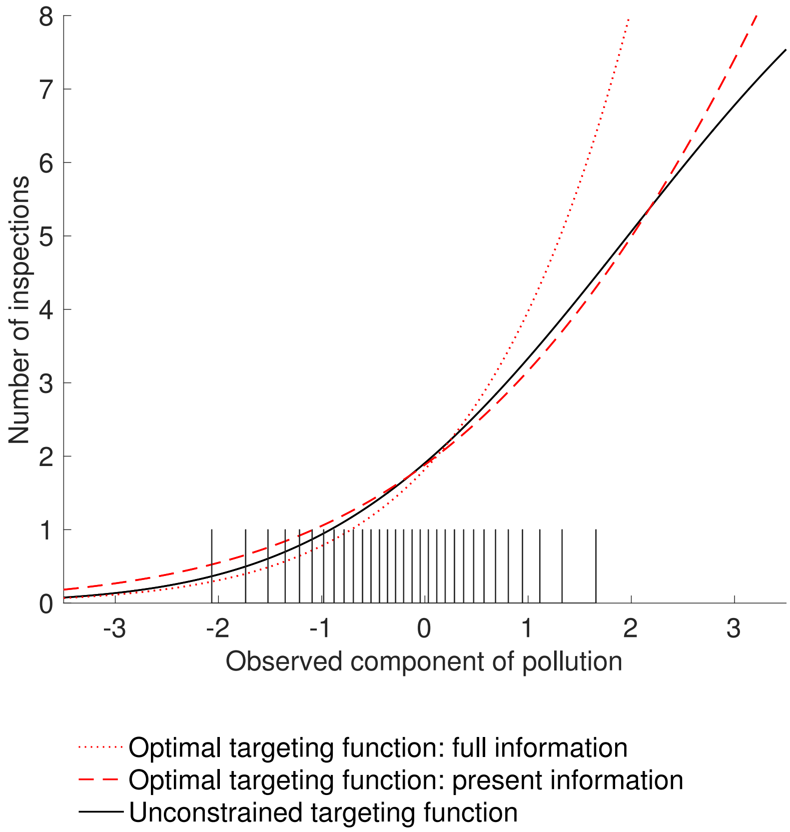

The figure shows several inspection targeting rules, which assign plants an annual rate of inspection as a function of the component of pollution observed by the regulator. Two of the targeting rules are based on the estimates of Table 6: the dashed (red) line gives the estimated optimal targeting function from the constrained estimates in the table, and the solid (black) line gives the estimated targeting function from the unconstrained estimates. The third, dotted (red) line gives the optimal targeting rule under an alternative regime where the regulator observes all the variance in pollution. The vertical spikes on the horizontal axis represent a normal distribution of pollution shocks, centered around the mean value of the observable characteristics $X_j'\beta_1$ on which inspections are targeted.

Figure 7: Model Fit to Inspections and Pollution

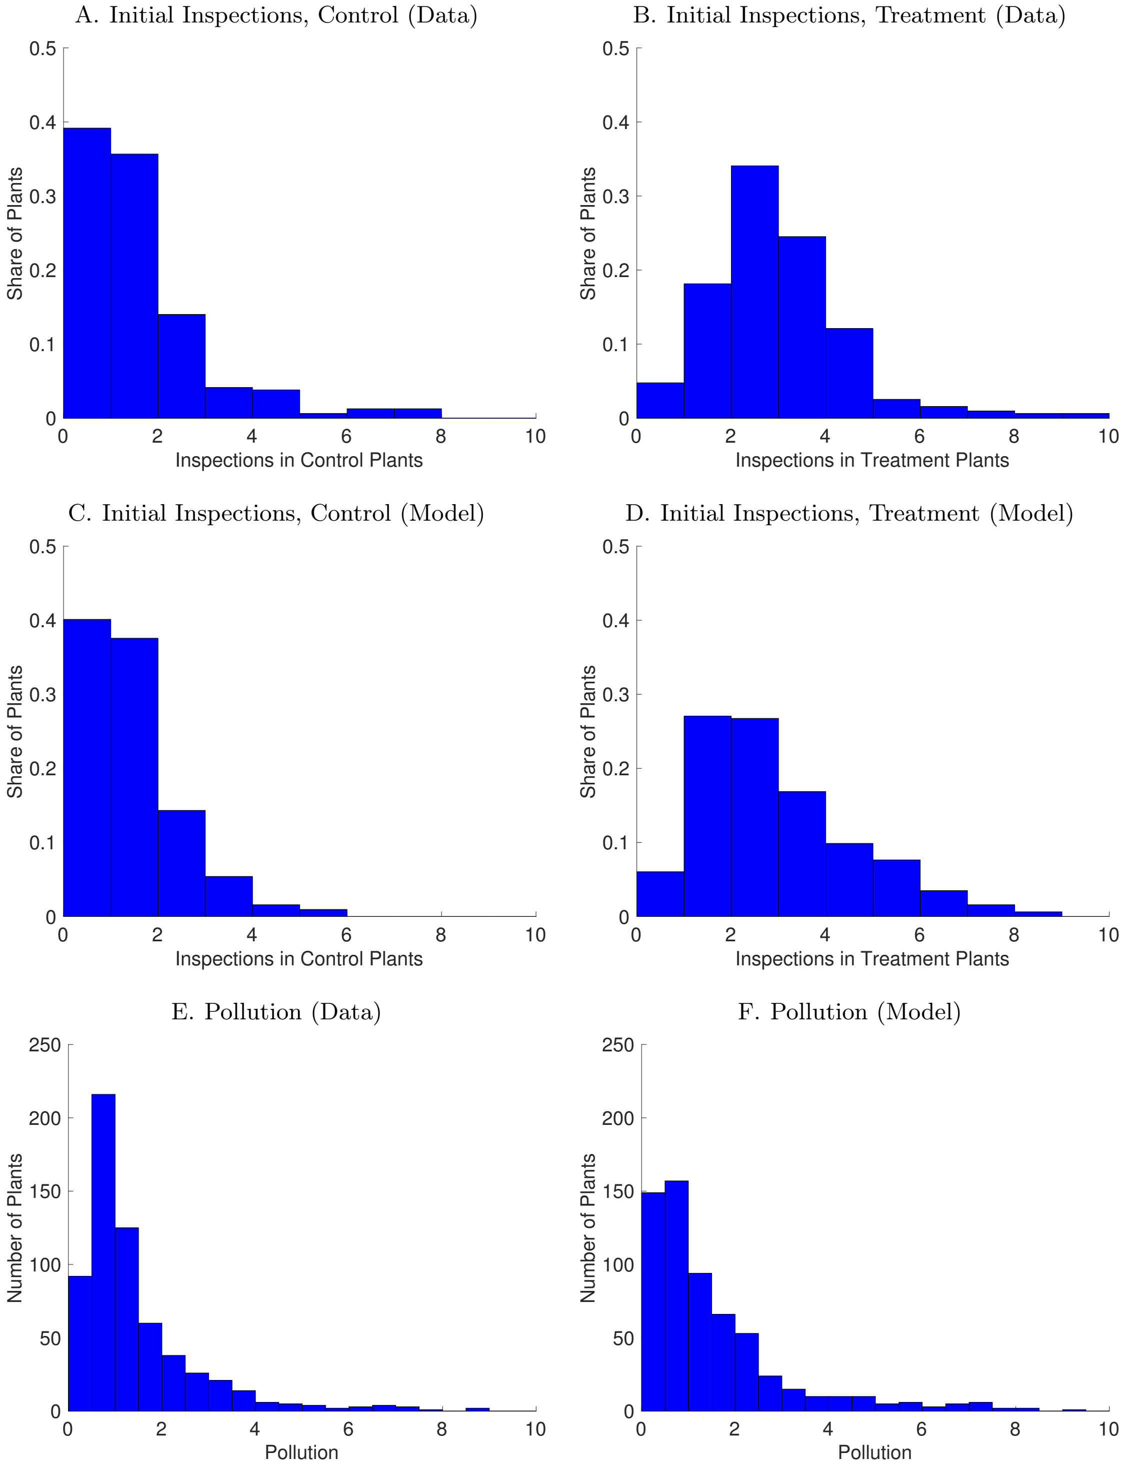

The figure compares the distributions of inspections and pollution in the model to those in the experimental data. Panels A through D show the distributions of the annual inspection rate (i.e., inspections per year). Inspections include only initial inspections and not follow-ups. Panels A and B give the distributions in the data in the control and treatment groups, respectively, using administrative records of inspection reports. Panels C and D give the same distributions in the model. Panels E and F give the distribution of pollution in the data and in the model, respectively. The units of pollution are units of the regulatory standard $\overline{p}$, such that a value of 2 represents pollution at twice the standard, etc.

Figure 8: Value of Discretion: Abatement by Information Regime and Budget Constraint

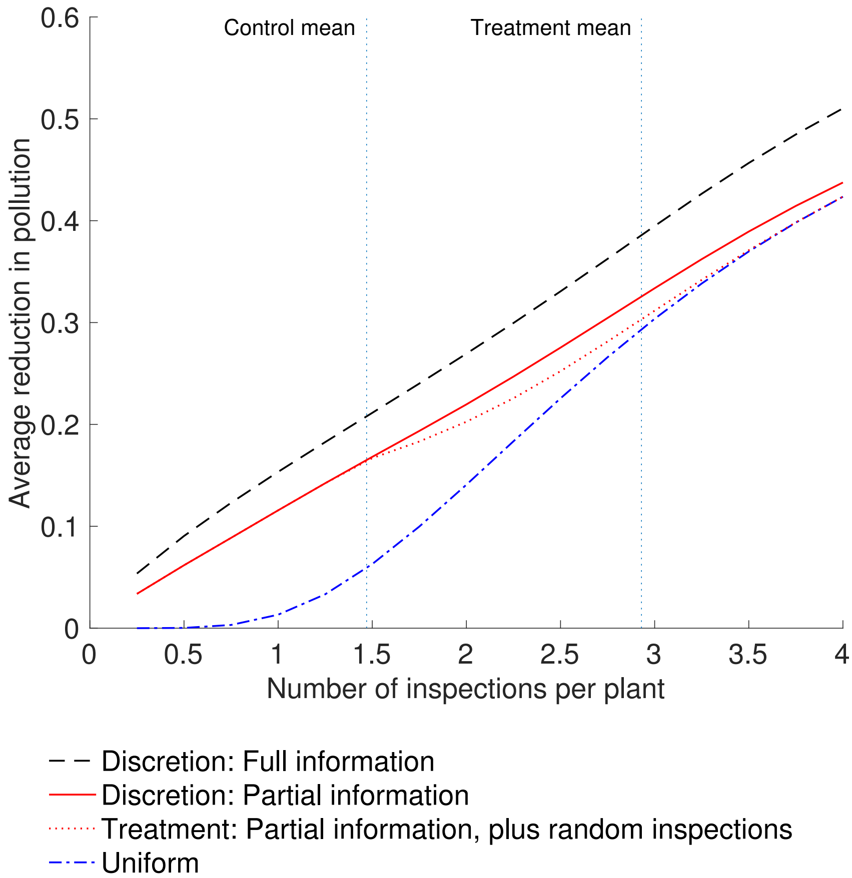

The figure shows the amount of pollution abatement achieved, in units of regulatory standards of abatement per plant, under different counterfactual inspection regimes. Regimes vary in the information used by the regulator and the inspection budget available per plant. Each line shows average abatement per plant against the total inspection budget per plant per year (horizontal axis) under a different regime. The dashed and dotted (blue) line shows abatement under a minimum threshold rule that requires all plants to be inspected the same number of times. The solid (red) line shows abatement under the optimal targeting rule where targeting is based on the observed component of pollution. The dashed (black) line shows abatement under the same optimal targeting rule where the regulator is assumed to observe all variation in pollution. The dotted (red) line is a hybrid regime, meant to reflect the experimental treatment, where inspections up to the control level of inspection (1.5 inspections per plant per year) are allocated with discretion, according to the optimal rule, and additional inspections beyond that level are allocated evenly across all plants.

X Tables

Table 1: Regulatory Interactions During Experiment

	Control	Treatment	Difference
Panel A. Inspections by Treatment Status			
Number inspections assigned in treatment, annual	0	2.12	2.12***
	[0]	[0.57]	(0.026)
Total inspections, annual over treatment	1.40	3.11	1.71***
	[1.59]	[1.77]	(0.11)
Initial inspections, annual over treatment	1.28	2.79	1.50***
	[1.38]	[1.52]	(0.094)
Observations	480	480	
Panel B. Perceived Inspections by Treatment Status			
Perceived Inspections, 2008	2.53	2.66	0.13
	[1.42]	[1.40]	(0.10)
Perceived Inspections, 2009	2.78	3.16	0.38***
	[1.44]	[1.37]	(0.100)
Perceived Inspections, 2010	2.92	3.62	0.71***
	[1.58]	[1.46]	(0.11)
Total perceived notices and closures received, 2010	0.27	0.30	0.025
	[0.64]	[0.70]	(0.048)
Observations	388	403	
Panel C. Regulatory Actions by Treatment Status			
Pollution reading ever collected at plant (=1)	0.38	0.60	0.21***
	[0.49]	[0.49]	(0.032)
Any pollution reading above limit at plant (=1)	0.34	0.55	0.22***
	[0.47]	[0.50]	(0.031)
Number of pollution readings above limit at plant	1.17	2.84	1.67***
	[2.58]	[3.67]	(0.20)
Total citations	0.15	0.35	0.20***
	[0.42]	[0.69]	(0.037)
Total water citations	0.046	0.12	0.071***
	[0.22]	[0.37]	(0.020)
Total air citations	0.021	0.042	0.021*
	[0.14]	[0.20]	(0.011)
Total closure warnings	0.094	0.17	0.077***
	[0.34]	[0.48]	(0.027)
Total closure directions	0.16	0.20	0.042
	[0.48]	[0.54]	(0.033)
Total bank guarantees	0.060	0.065	0.0042
	[0.27]	[0.25]	(0.017)
Total equipment mandates	0.027	0.040	0.013
	[0.19]	[0.23]	(0.014)
Total utility disconnections	0.040	0.042	0.0021
	[0.22]	[0.20]	(0.013)
Observations	480	480	

The table shows differences in actual inspection rates (Panel A), perceived inspection rates (Panel B), and other regulatory actions (Panel C) between the treatment and control groups of plants during the treatment period of approximately two years. Columns (1) and (2) show means with standard deviations in brackets. Column (3) shows the coefficient from regressions of each variable on treatment, where each regression includes region fixed effects and a control for the audit sample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Endline Pollution and Compliance on Treatments

	(1)	(2)	(3)	(4)
	Panel A. Plant-level Costs			
	Capital costs		Maintenance costs	
	(USD '000s)	Any (=1)	(USD '000s)	Any (=1)
Inspection treatment (=1)	-0.221	0.0213	0.838*	0.00974
	(0.453)	(0.0344)	(0.499)	(0.0224)
Plant characteristics	Yes	Yes	Yes	Yes
Audit experiment	Yes	Yes	Yes	Yes
Control Mean	2.050	0.567	0.264	0.108
Observations	791	791	791	791
	Panel B. Plant-by-pollutant level Pollution			
	Pollution		Compliance	
Inspection treatment (=1)	-0.105		0.0366*	
	(0.0839)		(0.0213)	
Audit treatment (=1)	-0.187**		0.0288	
	(0.0849)		(0.0258)	
Audit × inspection treatment (=1)	0.286**		-0.0365	
	(0.142)		(0.0353)	
Control mean	0.682		0.614	
Observations	4168		4168	

The table shows intent-to-treat effects of inspection treatment assignment on plant costs and pollution outcomes. Panel A shows regressions for plant costs estimated at the plant level. Costs are divided into capital and maintenance costs based on descriptions of each expenditure (See Appendix A). Cost amounts are in USD thousands. Capital costs, which are reported as lump-sum in the survey, are amortized to an equivalent constant annual expenditure (using an interest rate of 20 percent and a 10-year equipment lifespan). Plant characteristic controls include dummies for size, use of coal or lignite as fuel, high waste water generated, and all regions. Audit experiment includes dummies for audit treatment and audit sample. Robust standard errors in parentheses. Panel B shows regressions for pollution and compliance at the plant-by-pollutant level. Pollution consists of air and water pollution readings for each plant, taken during the endline survey, where each pollutant is standardized by dividing by its standard deviation. Compliance is a dummy for each pollutant being below its regulatory standard. Controls include region fixed effects and a dummy for the audit sample. Standard errors clustered at the plant level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Structure of Penalty Stage Actions

Round	Regulatory Action				Plant Action			
	Inspect	Warn	Punish	Accept	Ignore	Comply	N	% left
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1	100.0	0.0	0.0	0.0			7423	100.0
2					99.6	0.4	7423	
3	1.0	9.5	2.2	87.3			7423	100.0
4					92.8	7.2	941	
5	23.3	4.8	5.3	66.6			941	12.7
6					91.1	8.9	314	
7	18.8	11.8	9.9	59.6			314	4.2
8					83.5	16.5	127	
9	21.3	5.5	18.1	55.1			127	1.7
10					82.5	17.5	57	
11	26.3	3.5	10.5	59.6			57	0.8
12					87.0	13.0	23	
13	26.1	4.3	8.7	60.9			23	0.3
14					77.8	22.2	9	
15+	16.7	8.3	0.0	75.0	100.0	0.0	9	0.1
Total without inspections	0.0	4.6	1.6	42.7	50.2	0.9	7824	
Total	31.0	3.2	1.1	29.4	34.6	0.6	25217	

The table reports actions taken by the regulatory machine and by the firm in the penalty stage using administrative data. Figure 4 defines actions and their payoffs and Table 1 maps them to regulatory documents. Each column (1) - (6) of the table gives the probability, within that row, of the party moving at that round taking the action indicated in the column header. Column (7) gives the total number of actions observed in that round and column (8) gives the percentage of penalty stages that continue up to at least that round. The penalty stage always starts with an inspection. Action rounds within the stage then alternate between actions of the regulatory machine and sample plants. The penalty stage ends when the machine *Accepts*. Rounds after the 15th round are not shown and the row 15+ summarizes these rounds; 6 chains go at least 17 rounds and 4 chains go 19 rounds.

Table 4: Multinomial Logit Model of Action Choice Conditional on State

Party to move:	Regulatory Machine			Plant
	Inspect	Warn	Punish	Comply
	(1)	(2)	(3)	(4)
Lagged regulatory actions				
Warn, lag 1	0.33	-2.05***	-2.10***	-0.23
	(0.23)	(0.32)	(0.31)	(0.30)
Punish, lag 1	1.80***	-2.22***	-0.53*	1.29***
	(0.23)	(0.56)	(0.30)	(0.26)
Lagged plant actions				
Firm: Comply, lag 1	-1.80***	-1.03**	-0.82**	-0.53
	(0.32)	(0.47)	(0.37)	(0.66)
Last observed pollution reading				
0-1x	-0.38	-0.25	0.052	-0.18
	(0.23)	(0.16)	(0.24)	(0.38)
1-2x	-0.20	0.55***	0.37**	0.39*
	(0.16)	(0.098)	(0.18)	(0.23)
2-5x	-0.17	0.84***	0.70***	0.74***
	(0.17)	(0.10)	(0.17)	(0.22)
5x+	0.27	0.63***	1.15***	0.90***
	(0.21)	(0.16)	(0.21)	(0.26)
Period				
Constant	-4.41***	-2.47***	-3.91***	-5.71***
	(0.13)	(0.057)	(0.11)	(0.21)
t > 3	2.91***	1.26***	2.56***	2.59***
	(0.25)	(0.28)	(0.27)	(0.33)
t > 5	0.073	-0.35	-0.50	0.18
	(0.21)	(0.32)	(0.30)	(0.28)
t > 7	0.059	-0.55	0.55*	0.50*
	(0.24)	(0.37)	(0.29)	(0.28)
N	8897			8897

The table reports coefficients from multinomial logit models for the action choice probabilities of the regulatory machine and the plant conditional on the state within the penalty stage. See Table A1 for action definitions. Plant and regulatory actions are reported in administrative data by the regulator. The omitted action for the regulator is *Accept* and for the plant is *Ignore*, so the coefficients are to be interpreted as the effect of each component of the state on the party taking the specified column action relative to the omitted action. Pollution readings are taken during inspections throughout the treatment period. The omitted pollution reading is null, which occurs when the regulator inspects but does not take a pollution reading. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Estimates of Plant Utility Parameters (US$ 1000s)

	Whether bribes given if no compliance?			
	No bribes		On inspection	
	(1)	(2)	(3)	(4)
	Parameters of penalty function			
τ_0	53.54		28.12	36.71
	(24.68)		(20.88)	(22.92)
τ_1		39.57		
		(28.17)		
τ_2		54.11		
		(27.43)		
τ_3		41.51		
		(19.15)		
	Parameters of bribe function			
ν_0			9.67	
			(3.07)	
ν_1				10.93
				(3.48)
ν_2				9.72
				(3.99)
ν_3				5.83
				(4.96)
	Standard deviation of action shock			
σ	5.02	5.88	5.11	4.83
	(0.46)	(0.30)	(0.39)	(0.43)
Observations	1474	1474	1474	1474

The table presents pseudo-maximum likelihood estimates of the parameters of the plant profit function, from the estimation of the plant's dynamic problem in the penalty stage. The four columns represent different specifications for the penalties and bribes the plant must pay. The parameters τ give the value of penalties applied by the regulator, by choosing the action *Punish*, conditional on the pollution component of the state being between the standard and twice the standard (τ_0), between twice and five times the standard (τ_1) and above five times the standard (τ_2). The column 3 and 4 estimates also include estimates of bribes, in addition to formal penalties. The parameters ν, for which estimates are reported in columns (3) and (4), give the value of bribes given by the plant in penalty specifications where the plant is assumed to give bribes, if it has not already *Complied* in the stage, at the second inspection and later inspections. The final parameter σ is the standard deviation of the plant's action-specific payoff shock. Observations are those at which the plant moves in rounds t=4 and onwards; t=2 is omitted because a large number of actions in that round are imputed (see text). Inference is by the bootstrap over 100 samples with replacement, where samples are taken at the level of the plant-chain (i.e., series of interactions) stratified on the maximum pollution reading observed in the chain. Standard errors equal to the standard deviation of bootstrap estimates are in parentheses.

Table 6: Estimates of Targeting Stage Parameters

	Constrained		Unconstrained	
	Initial Inspections (1)	Log Pollution (2)	Initial Inspections (3)	Log Pollution (4)
Panel A. Targeting and Pollution Equations				
Inspection treatment	0.095		0.162	
	(0.009)		(0.025)	
Run equipment (=1)		-1.902		-0.711
		(0.160)		(0.308)
Inspection targeting shift	-0.395		-0.220	
parameter (λ_1)	(0.003)		(0.066)	
Inspection targeting level	33.022		10.064	
parameter (λ_2)	(1.876)		(3.137)	
Constant		0.212		-0.009
		(0.109)		(0.102)
Panel B. Distributions of Pollution and Maintenance Cost Shocks				
Standard deviation of observed	0.069		0.111	
pollution shock (σ_1)	(0.003)		(0.022)	
Standard deviation of unobserved	1.033		0.864	
pollution shock (σ_2)	(0.047)		(0.042)	
Mean of log maintenance	2.388		1.833	
cost (μ_c)	(0.061)		(0.334)	
Panel C. Test of Targeting Optimality Constraints				
Distance metric test statistic χ^2_2	16.1039			
Test p-value	0.0003			

The table reports parameters of the targeting stage of the model. The first two columns report estimates from the constrained model where the regulator is constrained to target optimally based on observed pollution shocks, and the second two columns report estimates from the unconstrained model. Within each pair of columns, the first column (1 and 3, respectively) reports the coefficients on the inspection equation and the second column (2 and 4) reports coefficients on the pollution equation. Panel B reports estimates of the distributional parameters for pollution and cost shocks under each model. Panel C reports the results of a test of the constraints that require optimal targeting (equations (16)-(18) in Appendix C.3). The data set is a cross-section of all sample plants. Inspections are calculated as the number of initial inspections per year for each plant over the course of the approximately two-year experiment, and pollution is the maximum (over pollutants) logged standardized endline survey pollution reading, where the standardization is in units of the regulatory standard for each pollutant (e.g., twice the pollutant-specific standard is a value of log(2)). Both the inspection and the pollution models additionally include the observable characteristics of audit treatment assignment, audit × inspection assignment, an audit sample dummy and region dummies (not reported). The parameters are estimated via generalized method of moments using a mix of analytic and simulated moments as described in Section V. Standard errors are bootstrapped with B = 100, S = 5,000.

A Data and Supplementary Analysis

A.1 Mapping documents to regulatory and plant actions

Section II.C. describes the regulatory documents we use to assign actions to the regulator and plants in our model of regulation. This Appendix gives more detail on the documents involved, the mapping of these documents to actions, and the linkage of these actions into chains of related actions.

Table A1: Mapping of Player Actions to Raw Documents

Action	Document	Description
	Panel A. Regulatory Actions	
Inspect	Inspection report	Analysis of air and water samples; report on plant characteristics.
Warn	Letter to plant	Non-threatening letter ordering improvement in pollutant concentrations.
	Citation	Threatening letter demanding explanation for high pollution levels, missing permit to operate, or missing pollution abatement equipment.
	Closure notice	Notice that the plant will be ordered to close in 15 days if the plant does not take action to improve pollution.
Punish	Closure direction	Order to plant to close immediately.
	Utility notice	Notice that water or electricity has been disconnected.
Accept	Revocation of closure direction	Permission to start operation.
	None	No further regulatory action within 60 days.
	Panel B. Firm Actions	
Comply	Equipment installed	Notice that pollution abatement equipment has been installed.
	Process installed	Notice that a process has been installed.
	Bond posted	Letter from bank to GPCB explaining that the firm has posted a guarantee against future misconduct.
Ignore	Letter to regulator	Letter of protest to GPCB. Challenges parameter readings and other directives.
	None	Implied by consecutive GPCB actions without plant response.

Our approach is based on rules that the Gujarat Pollution Control Board follows as a regulatory agency. All regulatory actions that GPCB takes regarding regulated plants must be documented. These documents are of distinct types, are all dated and are often cross-referenced;

that is, a document taking a regulatory action will often cite explicitly a prior violation of pollution standards as grounds for that action. Plants are not, like the regulator, obligated to follow rules in their correspondence. However, plants often do respond in writing to regulatory orders, in particular when they want to document that they have made a costly abatement investment.

Table A1 and the discussion in Section II.C. covered the basic types of documents and their mapping to actions. Many of the documents have a one-to-one correspondence to actions. For the regulator, the action *Inspect* is always indicated by an Inspection Report, which is a particular type of regulatory document. The action *Punish* is always indicated by a Closure Direction sent to plants (or the plants' utility). We group multiple documents of varying severity under the action *Warn*, including simple citations noting violations of pollution standards and more severe warnings that threaten closure. These documents all show that the regulator has found a violation, and often threaten action if the plant does not remediate, but they have in common that they do not impose a direct cost on plants. For plants, the main action of interest is whether a plant chooses to *Comply* by installing abatement equipment. This action is typically documented by a invoice or other record from an environmental consultant or vendor that installed the equipment.

Regulatory interactions with plants occur in groups of related actions we call chains. Many actions are responses to earlier actions. For example, if the regulator inspects a plant and finds a violation, this may result in a later action of *Warn* or *Punish*, or in the plant possibly choosing to *Comply* to avoid punishment. We link actions in a chain, for the same plant, using explicit references in the documents underlying each action and the dating of those documents, as follows:

1. *Link documents that reference one another.*
   1.1. Link documents if one document explicitly cites (by document number or date) an inspection or other earlier document.
   1.2. If there is no exact match, link documents if a near match exists that differs by at most one digit.
2. *Link documents dated close to one another.*
   2.1. Keep all documents already linked in (1.) together (we may call these sub-chains).

2.2. Link remaining documents (or groups of documents) to those that happened soon after and are plausibly logically related, using the following rules for each pair of candidate documents:

- Documents to link if within 30 days:
  i. Any regulator action followed by plant *Ignore*.
  ii. Any regulator action followed by regulator *Inspect*.
  iii. Regulator *Punish* (closure notice, closure notice, utility confirmation of action) followed by plant *Ignore* or *Comply* (equipment installation, process installation, bank guarantee posted).
- Documents to link if within 60 days:
  i. Regulator *Inspect* followed by *Warn* (letter, show cause notice, closure notice, closure direction).
  ii. Regulator *Inspect* followed by regulator *Accept*.
  iii. Plant *Comply* followed by regulator *Accept* (revocation of closure direction).

3. *Impute missing actions to enforce structure of the plant's problem.*

3.1. *Chains start with the action Inspect.* If a chain does not start with an inspection, append earlier inspections that occured within 1 month of chain start. If no such recent inspection occurred, truncate the chain before first inspection.

3.2. *Chains end with the action Accept.* Impute regulatory action *Accept* at the end of the chain if there are no further follow-up actions within 60 days.

3.3. *Chains alternate moves between the plant and regulatory machine.* Impute the plant move *Ignore* between any consecutive GPCB actions.

In stage (3.) we impute actions to enforce an alternate-move structure to the plant's problem. The main assumption in imputing moves is that the plant had an opportunity to respond to all regulatory actions with, even if it in fact did not respond. That is, we take the absence of a response as the action *Ignore*. Similarly, we assume the regulator could have continued pursuing plants at the end of each chain, and that if it does not this represents the action *Accept*. This assumption is weak because our data is comprehensive and all regulatory actions against plants are documented, so that if the regulator did later *Punish* or *Warn* a plant, this action would have been observed.

Table A2: Sample Chain

Round (1)	Player (2)	Action (3)	Document (4)	Date (5)
1	GPCB	Inspect	Inspection report	2008-09-05
2	Plant	Ignore		2008-09-05
3	GPCB	Punish	Closure Direction	2009-01-12
4	Plant	Comply	Equipment installed	2009-01-28
5	GPCB	Inspect	Inspection report	2009-01-31
6	Plant	Ignore		2009-01-31
7	GPCB	Inspect	Inspection report	2009-02-04
8	Plant	Ignore		2009-02-04
9	GPCB	Punish	Closure Direction	2009-05-22
10	Plant	Comply	Process installed	2009-05-30
11	GPCB	Inspect	Inspection report	2009-06-16
12	Plant	Comply	Process installed	2009-06-16
13	GPCB	Accept	Revocation of Closure Direction	2009-06-24

The table displays a 13-round chain of interactions between GPCB and one plant during the experiment. Column (3) indicates the category of action, while Column (4) reports the underlying document to which the action corresponds. *Ignore* actions by the plant in Rounds 2, 6 and 8 have been imputed based on adjacent actions in the chain and hence Column (4) is left blank in these rounds. All chains begin with a regulatory inspection, *Inspect*. The players then alternate moves until the regulator decides to *Accept* the plant's compliance for the time being, which terminates the chain. Table 1 in the paper describes the way in which the actions are mapped to the underlying documents, and the Data Appendix provides a full explanation of the rules used to construct the chains.

Table A2 gives an example of the result of this linking process. GPCB initially inspects a plant and the plant does nothing. GPCB, presumably on reviewing the inspection report and pollution samples, then orders that the plant be closed. The plant responds, several weeks later, by installing abatement equipment. GPCB repeatedly inspects the plant over the next several months, apparently finds the remedy inadequate, and again orders the plant closed. The plant responds by modifying its production process. GPCB is then satisfied and revokes the closure order. This example gives a sense of the rich back-and-forth that is possible to observe in the chained action data.

A.2 Abatement costs

Measures of plant abatement cost come from the endline survey. The survey asked plants to describe each piece of pollution control "equipment installed or upgrades made (including routine maintenance such as filter changes, etc.)" since the start of the experiment, to record the corresponding expenditure, and to verify the equipment or upgrade took place.

We flag these abatement expenditures as either capital or maintenance expenditures and sum them to the plant level. The determination of whether an expenditure is a capital or maintenance expenditure is based on string matching with the text description of each investment. Enumerators indicated maintenance using the word maintenance or change, as in the action of changing a filter or other replaceable part. Expenditures with descriptions containing any of the following strings are therefore coded as maintenance: "chang," "mainten," "maintain," "maintan," "(bag.*bag)." The "(bag.*bag)" string captures variations on "Change of bag filter bag," a common maintenance activity for air pollution control devices. Expenditures containing none of these strings are coded as capital expenditures.

For the purposes of Table 2, Panel A, which compares capital and maintenance costs, we amortize capital costs, which are reported as lump sums, into an equivalent annual expenditure. We calculate the constant annual expenditure such that the sum of the present value of the expenditure over the equipment lifespan equals the observed up-front capital expenditure (with an interest rate of $r = 0.20$ and a 10-year equipment lifespan).

A.3 Experimental integrity: covariate balance and attrition

This subsection verifies the integrity of the experiment both *ex ante* and *ex post*. First, we check the balance of covariates before the experiment started. Then, we check for differential attrition during the experiment. Finally, we check our model assumption that the regulatory acted similarly against treatment and control plants conditional on the results of an inspection.

Table A3: Experimental Design: Treatment Assignments

	Inspection control	Inspection treatment	Total
Audit control	120	120	240
Audit treatment	116	117	233
Not audit eligible	244	243	487
Total	480	480	960

The table reports the number of plants assigned to each combination of the inspection treatment and the audit treatment of Duflo et al. (2013). Inspection treatment status is either control or treatment. With respect to audit, only some plants are audit-eligible (see text). Conditional on being eligible for audit, plants are assigned to inspection treatment or control.

Table A3 describes the cross-cutting experimental design. Plants were assigned to inspection treatment status conditional on audit treatment status (Duflo et al., 2013). Since only certain plants are eligible for audits, there are three possible audit treatment statuses: not audit eligible, and, conditional on being eligible, audit control or treatment. The inspection treatment status is therefore orthogonal to audit treatment status and eligibility.

Table A4 compares the balance of the inspection treatment assignment on fixed plant characteristics, in Panel A, and on regulatory interactions between plants and the regulator, like inspections and violations of pollution standards, in Panel B. The table uses administrative data and the regulatory interactions are measured over the last full calendar year (2008) prior to the experiment starting in August, 2009. Each row considers a separate plant characteristic; columns 1 and 2 report the means for control and treatment plants, respectively, using administrative data for the year prior to the experiment. Column 3 reports the coefficient α_2 on the inspection treatment dummy T_j from the following regression, where for each outcome Y_j for plant j in region r,

$$Y_{jr} = \alpha_r + \alpha_1 AuditSample_j + \alpha_2 T_j + \epsilon_j \tag{7}$$

where α_r are region effects and $AuditSample_j$ is a dummy for a plant belonging to the audit sample (i.e., being audit eligible, rather than being assigned to the audit treatment).[25]

We find that inspections, pollution readings and citations are balanced by treatment assignment. Of 18 baseline measures reported, there is a significant difference between the treatment and control groups at the ten percent level on only one measure.

Table A5 reports overall levels of attrition in the experiment. About 18 percent of plants did not report complete the endline survey. Most of this attrition, 13 percentage points, was due to plants that closed during the experiment. Table A6 compares attrition across the treatment and control groups. One may be concerned that the extra scrutiny of treatment plants would drive them out of business. We find, to the contrary, that attrition is not differential across treatment arms (the point estimate for the effect of treatment on attrition is negative).

[25]Since only one region, Ahmedabad, contains both audit-eligible and audit-ineligible plants, this specification is equivalent to a full set of region $\times$ eligible effects.

Table A4: Inspection Treatment Covariate Balance

	Control (1)	Treatment (2)	Difference (3)
Panel A. Plant Characteristics			
Capital investment Rs. 50m to Rs. 100m (=1)	0.087	0.071	-0.017
	[0.28]	[0.26]	(0.017)
Located in industrial estate (=1)	0.33	0.37	0.032
	[0.47]	[0.48]	(0.027)
Textiles (=1)	0.45	0.45	-0.0092
	[0.50]	[0.50]	(0.020)
Dyes and Intermediates (=1)	0.13	0.16	0.027
	[0.34]	[0.36]	(0.022)
Effluent to common treatment (=1)	0.37	0.35	-0.021
	[0.48]	[0.48]	(0.031)
Waste water generated (kl / day)	192.1	196.8	4.30
	[310.9]	[316.4]	(16.2)
Air emissions from boiler (=1)	0.50	0.52	0.019
	[0.50]	[0.50]	(0.020)
Panel B. Regulatory Interactions in Year Prior to Study			
Number of inspections	1.22	1.25	0.026
	[1.32]	[1.32]	(0.079)
Inspections below prescribed (=1)	0.42	0.39	-0.031
	[0.49]	[0.49]	(0.029)
Number of pollution readings	3.64	3.92	0.28
	[5.65]	[5.58]	(0.31)
Pollution reading ever collected (=1)	0.40	0.44	0.048*
	[0.49]	[0.50]	(0.027)
Any pollution reading above limit (=1)	0.34	0.38	0.031
	[0.48]	[0.48]	(0.026)
Citations	0.22	0.20	-0.023
	[0.51]	[0.55]	(0.034)
Closure warnings	0.056	0.052	-0.0044
	[0.31]	[0.32]	(0.020)
Closure directions	0.075	0.077	0.0019
	[0.31]	[0.34]	(0.021)
Bank guarantees posted	0.019	0.029	0.010
	[0.15]	[0.21]	(0.012)
Equipment mandates	0.24	0.25	0.0082
	[0.54]	[0.53]	(0.029)
Any utility disconnection (=1)	0.010	0.0021	-0.0083
	[0.10]	[0.046]	(0.0051)
Observations	480	480	

The table tests for the balance of covariates by inspection treatment status using administrative data from the regulator covering the year prior to the experiment. Columns (1) and (2) show means with standard deviations in brackets. Column (3) shows the coefficient on treatment from regressions of each characteristic on treatment and region fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A5: Attrition in the Endline Survey

	N (1)	% (2)
Survey completed	791	82.4
Plant closed	124	12.9
Plant refused survey	5	0.5
Other	40	4.2
Total	960	100.0

The table shows how many plants completed the endline survey, and the reasons for attrition for those that did not. *Plant closed* includes plants that were permanently closed (111), plants that were temporarily closed, and plants where production was temporarily suspended. *Refused survey* includes plants that were operating at the time of the visit, but that refused to respond to the questions in the survey. *Other* includes plants that moved to an unknown address, and plants for which an incorrect address had been recorded

Table A6: Endline Attrition by Inspection Treatment Status

	Treatment (1)	Control (2)	Difference (3)
Survey completed	0.840	0.808	0.031
	[0.367]	[0.394]	(0.024)
Plant closed	0.123	0.135	-0.013
	[0.329]	[0.343]	(0.022)
Plant refused survey	0.008	0.002	0.006
	[0.091]	[0.046]	(0.005)
Other	0.029	0.054	-0.025*
	[0.168]	[0.227]	(0.013)
Observations	480	480	

The table shows differences in endline responses and reasons for attrition between the treatment and control groups. Columns (1) and (2) show means with standard deviations in brackets. Column (3) shows the coefficient on treatment from regressions of each characteristic on inspection treatment assignment, region fixed effects, and audit sample control. Reported are treatment effects, with region controls. *Plant closed* includes plants where production. was temporarily suspended, and plants that were temporarily or permanently closed. *Refused survey* includes plants that were in production at the time of the visit, but that refused to respond to the questions in the survey. *Other* includes plants that moved to an unknown address, and plants for which an incorrect address had been recorded. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A7: Probability of Regulator Acceptance on Treatment and Round

	(1)	(2)	(3)	(4)	(5)	(6)
Inspection treatment assigned (=1)	0.0129	0.00779	0.0127	0.0164	0.0106	-0.0280
	(0.0112)	(0.00926)	(0.0110)	(0.0137)	(0.00835)	(0.0207)
Constant	0.827***	0.869***	0.863***	0.861***	0.891***	0.910***
	(0.00831)	(0.00701)	(0.00894)	(0.00974)	(0.00700)	(0.0134)
Period						
t > 3		-0.219***			-0.325***	-0.323***
		(0.0252)			(0.0350)	(0.0557)
t > 5		-0.0367			0.0293	0.0548
		(0.0479)			(0.0366)	(0.0610)
t > 7		-0.0571			-0.0224	-0.125*
		(0.0634)			(0.0396)	(0.0682)
Period × Treatment						
t > 3 × Inspection Treatment		0.0219				0.0322
		(0.0344)				(0.0464)
t > 5 × Inspection Treatment		-0.0610				-0.0248
		(0.0638)				(0.0790)
t > 7 × Inspection Treatment		0.0757				0.0920
		(0.0821)				(0.0961)
Lagged regulatory actions						
Warn, lag 1					0.169***	0.150***
					(0.0364)	(0.0477)
Punish, lag 1					-0.129***	-0.101*
					(0.0413)	(0.0535)
Lagged plant actions						
Firm: Protest, lag 1					0.0482*	0.0535
					(0.0272)	(0.0363)
Firm: Comply, lag 1					0.248***	0.259***
					(0.0329)	(0.0416)
Last pollution reading						
0-1x			-0.0102	0.0140	0.000305	
			(0.0158)	(0.0211)	(0.0119)	
1-2x			-0.0661***	-0.0534***	-0.0524***	-0.0661***
			(0.0134)	(0.0207)	(0.0105)	(0.0205)
2-5x			-0.109***	-0.120***	-0.0896***	-0.119***
			(0.0151)	(0.0251)	(0.0120)	(0.0216)
>5x			-0.160***	-0.178***	-0.116***	-0.144***
			(0.0260)	(0.0375)	(0.0189)	(0.0308)
Pol Reading × Treatment						
0-1x × Inspection Treatment				-0.0424		
				(0.0309)		
1-2x × Inspection Treatment				-0.0221		0.0239
				(0.0271)		(0.0280)
2-5x × Inspection Treatment				0.0186		0.0497*
				(0.0313)		(0.0291)
>5x × Inspection Treatment				0.0330		0.0452
				(0.0516)		(0.0428)
p-value for F-test of no × Inspection Treatment terms	0.249	0.608	0.245	0.388	0.204	0.315
Inspection control mean	0.827	0.827	0.827	0.827	0.827	0.827
Observations	8897	8897	8897	8897	8897	4089

Does not include region fixed effects. Omitted actions are Ignore (for Regulator) and Inspect (for plant). Omitted pollution reading for column (6) is No Pollution Reading Taken. Standard errors clustered at plant level in parentheses with * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In the model we assume that the regulatory machine acts similarly against both treatment and control plants (Section V). This assumption is based on the design of the experiment, since the officials that decide to act against plants were not informed of whether an inspection report came from a treatment or control plant. Here we present additional statistical evidence supporting that this assumption held in practice.

Table A7 regresses a dummy for regulatory machine *Acceptance* (i.e., leaving the plant alone) in a given round on the observable characteristics noted down during an inspection, treatment status, and interactions of treatment status and observables. The columns 1 through 6 add progressively richer specifications of observables and interactions. In column 1, there are no controls. Column 2 controls for round of the game and its interactions with treatment, column 3 controls for pollution, column 4 controls for pollution and its interactions with treatment, and column 5 and 6 additionally control for past regulatory and plant actions. In each column, we report at bottom the p-value from an F-test that the coefficients on treatment and all treatment interactions with observables are jointly zero. If the regulator followed-up with treatment plants differently, we would expect the treatment main effect or the interaction of the treatment and some observable to be significant (for example, if the regulator did not pursue treatment plants after finding high pollution readings). We fail to reject the joint null with p-values between approximately 0.20 and 0.60. We conclude that the regulatory machine treats plants similarly conditional on the facts observed in an inspection.

A.4 Letter treatment

Table A8 reports the results of a treatment that sent plants a letter reminding them of their obligations to meet emissions limits, shortly before the endline survey. The letter had no significant effect on pollution or compliance.

A.5 Compliance placebo checks

Table A9 reports the results of placebo regressions showing treatment effects on compliance at various compliance thresholds. Column 1 shows the regression of compliance on treatment where compliance is defined as a pollution reading below the true standard $\overline{p}$, and columns 2 through 4 show the same regression with placebo standards set at several multiples of the true

Table A8: Endline Pollution and Compliance on Letter Treatment

	Pollution (1)	Compliance (2)
Inspection treatment assigned (=1)	-0.0160	0.0248
	(0.0866)	(0.0238)
Letter treatment assigned (=1)	-0.0482	0.0311
	(0.0928)	(0.0241)
Inspection treatment × Letter treatment (=1)	0.0326	-0.00340
	(0.130)	(0.0345)
Inspection and Letter control mean	0.652	0.595
Observations	4168	4168

The table shows regressions of pollution (column 1) and compliance (column 2) on inspection and letter treatment assignments. Observations are at the plant-by-pollutant level, where pollution consists of air and water pollution readings for each plant, taken during the endline survey, and each pollutant is standardized by dividing by its standard deviation. Compliance is a dummy for each pollutant being below its regulatory standard. The table regresses these outcomes on inspection treatment assignment, letter treatment assignment and inspection × letter treatment. The letter treatment was a letter sent by the regulator to plants shortly before the endline survey reiterating the terms of plants' environmental consents and reminding them of their obligations to meet emissions limits. Specifications also include region fixed effects. Standard errors clustered at the plant level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

level. There is a significant effect, at the ten percent level, only at the true threshold (column 1).

A.6 Regulatory targeting using endline survey data

Table A10 presents additional evidence on regulatory targeting, plant-level regression estimates for how inspections depend on a plant's underlying pollution. To get at targeting under the status quo, we restrict the regression to the control group. To measure the regulator's response to latent pollution levels, we run the regression in the year after the experiment ended, using our own endline survey readings, which were not reported to the regulator, as the measure of pollution. Columns 1 through 4 use a categorical dummy as the independent variable, where missing readings are coded zero (and indicated by a separate dummy variable), readings beneath the standard $\overline{p}$ are coded 1, between $(\overline{p}, 2\overline{p}]$ as 2, $(2\overline{p}, 5\overline{p}]$ as 3 and $5\overline{p}$ and above as 4; column 5 shows dummies for the underlying pollution categories.

Table A10, column 1 shows that, conditional on plant characteristics, higher endline survey pollution readings predict more inspections, even though the regulator does not observe these

Table A9: Placebo Check of Alternate Compliance Thresholds

	$[0, \overline{p})$ (1)	$[0, 2\overline{p})$ (2)	$[0, 5\overline{p})$ (3)	$[0, 10\overline{p})$ (4)
Inspection treatment assigned (=1)	0.0366*	0.0144	0.00323	-0.000368
	(0.0213)	(0.0193)	(0.0131)	(0.00824)
Audit treatment assigned (=1)	0.0288	0.0154	0.0123	0.0166*
	(0.0258)	(0.0238)	(0.0162)	(0.00917)
Audit × inspection treatment (=1)	-0.0365	-0.0245	-0.0109	-0.0106
	(0.0353)	(0.0316)	(0.0214)	(0.0116)
Inspection and audit control mean	0.614	0.813	0.928	0.975
Observations	4168	4168	4168	4168

The table presents regression estimates for compliance on inspection treatment assignment and audit treatment assignment, using pollution levels taken from the endline survey. Compliance is defined as pollution being below N times the limit $\overline{p}$, with N being 1, 2, 3, 5 and 10 respectively. Pollution standardized by backcheck standard deviation. Standard errors in parentheses. Includes audit treatment and treatment interaction controls, and year and region fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

readings. Each higher category of pollution results in 0.170 more inspections per year (standard error 0.0978), relative to a compliant plant.[26] Regulatory targeting on unobserved pollution remains as strong when adding controls for audit treatment status during the experiment (column 2) and recent regulatory penalties (column 3). Perhaps most strikingly, when adding recent pollution readings that the regulator itself took as controls, we find the coefficient on endline survey pollution readings remains unchanged (column 4). In column 5, we separate the components of the categorical endline pollution variable and find that belonging to the highest pollution bin is associated with 0.665 additional inspections per year (standard error 0.331). The dummies for pollution bins above the standard are jointly significant in predicting later inspections ($F_{4,367} = 3.46$, p-value < 0.01). The regulator appears to target inspections based on signals of plant pollution beyond its own past pollution readings.

[26] Plant observable characteristics like belonging to a dirtier sector, being of greater scale and generating more wastewater also predict more inspections (coefficients not reported).

Table A10: Regulatory Targeting on Unobserved Pollution in the Control Group

	Dependent variable: Number of inspections in one year after endline survey				
	(1)	(2)	(3)	(4)	(5)
Endline pollution bin (0-4)	0.170*	0.173*	0.182*	0.172*	
	(0.0978)	(0.103)	(0.101)	(0.103)	
Endline Pollution $\in [1\overline{p}, 2\overline{p})$ (=1)					0.459
					(0.278)
Endline Pollution $\in [2\overline{p}, 5\overline{p})$ (=1)					0.539
					(0.349)
Endline Pollution $\geq 5\overline{p}$ (=1)					0.665**
					(0.331)
Constant	2.058***	1.995***	1.898***	1.893***	1.943***
	(0.302)	(0.463)	(0.461)	(0.532)	(0.486)
Plant characteristics	Yes	Yes	Yes	Yes	Yes
Audit treatment assignment		Yes	Yes	Yes	Yes
Recent regulatory actions			Yes	Yes	Yes
Recent pollution readings				Yes	Yes
Mean dependent variable	1.392	1.392	1.392	1.392	1.392
F-stat p-value					0.00859
R^2	0.213	0.213	0.259	0.284	0.285
Observations	388	388	388	388	388

The table regresses the number of regulatory inspections in the year after the endline survey on pollution readings as measured during the survey, in the inspection control group of plants. As endline pollution readings were not reported to the regulator, these readings represent an unobserved component of pollution, from the regulator's perspective. Endline pollution bin is a categorical variable that takes the value of 0 for plants with no pollution readings, 1 if pollution is in $[0, \overline{p})$, 2 if in $[\overline{p}, 2\overline{p})$, 3 if in $[2\overline{p}, 5\overline{p})$, and 4 if above $5\overline{p}$. All specifications also include a dummy for a plant having no pollution reading. Column 5 separates pollution bin into a set of dummy variables and reports the coefficient on each; the omitted dummy is having endline pollution bin equal to one (i.e., pollution beneath the standard). The F-test reported is for the joint significance of the endline pollution bin dummies. Plant characteristics include dummies for size, use of coal or lignite as fuel, high waste water generated, dye sector, textile sector, and region. Recent regulatory actions include the number of regulatory actions of several types against the plant in the year before endline. Recent pollution readings include dummies for pollution bins at the most recent regulatory inspection before endline. Robust standard errors in parentheses with * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

B Simulation of Targeting Stage

This appendix describes the Monte Carlo simulation of optimal regulatory targeting.

B.1 Set-up

We set $N = 900$ plants and endow plants with exogenous observables $X_j = [X_1 \;\; X_2 \;\; X_3]$ where $X_1, X_2, X_3 \in \{0, 1\}$ and $Pr(X_1 = 1) = 0.25, Pr(X_2 = 1) = Pr(X_3 = 1) = 0.50$. We assume the parameter values in Table B11.

Table B11: Parameters for Targeting Model Simulation

Description	Parameter (1)	Value (2)
Pollution equation	ϕ	$[0.5 \quad 0.5 \quad -.1 \quad -1.5]'$
Inspection equation	β	$[0.3 \quad -.1 \quad 0.25]'$
Maintenance cost	μ_c	2
	σ_c	1
Pollution shocks	σ_1	$\in \{0.31, 0.43, 0.61\}$
	σ_2	$\in \{0.73, 0.66, 0.50\}$
Targeting parameter	ρ	1

The table gives the parameters for the Monte Carlo simulation of the targeting model. Pollution parameters are the coefficients on $[X_j \;\; Run]$ in the pollution equation and inspection parameters the coefficients on X_j. Other parameters are described in Section IV.

In each simulation, we:

1. Draw shocks for each plant (holding the draws fixed across simulations);
2. Construct latent levels of pollution if the plant does not abate;
3. Solve the regulator's targeting problem subject to a budget of $\overline{N} = 1.47$ inspections per plant.

The nine different simulations vary on the two dimensions of penalty function shape and regulatory information.

We use three different penalty functions that yield approximately the same average level of penalties but have different curvature. Figure B1 shows these alternate functions. The horizontal axis in the Figure is the pollution found at a plant on initial inspection, in units of multiples of the regulatory standard. The vertical axis is the expected discounted value of the penalty function, in USD thousands. All penalty functions intersect the same penalty values at pollution readings of

1× and 7.5× the regulatory standard. The three functions differ in their curvature. The dotted (blue) line is the penalty function as estimated in the penalty stage, interpolated between discrete values using a piecewise-cubic hermite interpolating polynomial (pchip), a type of spline that preserves monotonicity between knots. The estimated penalty function shows sharply increasing marginal penalties at the regulatory standard and decreasing marginal penalties beyond that. The solid (black) line is a linear penalty function (constant marginal penalty) and the dashed (red) line is a quartic penalty function (increasing marginal penalty).

Figure B1: Alternate penalty function forms

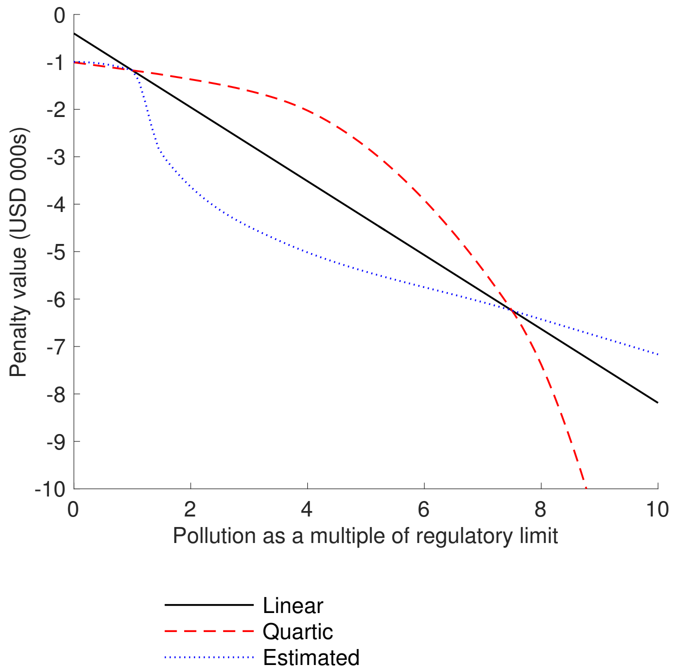

The figure shows alterate functional forms for the penalty function, which gives the expected discounted value of the penalty stage, at the time of an initial inspection, as a function of plant pollution. The value is measured in USD thousands and polluion is measured as a multiple of the regulatory limit; i.e., a plant with pollution equal to two has a reading double the limit. The solid (black) line is a linear penalty function (constant marginal penalty), the dashed (red) line is a quartic penalty function (increasing marginal penalty), and the dotted (blue) line is the penalty function as estimated in the penalty stage. The estimated penalty function is interpolated between discrete values using a piecewise-cubic hermite interpolating polynomial (pchip), a type of spline that preserves monotonicity between knots. All penalty functions are set to intersect the estimated penalty function at pollution values of 1.0 (the standard) and 7.5 (roughly the limit of the pollution values observed).

The second dimension on which the simulations vary is the information available to the regulator. We keep $\sigma_1^2 + \sigma_2^2 = 0.625$ across simulations and vary the fraction of the variance in total pollution that is observable to the regulator, setting σ_1 such that $\sigma_1^2/(\sigma_1^2+\sigma_2^2) \in \{0.15, 0.30, 0.60\}$.

In total there are therefore nine different simulation runs with each combination of regulatory information and penalties.

B.2 Simulation results

This section presents results from simulations of the targeting model to show how regulatory penalties and information shape optimal inspection targeting.

The optimal targeting rule will depend on several aspects of the model. First, since the plant's value of regulation from the penalty stage varies with pollution, the regulator can induce more abatement by allocating inspections to plants with high pollution and therefore high expected penalties. Second, plant reductions in pollution are proportional to the level of pollution, so allocating inspections to higher-polluting plants will have higher yield in abatement, if those plants do abate. These forces both suggest allocating more inspections to plants with higher observable pollution shocks. Third, however, plants have idiosyncratic maintenance costs, and if inspections are very concentrated on the plants with the highest observable pollution shocks, the regulator may miss chances to induce lower-cost plants to *Run* their equipment. The relative strength of these forces—how much the regulator should go after the plants it thinks are the dirtiest—will depend on the accuracy of the regulator's signal of pollution and the shape of the penalty function with respect to measured pollution.

Figure B2 shows the optimal targeting functions, the key endogenous object of interest, from each of these nine simulations. The three panels vary in the share of pollution observable to the regulator, and the three lines within each panel vary in the shape of the penalty function. We plot all the optimal targeting functions for a hypothetical plant for which the argument of the targeting function $X_j'\beta = 0$, aside from the contribution of the observed pollution shocks.

Consider first the variation in targeting functions induced by the regulator's information. The solid (black) line in all figures represents the optimal targeting function, of the form (5), given a linear value function of the penalty stage (constant marginal penalties for higher initial pollution). In Panel A, when the share of variation in pollution observed by the regulator is small, the optimal penalty function is steep. The maximum inspection parameter $\lambda_2 = 28.72$ indicates that a plant with an arbitrarily high u_{1j} shock would receive at most 28.72 inspections per year, and the shift parameter $\lambda_1 = -1.82$ implies that plants would have to draw extremely

Figure B2: Optimal Targeting Rules as Observable Share of Pollution Varies

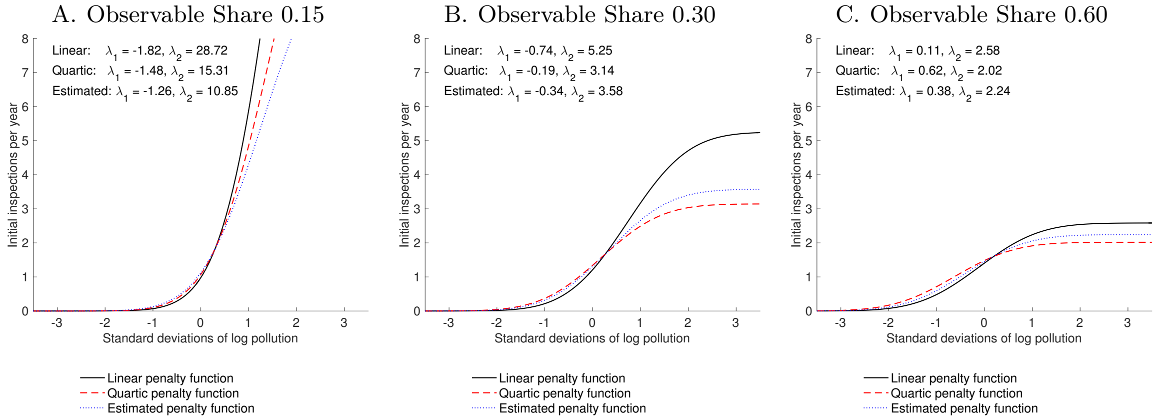

The figure compares the optimal inspection targeting function in Monte Carlo simulations of the targeting stage of the model that vary (i) the share of variation in pollution that is observable by the regulator (across panels) and (ii) the shape of the penalty function (across curves within a panel). The share of variation observable to the regulator is equal to $\sigma_1^2/(\sigma_1^2+\sigma_2^2) \in \{0.15, 0.30, 0.60\}$ across panels. Each curve within a panel shows the optimal targeting rule, giving inspections as a function of the standard deviations of observed pollution, that solves the regulator's problem of minimizing pollution subject to an inspection budget, in a Monte Carlo simulation of the model. The curves within a panel differ in the shape of the regulatory penalty function. The alternate shapes of the penalty function are shown in Figure B1 and explained in the notes for that table.

high shocks to receive these high inspections—a plant with a $3\sigma_1$ shock and $X_j'\beta = 0$ would receive only about 4 initial inspections per year. The regulator knows little about pollution and puts its eggs in one basket by going after the plants observed to be dirtiest aggressively. As the share of pollution observable to the regulator decreases, in Panels B and C, the optimal targeting function (black curve) gets much flatter and shifts leftwards. The regulator, more confident in its signal and therefore that a plant with moderately high observed pollution will *Run*, spreads inspections around to pick up a broader set of plants that may have low abatement costs.

The variation in the shape of the penalty function is not as important as information for the optimal targeting rule, though it does change the rule somewhat. For example, if the regulator has very little information (Panel A), the optimal targeting rule is somewhat more concentrated in high polluting firms under a linear penalty function then under the estimated penalty function, which has decreasing marginal penalties at high pollution levels. This difference makes sense; under a linear penalty function, inspecting plants that are highly polluting has an extra kick, since marginal penalties do not taper off at high levels of pollution. We see a similar ordering for higher levels of information in Panels B and C. Perhaps most subtly, the estimated penalty function (dashed) with increasing, then decreasing marginal penalties, yields less concentrated inspections than the quartic penalty function under low information but more concentrated inspections under high information. This reversal happens because, under low information, the regulator is inspecting mainly the plants with the highest observable shocks, so the region of the penalty function at very high pollution is relatively more important. At this high range the estimated penalty function has decreasing marginal penalties, unlike the quartic penalty function. Under high information, the regulator is inspecting a broader range of plants, so the region of the estimated penalty function near the standard—which has sharply increasing marginal penalties—comes into play, and it is better for the regulator to concentrate inspections slightly more.

The regulator's targeting problem captures the trade-offs involved in setting a targeting rule in an economically rich way. The parsimonious probit link form (5), governed by parameters λ, allows a range of interesting targeting rules, from those that are very steeply increasing, to rules close to linear in pollution, to rules that are nearly flat.

C Estimation

C.1 Penalty stage estimation

Here we provide further details of the preliminary estimation steps involved in forming the penalty stage likelihood.

1 . Actions and plant payoffs in a penalty stage round

In all even rounds, the plant may *Comply* or *Ignore* the regulatory machine. To *Comply*, a plant must pay a constant $\pi_j(a_{jt} = Comply|s_t) = -k$ to install abatement equipment. We assume all plants have costs for installing abatement capital equal to the average value of abatement capital costs observed in our sample, conditional on installation, which is $k =$ \$17,000. To *Ignore* the regulator costs the plant nothing today, but may increase the risk of future regulatory action.

If the machine *Punishes* the plant payoff takes one of two forms depending on the specification. Most simply, we estimate specifications where plant penalties are a constant $h(p_{jt}) = -h$. We also estimate specifications where plant penalties depend on pollution

$$\begin{aligned}\pi_j(a_{Rt} = Punish|s_t) &= -h(p_{jt}) \\ &= -\left(\tau_0 \mathbf{1}\{\overline{p} < p_{jt} < 2\overline{p}\} + \tau_1 \mathbf{1}\{2\overline{p} \leq p_{jt} < 5\overline{p}\} + \tau_2 \mathbf{1}\{5\overline{p} \leq p_{jt}\}\right),\end{aligned}$$

where p_{jt} is the pollution reading and the legally mandated pollution threshold $\overline{p}$. The payoff for punishment is the cost to the plant of temporary closure and any remediation.

If the machine *Inspects* the plant has payoff

$$\begin{aligned}\pi_j(a_{Rt} = Inspect|s_t) &= -b(p_{jt}, a_{j-}) \\ &= -(1 - \mathbf{1}\{a_{j-} = Comply\}) \times \\ &\quad \left(\nu_0 \mathbf{1}\{\overline{p} < p_{jt} < 2\overline{p}\} + \nu_1 \mathbf{1}\{2\overline{p} \leq p_{jt} < 5\overline{p}\} + \nu_2 \mathbf{1}\{5\overline{p} \leq p_{jt}\}\right).\end{aligned}$$

The subscripts in a_{j-} reference the prior action of the plant. If the regulator is to move at turn t, then a_{R-} will be the regulatory machine's prior action at $t-2$; if the plant is to move at t then a_{R-} will have been taken at $t-1$. The payoff for inspection captures the possible disruption posed by inspections and any bribes plants give on being inspected. This function specifies that inspections are costless for plants that have recently complied but, for plants that have not recently complied, inspections incur costs that depend on pollution emissions. We also estimate specifications where inspections are costless for all plants.

Warnings are costless to the plant but continue the stage and obligate the plant to respond. Finally, the machine may *Accept* that the plant is compliant, which costs the plant nothing and ends the penalty stage.

2 . States and state transitions

We specify the common state of the game as comprised of the pollution reading, the last two actions of the regulator and plant, and the penalty stage round:

$$s_t = \{p_{jt}, a_{j-}, a_{R-}, \mathbf{1}\{t > 2\}, \mathbf{1}\{t > 4\}, \mathbf{1}\{t > 6\}\}.$$

The pollution reading p_{jt} is the maximum reading for any pollutant observed in the most recent inspection; if no inspection occurs in round t, then pollution is recalled from the last inspection within the stage. We specify the round as entering with several dummies to allow regulatory actions to respond flexibly to the selection of plants that may occur across rounds.[27]

The state transition after the plant moves is wholly deterministic, because the plant affects only how its own action is recorded in the state: if it chooses to *Comply* today, then $a_{j-} = Comply$ tomorrow. The transition after the regulator moves has a deterministic part, for the machine's action, and a stochastic part, the current plant pollution reading. We use a simple count estimator for the pollution state transition when the machine moves.

$$Pr(p'|p_{jt}, a_{Rt}) = \frac{\sum_{j,c,t} \mathbf{1}\{p_{j,t+1} = p'|p_{jt}, a_{Rt}\}}{\sum_{j,c,t} \mathbf{1}\{p_{jt}, a_{Rt}\}}.$$

The pollution state may transition only if $a_{Rt} = Inspect$. We restrict the pollution transition to depend on past pollution and the machine's move, but not the plant's past moves.[28] When the machine *Accepts* the penalty stage ends and the firm draws a new u_{2jm}.

We treat each chain of interactions between the plant and the regulatory machine as independent, i.e. $u_{2j,m+1}$ is independent of u_{1j} and u_{2jm}. The regulator continues to observe part of

[27]In theory, the whole history of player actions could enter the state. We found that enriching the state in plausible directions, like including further lags, did not help predict regulatory actions.

[28]The count estimator may be biased for low-probability events in finite samples, so that conditioning on more past actions will leave many cells empty (e.g., the probability of pollution transitioning from above 5 times the standard to between 1 and 2 times given that the plant complied and the regulator inspected). Nonetheless, we find the count estimator preferable to smooth alternatives, such as an ordered logit model, because it does not restrict state transition patterns.

pollution, u_{1j}, but conditional on this observation does not, for example, use past readings from the penalty stage to determine targeting.[29]

3 . Regulatory machine conditional choice probabilities

The plant knows the machine action probabilities in each possible future state. We use a multinomial logit model to estimate these conditional probabilities, where:

$$Pr(a_{Rt} = a|s_t) = \frac{\exp\left(q(s_t)'\omega_a\right)}{\sum_{a'} \exp\left(q(s_t)'\omega_a\right)}$$

ω_a is a vector of coefficients for each action and $q(s_{it})$ is a vector of state values: dummies for the possible most recent actions, categorical bins for the observed pollution level p_{it}, and dummies for the stage of the game.

4 . Action-specific values

We calculate action-specific values for the plant using backwards induction. We assume the game is finite and that the regulatory machine will always accept in period $T = 35$, which is well beyond the ultimate round of $t = 19$ actually observed in the data.[30] We further assume the plant does not anticipate any change in future value from actions beyond the current penalty stage. In the plant's problem, which is then finite, we infer action-specific values using the state transitions and choice probabilities, starting at the final round.[31] We use a discount factor of $\delta = 0.991$ between rounds, which has been calibrated, given the average round duration, to match the annual returns on capital for Indian firms found by Banerjee and Duflo (2014).

[29] This assumption is for tractability but is also empirically reasonable. The average time between chains, about five months, is much larger than the average time between actions within a chain, two weeks. Recent pollution readings do not change regulatory targeting of inspections (Appendix Table A10, column 4). Lastly, the regulator has a short memory in practice; of those actions that explicitly cite a prior inspection, 93% of the time the inspection cited is the most recent prior inspection.

[30] Given that the probability of regulatory machine's acceptance in any given round acts like a discount factor, this assumption on the game length is conservative in that late rounds matter very little for plants' expected values.

[31] When $t = T$, then $v_j(a_{Rt}|s_t) = 0$. At $t = T - 1$, the plant's value equals its one-period profit plus an action-specific shock, $v_j(a_{jt}|s_t) = \pi_j(a_{jt}|s_t) + e_j(a_{jt}, s_t)$. The regulator always moves as estimated in the data. At moves $t = T - 3$ and all earlier moves of the plant, the plants action-specific value is found with the empirical analogue to equation (6), where the plant's profit in a given round depends on the parameters θ_{uP}. For the estimation, we restrict the sample to all plant actions taken in round $t = 4$ and after, omitting $t = 2$ on the grounds that we believe the plant often does not have a chance to respond to the regulator in $t = 2$ before another regulatory action is observed (See discussion in Section II).

C.2 Targeting stage estimation

Section V.B. describes the moments used to estimate the targeting stage of the model. This Appendix derives the expressions for these moments in the model. Given the form of the inspection targeting function and the assumed distributions of cost and pollution shocks, most of the model moments have concise analytic forms.

- **Pollution equation**. We use T_j as an instrument for N_j in the pollution equation (1):

$$g_{1j}(\phi) = Z_j'(\log P_j - \phi_0 - X_j'\phi_1 - Run_j\phi_2) \quad (8)$$

  where $Z_j = [1\ X_j'\ T_j]'$.

- **Inspection equation**. The targeting rule allows for an analytic expression of the expected number of inspections, given plant characteristics X_j. We calculate the expected number of initial inspections for a plant with certain observable characteristics as

$$\begin{aligned}
\mathbb{E}[\mathcal{I}(u_{1j}|X_j,T_j,\lambda,\beta,\rho)] &= \int \phi(u_1/\sigma_1)\mathcal{I}(u_{1j}|X_j,T_j,\lambda,\beta,\rho)dF(U_1) \\
&= \int \phi(u_1/\sigma_1)\lambda_2\Phi\left(\frac{\lambda_1 + X_j'\beta_1 + T_j\beta_2 + u_1}{\rho}\right)dF(U_1) \\
&= \lambda_2\Phi\left(\frac{\lambda_1 + X_j'\beta_1 + T_j\beta_2}{\sqrt{\rho^2+\sigma_1^2}}\right),
\end{aligned}$$

  using the distribution of u_1.[32] Therefore, for any X_j, T_j and candidate parameters, we can calculate the expected number of initial inspections by integrating over the observable pollution shocks. With this expression, we form moments

$$g_{2j}(\beta) = Z_j'(\mathbb{E}[\mathcal{I}(u_{1j}|X_j,T_j,\lambda,\beta,\rho)] - I_j). \quad (9)$$

  Similarly, we can calculate the expected value of squared inspections as

$$\begin{aligned}
\mathbb{E}[\mathcal{I}^2(u_{1j}|X_j,T_j,\lambda,\beta,\rho)] &= \int \phi(u_1/\sigma_1)\mathcal{I}^2(u_{1j}|X_j,T_j,\lambda,\beta,\rho)(u_1|X_j,T_j;\beta)dF(U_1) \\
&= \lambda_1^2\int \phi(x)\Phi^2\left(\frac{x-a}{b}\right)dx
\end{aligned}$$

  for $x = u_1/\sigma_1 \sim \mathcal{N}(0,1)$, $a_j = [-(\lambda_1 + X_j'\beta_1 + T_j\beta_2)/\sigma_1]$, $b = \sqrt{\rho^2+\sigma_1^2}$. This integral can

[32] Let $X \sim \mathcal{N}(0,1)$. Then

$$I(a,b) = \int \phi(x)\Phi\left(\frac{x-b}{a}\right)dx = 1-\Phi\left(\frac{b}{\sqrt{a^2+1}}\right) = \Phi\left(\frac{-b}{\sqrt{a^2+1}}\right).$$

Recall that $u_1 \sim \mathcal{N}(0,\sigma_1)$ so $u_1/\sigma_1 \sim \mathcal{N}(0,1)$.

be represented as a bivariate normal cumulative distribution function $\mathcal{N}_2$. If we let

$$Z_1, Z_2 \sim \mathcal{N}\left(\mu = \begin{bmatrix} -a_j \\ -a_j \end{bmatrix}, \Sigma = \begin{bmatrix} \rho^2 + \sigma_1^2 & \sigma_1^2 \\ \sigma_1^2 & \rho^2 + \sigma_1^2 \end{bmatrix}\right)$$

and represent the joint CDF of these variables as $Pr(Z_1 \leq z \cup Z_2 \leq z) = \mathcal{F}(z, \mu, \Sigma)$, then the expected value of inspections squared is

$$\mathbb{E}[\mathcal{I}^2(u_{1j}|X_j, T_j, \lambda, \beta, \rho)] = \lambda_2^2 \mathcal{F}([0\ 0]', \mu, \Sigma).$$

We form moments as

$$g_{3j}(\beta) = Z_{3j}'(\mathbb{E}[\mathcal{I}^2(u_{1j}|X_j, T_j, \lambda, \beta, \rho)] - I_j^2). \tag{10}$$

where $Z_{3j} = [1\ T_j]'$.

- **Abatement cost moments**. We use the observed probabilities of *Run* and expected abatement costs, conditional on running, to form:

$$g_{4j}(\phi, \mu, \sigma) = Pr(Run = 1|\phi, \mu, \sigma) - \mathbf{1}\{c_j > 0\} \tag{11}$$

$$g_{5j}(\phi, \mu, \sigma) = \mathbb{E}[c_j|Run, \phi, \mu, \sigma] - \mathbf{1}\{c_j > 0\}c_j. \tag{12}$$

The probability of running and expected cost, conditional on running, are both functions of the distributions of c_j, u_1 and u_2. There is not a convenient analytic form for the truncated distribution of c_j, from which to form moments, since this would be the expectation of a truncated sum of log-normal and normal components, passed through a non-linear penalty function $V_0(\cdot)$. Therefore, we partly simulate the relevant probability and expectation.

Let $\bar{c}_j = I_j(V_0(P_j) - V_0(P_j^*))$ so that $Run = \mathbf{1}\{c_j < \bar{c}_j\}$. The threshold $\bar{c}_j$ is a function of plant pollution shocks u_1, u_2 and model parameters (including parameters on targeting, to determine I_j). For candidate parameters and a draw of shocks, we can calculate $\bar{c}_j$ for each plant. Then the probability of *Run* in a simulation s is

$$\begin{aligned} Pr^s(Run = 1|\phi, \mu, \sigma) &= Pr(c_j < \bar{c}_j^s|\phi, \mu, \sigma) \\ &= \Phi\left(\frac{\log \bar{c}_j^s - \mu_c}{\sigma_c}\right). \end{aligned}$$

Across S simulation draws for each pollution shock, we calculate

$$Pr(Run = 1|\phi, \mu, \sigma) = \sum_{s=1,\ldots,S} \Phi\left(\frac{\log \bar{c}_j^s - \mu_c}{\sigma_c}\right)/S.$$

We have used simulation over two dimensions of the shock, and an analytic expression for the truncated moment over the third.

For the expected value of maintenance costs, conditional on maintenance, we use the moments of the truncated log-normal distribution.[33] We form the expected cost of maintenance, conditional on running equipment, as

$$\mathbb{E}[c_j|Run, \phi, \mu, \sigma] = \exp(\mu_c + \sigma_c^2/2)\frac{\Phi(-\sigma_c + b_{0j})}{\Phi(b_{0j})},$$

where $b_0 = (\log \overline{c}_j^s - \mu_c)/\sigma_c$. Again, because the value $\overline{c}_j^s$ depends on the simulation draws, we take the expected value of c_j as the mean of this expression over simulation draws.

This simulation is somewhat complicated: plants only decide on whether to *Run* abatement equipment based on their expectation of the regulator's targeting, and this targeting rule depends on parameters. Thus the simulation involves first solving the targeting rule for given parameters, then simulating $\overline{c}_j^s$ for each draw of pollution shocks, and finally calculating the probability of abatement and expected value of abatement over simulation draws.

- **Variance of pollution shocks**. We wish to estimate the components of σ for the standard deviation of the observed and unobserved pollution distributions.

  From the pollution equation, the variance of log pollution is equal to the sum of the variances of the two independent shocks (recall, $\varepsilon_2 = u_1 + u_2$). Therefore

  $$\begin{aligned} g_{6j}(\beta, \phi, \sigma) &= \mathbb{E}[\varepsilon_2^2|\beta, \phi, \sigma] - \hat{\varepsilon}_2^2 \\ &= \sigma_1^2 + \sigma_2^2 - \hat{\varepsilon}_2^2. \end{aligned}$$

  where we can form our empirical estimate of the pollution residual as $\hat{\varepsilon}_2 = P_j - \hat{\phi}_0 - X_j'\hat{\phi}_1 - Run_j\hat{\phi}_2$. This moment identifies the sum of the variances of the observed and unobserved shocks.

- **Covariance of pollution and inspection shocks**. Finally, we are interested to separate the effect of observed and unobserved pollution shocks. The key idea is that only observable

[33] For truncation from above with $x \sim \mathcal{N}(\mu, \sigma)$ and $y = e^x$, these are

$$\begin{aligned} \mathbb{E}[y|y \le b] &= \exp(\mu + \sigma^2/2)\frac{\Phi(-\sigma + b_0)}{\Phi(b_0)} \\ \mathbb{E}[y^2|y \le b] &= \exp(2\mu + 2\sigma^2)\frac{\Phi(-2\sigma + b_0)}{\Phi(b_0)}. \end{aligned}$$

where $b_0 = (\log b - \mu)/\sigma$.

pollution shocks result in more inspections. We form an additional moment

$$g_{7j}(\beta, \phi, \sigma) \quad = \quad \mathbb{E}[\varepsilon_2 \cdot \mathcal{I}|\theta] - \hat{e}_{2j} \times I_j.$$

This moment is related to the covariance of the pollution shock and the level of inspections. Intuitively, if the inspection decision is based upon a part of the pollution error not observed by the econometrician, then the pollution residual and inspections will covary.

We derive a prediction for $\mathbb{E}[\varepsilon_2 \cdot \mathcal{I}|\theta]$ in the model.

$$\begin{aligned}
\mathbb{E}[\varepsilon_2 \cdot \mathcal{I}|\theta] &= \mathbb{E}[(u_1 + u_2) \cdot \mathcal{I}(u_1|\theta)] \\
&= \mathbb{E}[u_1 \cdot \mathcal{I}(u_1|\theta)] + \mathbb{E}[u_2 \cdot \mathcal{I}(u_1|\theta)] \\
&= \mathbb{E}[u_1 \cdot \mathcal{I}(u_1|\theta)].
\end{aligned}$$

Where the third line follows because the pollution shock u_2 is unobserved by the regulator and cannot affect inspection decisions. This is the key idea for identifying targeting; the inspection targeting function depends only on the observable part of the shock.

We proceed by substituting in the targeting function and integrating by parts to yield the desired moment in the model

$$\mathbb{E}[\varepsilon_2 \cdot \mathcal{I}|\theta] \quad = \quad \lambda_2 \frac{\sigma_1}{\rho} \phi \left(\frac{-(\lambda_1 + X_j' \beta_1 + T_j \beta_2)}{\sqrt{\rho^2 + \sigma_1^2}} \right).$$

This moment will depend on the observable characteristics of each plant. Because $\lambda_2 > 0$, $\sigma_1 > 0$ and $\rho > 0$, this correlation is expected to be positive: plants with high pollution shocks are expected to have higher inspections, due to regulatory targeting.

C.3 First-order conditions of regulator's targeting problem

The regulator's objective function given parameters and firm characteristics is

$$\begin{aligned}
\lambda_0^*, \lambda_1^* \quad \in \quad \underset{\lambda_1, \lambda_2}{\arg\min} \sum_{j=1,\ldots,N} \int \int \mathcal{F}\left(\mathcal{I}(u_{1j}|X_j, T_j, \lambda, \beta, \rho)(V_0(P_j) - V_0(\tilde{P}_j))|\phi, \mu_c, \sigma_c\right) \\
\times \tilde{P}(1 - e^{\phi_2}) dF(U_2) dF(U_1) \qquad (13)
\end{aligned}$$

such that

$$\sum_{j=1,\ldots,N} \int \mathcal{I}(u_{1j}|X_j, T_j, \lambda, \beta, \rho) dF(U_1) = N \cdot \bar{I}. \qquad (14)$$

The optimal targeting rule is of the form

$$\mathcal{I}^*(u_{1j}|X_j, T_j, \lambda, \beta, \rho) \quad = \quad \lambda_2^* \Phi \left(\frac{\lambda_1 + X_j' \beta_1 + T_j \beta_2 + u_1}{\rho} \right). \qquad (15)$$

In practice, the integrals over pollution shocks are approximated with draws of u_1, u_2. Therefore we write this objective function as a sum over simulations $s = 1, \ldots, S$. Let $\mathcal{A}(\lambda|\cdot)$ represent the amount of abatement achieved with targeting parameters λ. The cost distribution is log-normal, so

$$\mathcal{A}(\lambda|X_j, T_j, \beta, \phi, \sigma, u_s) = \sum_j \sum_s \Phi\left(\frac{\log[\mathcal{I}(u_{1j}|X_j, T_j, \lambda, \beta, \rho)\Delta_{js}] - \mu_c}{\sigma_c}\right) \times \tilde{P}(1 - e^{\phi_2}),$$

where $\Delta_{js} = V_0(P_j) - V_0(\tilde{P}_j)$ is the reduction in expected penalties from taking the abatement action *Run*. We omit arguments and substitute the form of the targeting function.

$$\begin{aligned}
\mathcal{A}(\lambda) &= \sum_j \sum_s \Phi\left(\frac{\log[\mathcal{I}(u_1)\Delta_{js}] - \mu_c}{\sigma_c}\right) \times \tilde{P}(1 - e^{\phi_2}) \\
&= \sum_j \sum_s \Phi\left(\frac{\log[\lambda_2\Phi((\lambda_1 + \beta_1 X_j + \beta_2 T_j + u_1)/\rho))\Delta_{js}] - \mu_c}{\sigma_c}\right) \times \tilde{P}(1 - e^{\phi_2}).
\end{aligned}$$

Let $Z_{js} = (X'_J\beta_1 + T_j\beta_2 + u_1)/\rho$. Then we have

$$\begin{aligned}
\mathcal{A}(\lambda) &= \sum_j \sum_s \Phi\left(\frac{\log[\lambda_2\Phi(\lambda_1/\rho + Z_{js})\Delta_{js}] - \mu_c}{\sigma_c}\right) \times \tilde{P}(1 - e^{\phi_2}) \\
&= \sum_j \sum_s \Phi\left(\frac{\log\lambda_2 + \log[\Phi(\lambda_1/\rho + Z_{js})] + \log[\Delta_{js}] - \mu_c}{\sigma_c}\right) \times \tilde{P}(1 - e^{\phi_2}).
\end{aligned}$$

Let the argument of the cost distribution be denoted by C_{js}. Taking the derivative with respect to λ_2,

$$\partial\mathcal{A}(\lambda)/\partial\lambda_2 = \sum_j \sum_s \phi(C_{js})\frac{1}{\lambda_2\sigma_c} \times \tilde{P}(1 - e^{\phi_2}).$$

With respect to λ_1,

$$\partial\mathcal{A}(\lambda)/\partial\lambda_2 = \sum_j \sum_s \phi(C_{js})\frac{1}{\sigma_c}\frac{1}{\Phi(\lambda_1/\rho + Z_{js})}\phi(\lambda_1/\rho + Z_{js})\frac{1}{\rho} \times \tilde{P}(1 - e^{\phi_2}).$$

We can write the Lagrangian of the optimal targeting problem. Omitting arguments,

$$\begin{aligned}
\mathcal{L}(\lambda_1, \lambda_2) &= \sum_{j=1} \int\int \mathcal{F}\left(\mathcal{I}_j(u_1)(V_0(P_j) - V_0(\tilde{P}_j))\right) \times \tilde{P}_j(1 - e^{\phi_2})dF(U_1)dF(U_2) \\
&\quad -\gamma\left(\sum_{j\in\mathcal{C}} \int \mathcal{I}_j(u_1)dF(U_1) - N \cdot \overline{I}\right) \\
&= \sum_j \sum_s \Phi\left(\frac{\log[\mathcal{I}(u_{1j}|X_j, T_j, \lambda, \beta, \rho)\Delta_{js}] - \mu_c}{\sigma_c}\right) \times \tilde{P}_j(1 - e^{\phi_2}) \\
&\quad -\gamma\left(\sum_{j\in\mathcal{C}} \lambda_2\Phi\left(\frac{\lambda_1 + X'_j\beta_1 + T_j\beta_2}{\sqrt{\rho^2 + \sigma_1^2}}\right) - N \cdot \overline{I}\right).
\end{aligned}$$

Here we denote by $\mathcal{C}$ the set of plants j in the control group. Thus the optimality constraint is that the targeting rule is optimal in the control group. We then calculate the first-order conditions of the constrained problem.

$$\frac{\partial \mathcal{L}}{\partial \lambda_2} = \frac{\partial \mathcal{A}}{\partial \lambda_2} - \gamma \sum_j \Phi \left(\frac{\lambda_1 + X_j' \beta_1 + T_j \beta_2}{\sqrt{\rho^2 + \sigma_1^2}} \right) = 0 \tag{16}$$

$$\frac{\partial \mathcal{L}}{\partial \lambda_2} = \frac{\partial \mathcal{A}}{\partial \lambda_1} - \gamma \sum_j \lambda_2 \phi \left(\frac{\lambda_1 + X_j' \beta_1 + T_j \beta_2}{\sqrt{\rho^2 + \sigma_1^2}} \right) \frac{1}{\sqrt{\rho^2 + \sigma_1^2}} = 0 \tag{17}$$

$$\frac{\partial \mathcal{L}}{\partial \gamma} = \sum_{j \in \mathcal{C}} \lambda_2 \Phi \left(\frac{\lambda_1 + X_j' \beta_1 + T_j \beta_2}{\sqrt{\rho^2 + \sigma_1^2}} \right) - N \cdot \bar{I} = 0. \tag{18}$$

This optimal targeting parameters λ_1, λ_2 satisfy these first-order conditions with Lagrange multiplier γ representing the shadow value of the inspection budget constraint.

D Robustness and Sensitivity Analysis

This Appendix studies the robustness of the structural estimates. Section D.1 considers the robustness of estimated targeting parameters with respect to the two calibrated parameters σ_c and ρ. Section D.2 considers the sensitivity of a broader set of parameters of interest with respect to variation in the underlying moments in the data, using the measure of Andrews et al. (2017).

D.1 Robustness to Calibrated Parameters

The targeting stage estimation fixed the values of σ_c and ρ. Fixing these values reduces the number of free parameters in the maintenance cost distribution from two to one and in the targeting function from three to two. This sub-section studies how the values of these parameters affect the values of estimated parameters.

Table D12 shows the baseline targeting stage estimates in column 1 and estimates with alternate values of the fixed parameters in columns 2 through 4. The calibrated parameters are shown in the column headers. Columns 2 and 3 change the value of σ_c from the baseline value of $\sigma_c = 0.50$ to $\sigma_c = 0.25$ and $\sigma_c = 1.00$, respectively, and columns 4 and 5 change the value of ρ from the baseline value of $\rho = 0.25$ to $\rho = 0.15$ and $\rho = 0.35$, respectively.

First consider the effect of altering σ_c on the parameter estimates. Since σ_c is a cost parameter, and affects the willingness to *Run* maintenance equipment to reduce pollution, we expect and indeed find that there are very small effects of this parameter on the targeting equation coefficients. There are two main effects of altering σ_c: first, a higher σ_c in column 3, relative to column 1, increases the estimated μ_c and increases the efficacy of abatement (*Run* equipment coefficient). The model accommodates higher dispersion of cost by moving the cost distribution up and increasing the efficacy of abatement, raising both the gross costs and benefits of abatement, so that the model can still match the moments on the share of plants that are willing to *Run* and their cost conditional on *Run*. The changes in estimates are noticeable but not very large; multiplying σ_c by a factor of four, from column 2 to 3, increases the efficacy of abatement by a factor of $-1.07/-0.60 = 1.78$, and both estimates are within one standard error of our baseline estimate in column 1. For comparison, in Table 6, the change in the estimated abatement efficacy from imposing the constraint of optimal targeting is much larger.

Next consider the change from altering ρ on the parameter estimates. Since ρ is an inspection targeting parameter, we expect and find that there are very small effects of changing this parameter (in columns 4 and 5) on the coefficients in the pollution equation (relative to column 1 baseline estimates). Changing ρ, which is the denominator of the argument of the targeting function, has predictably larger effects on the coefficients β and λ in the targeting function. In particular, for a smaller ρ, in column 4, we see smaller β estimates for the targeting equation, and for a larger ρ, in column 4, we see larger β estimates. These changes are roughly but not exactly proportional; for example the inspection treatment coefficient normalized by ρ is $0.162/0.25 = 0.648$ in column 1 and $0.121/0.15 = 0.807$ in column 4. The argument of the targeting functions, for a plant the same observables and u_1 shock, therefore changes with ρ somewhat, and the estimated parameters λ also change to fit the moments in the data. For example, in the column 4 estimates the maximum inspections parameter λ_2 moves down for a smaller ρ and the shift parameter λ_1 becomes less negative, which offset the effects of changes in the β/ρ coefficients on inspections.

These counteracting shifts in estimated parameters will matter to the extent that the targeting function fits the data moments differently with different values of ρ. To get a sense of the net effect of these changes, in Table D13 we give the values of the main targeting moments, expected inspections and expected inspections squared, from the model, calculated under each set of fixed and estimated parameters. The column headers are the same as in Table D12.

Looking across the columns of Table D13 we see that the expected inspections moment ranges from 2.15 in the baseline case to as high as 2.18 and as low as 2.13. The range of expected inspections hardly varies depending on whether the targeting function has a moderate ρ, or, say, a lower ρ with smaller β estimates and shifts in λ, as in column 4. The expected inspections squared moment is only somewhat more variable with a range from 6.95 to 7.36 depending on the fixed value of ρ. Because the targeting moments can be fit about equally well with different values of ρ, offset by changes in the estimated $(\hat{\beta}, \hat{lambda})$, the values of these parameters are not separately well-identified, and estimation runs allowing a free ρ yielded very imprecise targeting parameters.

Finally, consider the effects of both calibrated parameters on the estimated standard deviations of pollution shocks in Table D12, Panel B. The baseline estimate of $\hat{\sigma}_1 = 0.111$. Doubling or halving σ_c has nearly no effect on the estimated σ_1 ($\hat{\sigma}_1 = 0.110$ and $\hat{\sigma}_1 = 0.117$, respectively) and decreasing and increasing ρ also has small effects ($\hat{\sigma}_1 = 0.093$ and $\hat{\sigma}_1 = 0.112$, respectively). All these changes are within one standard error of the original estimate and typically far smaller. The effects of the fixed parameters on the standard deviations of unobserved pollution σ_2 are also small.

We conclude that the interpretation of the parameter estimates is robust to variations in the fixed parameters σ_c and ρ. In particular, the significance of direct changes in the ρ parameter for targeting appear to be offset by changes in the estimated λ and β parameters. The fixed parameters do not affect the key finding that the regulator has little information on pollution.

D.2 Sensitivity of Parameter Estimates to Moments

1 . Sensitivity Matrix: Definition

We estimate the targeting function through the generalized method of moments with a mix of analytic and simulated moments (Section V.B. and Appendix C.2). Andrews et al. (2017) defines the sensitivity matrix Λ for any estimator $\hat{\theta}$ that minimizes a criterion function $\hat{g}(\theta)'\hat{W}\hat{g}(\theta)$, where $\hat{g}(\theta)$ is a vector of moments or other statistics and $\hat{W}$ is a weight matrix.

Assume that $\sqrt{n}\hat{g}(\theta_0)$ converges in distribution to $\tilde{g}$ such that $\mathbb{E}[\tilde{g}] = 0$ under the model. For alternative specifications of the model this may not be the case. Andrews et al. (2017) define local perturbations of the maintained model and show that, under these perturbations, $\sqrt{n}(\hat{\theta} - \theta_0)$ converges in distribution to a random variable $\tilde{\theta}$ such that $\tilde{\theta} = \Lambda\tilde{g}$. Then the estimator $\hat{\theta}$ has first-order asymptotic bias:

$$\mathbb{E}[\tilde{\theta}] = \Lambda\mathbb{E}(\tilde{g}),$$

where $\Lambda = -(G'WG)^{-1}G'W$, W is the probability limit of the weight matrix $\hat{W}$, and G is the Jacobian of the probability limit of $\hat{g}(\theta)$ at θ_0. We can estimate Λ using the standard plug-in estimates of W and G at little extra computational cost.

Because the moments are in different units, we scale Λ so that it can be read as the effect of a one-standard-deviation violation of the given moment condition on the asymptotic bias of the given parameter. The value of entry Λ_{jk} is therefore measured in the units of parameter k per standard deviation of moment j.

Sensitivity has two equivalent interpretations. Formally, sensitivity is defined as the asymptotic bias of the estimator under local misspecification. One can also think of sensitivity as approximating how a change in one data moment, such as would be generated by an alternative model, would locally affect the parameter estimates.

D.3 Sensitivity of Selected Parameters

Table D14 presents sensitivities for selected parameters (columns) with respect to the estimation moments. Each column of the table gives the sensitivity of the column parameter to the estimation moments listed down the rows. Table D15 condenses this information into the four moments to which each main parameter of interest is most sensitive.

1 . Sensitivity of inspection treatment coefficient

To illustrate the intuition for the sensitivity measure, consider the sensitivity of the inspection treatment coefficient in the targeting equation (Table D14, column 1). The sensitivity of the treatment coefficient to the product of the mean of inspections and the treatment dummy is 1.89 (Row Inspection mean $\times$ T). This sensitivity means that, if the product of inspections and the treatment dummy were higher by one standard deviation, the estimated inspection coefficient β_2 would increase by 1.89 (within the argument of the targeting function; the marginal effect of this change on mean inspections therefore depends on covariates). In this simple case, sensitivity is easy to understand. The higher the mean inspections for treated firms, relative to control, the higher the estimated treatment effect. Table D15 shows that the four most influential moments for β_2 are (i) the mean of the product of treatment with inspections squared, (ii) the mean of the product of treatment with inspections, (iii) mean inspections squared and (iv) mean inspections.

We now turn to consider the sensitivity of several parameters of interest to the estimation moments. We focus on the parameters of regulatory information, abatement cost and the efficacy of pollution abatement

2 . Sensitivity of inspection targeting parameters λ_1 and λ_2

The inspection targeting equation relates unobserved (by the econometrician) pollution shocks and plant observables to inspections. A higher level parameter λ_2 gives the maximum number of inspections a plant may receive and a higher shift parameter λ_1 increases the argument of the targeting function, so that all plants receive a higher level of inspections.

Table D15 shows that interactions of the treatment with the inspection distribution moments are influential for the targeting parameter estimates. In the λ_1 row, the column (2) sensitivity of -2.99 means that a one-standard-deviation higher mean product of inspections and treatment, conditional on other moments, would decrease the inspection shift parameter λ_1 by 2.99. If squared inspections in the treatment group were relatively higher—without changing mean inspections—the estimated inspection shift parameter λ_1 would decrease. This shift downwards would allow the model, with a steep targeting function, to match a *relatively* higher volatility of inspections in the treatment, since plants in the treatment group are more likely to be shifted out onto the steep part of the targeting function. The targeting parameters are also sensitive to other interactions of the treatment with the inspection distribution.

Table D12: Robustness of Targeting Estimates to Calibrated Parameters

$\sigma_c =$	0.50	0.25	1.00	0.50	0.50
$\rho =$	0.25	0.25	0.25	0.15	0.35
	(1)	(2)	(3)	(4)	(5)
Panel A. Targeting and Pollution Equations					
Pollution Equation					
Run equipment (=1)	-0.742	-0.604	-1.073	-0.703	-0.786
	(0.307)	(0.233)	(0.550)	(0.282)	(0.333)
Audit treatment	-0.102	-0.097	-0.111	-0.106	-0.099
	(0.085)	(0.084)	(0.089)	(0.085)	(0.086)
Audit treatment × inspection treatment	0.066	0.059	0.082	0.069	0.069
	(0.108)	(0.107)	(0.114)	(0.107)	(0.109)
Audit sample	0.613	0.607	0.624	0.618	0.601
	(0.137)	(0.135)	(0.143)	(0.137)	(0.138)
Region: Ahmedabad	-0.201	-0.186	-0.232	-0.214	-0.183
	(0.132)	(0.130)	(0.138)	(0.131)	(0.132)
Region: Surat	-0.371	-0.345	-0.426	-0.382	-0.351
	(0.164)	(0.161)	(0.177)	(0.163)	(0.166)
Constant	-0.004	-0.034	0.067	0.002	-0.010
	(0.103)	(0.096)	(0.129)	(0.100)	(0.105)
Targeting Equation					
Inspection targeting shift parameter (λ_1)	-0.219	-0.220	-0.203	-0.075	-0.457
	(0.066)	(0.066)	(0.062)	(0.034)	(0.137)
Inspection targeting level parameter (λ_2)	10.043	10.162	9.304	6.951	18.445
	(3.124)	(3.164)	(2.598)	(1.233)	(11.751)
Inspection treatment	0.162	0.161	0.168	0.121	0.182
	(0.025)	(0.025)	(0.025)	(0.017)	(0.036)
Audit treatment	-0.005	-0.005	-0.007	-0.006	-0.003
	(0.017)	(0.017)	(0.018)	(0.013)	(0.020)
Audit treatment × inspection treatment	0.016	0.015	0.018	0.015	0.012
	(0.021)	(0.020)	(0.022)	(0.016)	(0.023)
Audit sample	0.095	0.095	0.098	0.070	0.110
	(0.024)	(0.024)	(0.025)	(0.018)	(0.030)
Region: Ahmedabad	-0.221	-0.222	-0.232	-0.178	-0.237
	(0.044)	(0.044)	(0.045)	(0.033)	(0.057)
Region: Surat	-0.178	-0.179	-0.186	-0.144	-0.193
	(0.040)	(0.040)	(0.041)	(0.030)	(0.050)
Panel B. Distributions of Pollution and Maintenance Cost Shocks					
Standard deviation of observed pollution shock (σ_1)	0.111	0.110	0.117	0.093	0.112
	(0.022)	(0.022)	(0.023)	(0.016)	(0.028)
Standard deviation of unobserved pollution shock (σ_2)	0.866	0.855	0.899	0.862	0.871
	(0.042)	(0.037)	(0.073)	(0.041)	(0.045)
Mean of log maintenance cost (μ_c)	1.859	1.637	2.350	1.852	1.870
	(0.316)	(0.317)	(0.337)	(0.315)	(0.315)

The table reports estimates of the targeting stage of the model under alternate values of the calibrated parameters σ_c and ρ. Each column shows one set of estimates of the unconstrained targeting model, i.e. without imposing that the regulator's inspection rule is optimal. We use $S = 200$ simulations for each set of estimates. Column 1 shows the baseline estimates (These estimates differ very slightly from those reported in Table 10 of the paper because Table 10 uses $S = 5000$ simulations).

Table D13: Robustness of Expected Inspections to Calibration

$\sigma_c =$	0.50	0.25	1.00	0.50	0.50
$\rho =$	0.25	0.25	0.25	0.15	0.35
	(1)	(2)	(3)	(4)	(5)
$\mathbb{E}[Inspections]$	2.147	2.149	2.150	2.177	2.126
$\mathbb{E}[Inspections^2]$	7.111	7.136	7.131	7.358	6.949

The table reports how the model predictions for moments of expected inspections and expected inspections squared change depending on the value of the calibrated parameters σ_c and ρ. The rows show the values of the two moments and the columns show the predictions of the model at the estimated parameters with each set of calibrated parameters (shown in the column headers). The values of the expected inspections and expected inspections squared moments in the data are 2.20 and 7.53, respectively.

Table D14: Sensitivities for Selected Parameters

Moment	β_2 (1)	ϕ_2 (2)	λ_1 (3)	λ_2 (4)	ϕ_0 (5)	σ_1 (6)	σ_2 (7)	μ_c (8)
Pollution Resid.$\times T$	-0.360	3.114	-0.938	41.689	0.758	-0.362	0.444	1.579
Pollution Resid.	-0.873	-1.858	-2.406	105.770	3.901	-0.927	-0.119	-1.800
Pollution Resid.$\times X_1$	0.026	0.819	0.087	-3.703	0.272	0.033	0.106	0.402
Pollution Resid.$\times X_2$	-0.047	-1.464	-0.151	6.477	-0.458	-0.058	-0.186	-0.817
Pollution Resid.$\times X_3$	-0.185	0.081	-0.505	22.258	-0.055	-0.195	0.045	0.080
Pollution Resid.$\times X_4$	0.439	-0.206	1.195	-52.658	-2.651	0.461	-0.081	0.160
Pollution Resid.$\times X_5$	0.383	-0.015	1.042	-45.914	-2.829	0.402	-0.047	0.216
Inspection Mean$\times T$	1.888	-0.411	1.779	-112.948	-0.357	0.688	-0.075	0.254
Inspection Mean	-1.210	0.223	0.292	44.942	0.402	-0.583	0.027	0.223
Inspection Mean$\times X_1$	0.297	0.168	-0.061	-4.052	0.013	0.029	0.017	0.063
Inspection Mean$\times X_2$	-0.335	-0.133	0.102	3.278	-0.007	-0.027	-0.013	-0.074
Inspection Mean$\times X_3$	0.128	-0.080	0.260	-12.134	-0.004	0.033	-0.014	0.278
Inspection Mean$\times X_4$	-0.182	-0.283	-0.622	16.114	-0.083	-0.017	-0.020	-0.312
Inspection Mean$\times X_5$	-0.172	0.023	-0.644	14.989	-0.031	-0.014	0.011	-0.191
Insp. Squared Mean$\times T$	-2.085	1.746	-2.992	171.161	0.880	-1.205	0.220	0.336
Insp. Squared Mean	1.469	-2.163	1.283	-95.999	-1.003	1.044	-0.232	-0.742
Prob(Run)	-0.005	1.026	-0.014	0.603	0.172	-0.005	0.107	-1.356
$E[c_j \vert Run]$	-0.049	10.317	-0.138	6.066	1.729	-0.053	1.072	7.256
Var(Pollution)	-0.000	0.090	-0.001	0.053	0.015	-0.000	0.886	0.067
Cov(Pollution, Insp.)	0.872	-0.731	2.376	-104.703	0.185	0.917	-0.255	0.806

Note: This table presents values from the estimated sensitivity matrix Λ, scaled so that the entries can be read as the effect of a one-standard-deviation violation of the given moment condition on the asymptotic bias of the given parameter. For the column of corresponding to a given parameter, we can interpret the entries as the sensitivity of the estimated parameter to beliefs about the degree of misspecification of each moment, expressed in standard deviations. Alternatively, we can interpret the estimated sensitivities as a measure of how estimates will respond to changes in the underlying data. While sensitivity is computed with respect to the complete set of variables, the table shows only a subset of particular interest.

Table D15: Most Important Moments for Selected Parameters

Variable	Most Influential Moment	Value	2nd Most Influential Moment	Value	3rd Most Influential Moment	Value	4th Most Influential Moment	Value
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
insp. treatment (β_2)	Insp. Squared Mean$\times T$	-2.085	Insp. Mean$\times T$	1.888	Insp. Squared Mean	1.469	Inspection Mean	-1.210
Run (ϕ_2)	$E[c_j \| Run]$	10.317	Pollution Resid.$\times T$	3.114	Insp. Squared Mean	-2.163	Pollution Resid.	-1.858
Inspection Shift (λ_1)	Insp. Squared Mean$\times T$	-2.992	Pollution Resid.	-2.406	Cov(Pollution, Inspection)	2.376	Insp. Mean$\times T$	1.779
Inspection Level (λ_2)	Insp. Squared Mean$\times T$	171.161	Insp. Mean$\times T$	-112.948	Pollution Resid.	105.770	Cov(Pollution, Inspection)	-104.703
constant (ϕ_0)	Pollution Resid.	3.901	Pollution Resid.$\times X_5$	-2.829	Pollution Resid.$\times X_4$	-2.651	$E[c_j \| Run]$	1.729
SD Observed Shock σ_1	Insp. Squared Mean$\times T$	-1.205	Insp. Squared Mean	1.044	Pollution Resid.	-0.927	Cov(Pollution, Inspection)	0.917
SD Unobserved Shock (σ_2)	$E[c_j \| Run]$	1.072	Var(Pollution)	0.886	Pollution Resid.$\times T$	0.444	Cov(Pollution, Inspection)	-0.255
Mean Maintenance Cost (μ_c)	$E[c_j \| Run]$	7.256	Pollution Resid.	-1.800	Pollution Resid.$\times T$	1.579	Prob(Run)	-1.356

Note: Table 2 presents the names and values of the four moments with the largest absolute value of sensitivity for each parameter. The sensitivities are scaled so that the entries can be read as the effect of a one-standard-deviation violation of the given moment condition on the asymptotic bias of the given parameter. The magnitudes are thus comparable across moments for a given parameter, and can be interpreted as the degree of influence of a moment on the estimated parameter.

The units of sensitivity are not comparable across parameters, since they are measured in the units of each parameter. The sensitivities for λ_2 are so large because λ_2 gives the maximum number of inspections for an arbitrarily large observable pollution shock. Given the estimated shape of the targeting function, this maximum is not reached in the sample. Movements in λ_2 have a lesser effect on the changes in inspections for plants with in-sample covariates and more likely shocks.

The parameters λ_1 and λ_2 are sensitive to many of the same moments, but in opposite directions. The sensitivities of these parameters have different signs for seventeen of twenty moments (Table D14) and for all of the four most influential moments (Table D15). In addition to the moments of the inspection distribution, the inspection shift parameter is sensitive to the pollution residual and the covariance of the pollution and inspection residuals (Table D15, columns 4 and 6). This relationship arises since λ_1 is the constant in the argument of the targeting function, and u_1, the observed pollution shock, enters the same argument. The level of λ_1 must adjust to match the observed level and dispersion of inspections given that the observed pollution shock has mean zero. To put it another way, if the pollution residual increased and the inspection shift parameter did not change, the model would over-predict inspections on average.

3 . Sensitivity of abatement efficacy ϕ_2

The sensitivities show that several moments affect the estimated effect of running abatement equipment on pollution, ϕ_2. The most obvious *ex ante* are that the ϕ_2 estimate is sensitive to the pollution residual times the treatment and the pollution residual. If treatment pollution were higher by one standard deviation, then $\hat{\phi}_2$ would increase from -0.71 to -0.51, indicating a decline in abatement efficacy of 11 percentage points (efficacy being $1 - \exp(\phi_2)$). If the treatment had reduced pollution less than observed, the model would infer that abatement was less effective.

The efficacy of abatement depends on the pollution equation, as it would exclusively in a single-equation model, but also the cost of maintenance. Table D15 shows that Pollution Resid.$\times T$ is only the second most influential moment for ϕ_2, after the moment giving the abatement maintenance cost conditional on running equipment. If the mean maintenance cost conditional on running increased by 100 dollars, the $\hat{\phi}_2$ coefficient would increase from -0.71 to -0.53

(Panel ϕ_2, moment $\mathbb{E}[c_j|Run]$).

4 . Variances of observed and unobserved pollution shocks, σ_1 and σ_2

A goal of the model is to understand what the regulator knows about pollution. Sensitivity analysis can tell us how the model separates the pollution shock into an observed component with standard deviation σ_1 and an unobserved component with standard deviation σ_2.

Table D15, row σ_2 shows that the unobserved component of pollution is most sensitive to maintenance cost, the variance of pollution and the pollution residual interacted with treatment, in that order. Directly, if the variance of observed log pollution were higher then the unobserved pollution shock would be estimated to be higher also (column 4). More subtly, if expected maintenance cost is higher conditional on running, for example, then it must be that there is greater variance in unobserved pollution, in order to induce plants to be willing to run their costly equipment (given a fixed inspection targeting function, set by parameters λ).

The moments important for the observed component σ_1 are notably distinct from those that are important for σ_2. The moment to which σ_1 is most sensitive is not a pollution moment but the product of squared inspections and treatment. Since the observed pollution shock enters the targeting function, and treatment plants have a higher targeting function argument (i.e., higher inspections), the observable pollution shock is estimated to be larger if the dispersion of inspections is greater for treatment plants (column 2) and for control plants (column 4). The squared inspections moments are the uncentered second moments of the targeting function and so capture how much variation in pollution there is on which the regulator can target. If the covariance of pollution and inspection residuals is higher this increases the observable part of the pollution shock (row σ_1, column 8) but decreases the unobservable part (row σ_2, column 8). The model can therefore separate the two pollution shocks since only the observable shock influences the distribution of inspections, through targeting, and since they have effects of opposite sign on the covariance between pollution and inspections.