Econometrica, Vol. 86, No. 6 (November, 2018), 2123–2160

THE VALUE OF REGULATORY DISCRETION: ESTIMATES FROM ENVIRONMENTAL INSPECTIONS IN INDIA

ESTHER DUFLO
Department of Economics, Massachusetts Institute of Technology

MICHAEL GREENSTONE
Department of Economics, University of Chicago

ROHINI PANDE
Kennedy School, Harvard University

NICHOLAS RYAN
Department of Economics, Yale University

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 the 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.

KEYWORDS: Environmental regulation, rules versus discretion, regulatory inspections, development and pollution, industrialization.

1. 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, Ebenstein, Greenstone, and Li (2013), Greenstone, Nilekani, Pande, Ryan, Sudarshan,

Esther Duflo: eduflo@mit.edu
Michael Greenstone: mgreenst@uchicago.edu
Rohini Pande: rohini_pande@harvard.edu
Nicholas Ryan: nicholas.ryan@yale.edu

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 thank 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.

 https://doi.org/10.3982/ECTA12876

and Sugathan (2014), Ebenstein, Fan, Greenstone, He, and Zhou (2017)) and labor productivity (Graff-Zivin and Neidell (2012), Chang, Zivin, Gross, and Neidell (2016), Adhvaryu, Kala, and Nyshadham (2016), He, Liu, and Salvo (2016)). High pollution persists despite strict emission 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 due to 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 due to a lack of standards, as there are strict maximum limits on air and water emissions from industrial plants. Neither 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 the 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 2 years. All sample plants came from the highest category of pollution potential. The inspection treatment, assigned to half of plants, was cross-randomized with an audit reform experiment in the subset of plants that were eligible for environmental audits (Duflo, Greenstone, Pande, and Ryan (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 the control groups.

[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)).

[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.

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 5 years (from 2 years before the experiment through 1 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 end-line 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 followup. 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 removal of regulatory discretion over which plants to inspect. Data on status quo inspections and the process of regulatory sanctions, after an inspection, show 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 chooses which plants 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 noncompliant 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.

[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.

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 such as a mandated plant closure, and informal costs such as 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 1 inspection per year, while plants expected to be the dirtiest may receive 10. 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, Gentzkow, and Shapiro (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 2 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 \$2000. 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 \$6000—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 15% 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 30% 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)).

This 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 indepen-

[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

dently 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.

2. CONTEXT AND EXPERIMENTAL DESIGN

2.1. *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 responsible 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% of control

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, List, and Syverson (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, Burgess, and Redding (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, Hansen, Olken, Potapov, and Sieber (2012), Jia (2014)).

[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 50–100m ($1–2 million),

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 followups 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 toward 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.

2.2. *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 3455 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 3 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; 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 inspections 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]

---

and large, above INR 100m ($2 million) (throughout, we use an exchange rate of $1U.S. = 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.

[9]Toward the end of the inspection treatment, in the month prior to the end-line 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 S.VIII (Duflo, Greenstone, Pande, and Ryan (2018))).

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 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 followup 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 followup behavior in response, which our experimental estimates will not capture. Identical followup 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).

2.3. *Data*

The paper uses two sources of data: an end-line plant survey and GPCB administrative records. The end-line survey was conducted between April and July 2011 (the experiment ended in May 2011) by independent agencies, mainly engineering departments of local universities, supervised by the research organization Abdul Latif Jameel Poverty Action Lab—South Asia (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 S.V and S.VI).

The second source of data is 9624 GPCB documents on its interactions with plants. We categorize these documents by (a) whether they record an action of the regulator or a

[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, participated in both treatment and control inspections.

plant and (b) the type of action they record. Figure 4 shows the actions of the regulator and plant and Appendix Table S.I 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 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* means to accept that the plant is in compliance and is documented by the regulator revoking a prior action in writing 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) 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.

The resulting chains of interactions between the regulator and plants give a picture of the probabilities of punishment and plant compliance. Table I summarizes the structure of the chained interactions between the regulator and plants across rounds.[12] Columns 1 through 6 give the frequency of actions of the regulator or plant in that round, from regulatory records, and column 7 gives the total number of observations in that round. The stage begins with a regulatory inspection. The plant and the regulator then alternate moves until the regulator decides to *Accept* the plant's compliance.

The rapidly descending numbers of observations in column 7 shows that the regulator seldom penalizes plants: 87% of chains end after a single inspection, with the regulator accepting in the third round. Over 900 chains continue beyond that stage, and a handful go on for a dozen rounds or more as violating plants are inspected and punished. If the regulator does not immediately accept the state of the plant, then it is initially more likely to issue a warning: in round three, 9.5% of actions are warnings against 2.2% that are punishments. Thereafter, the regulator is increasingly more likely to punish. Conditional on reaching a given round the probability of punishment rises monotonically with every round from 2.2% in round 3 to 18.1% in round 9 before turning downwards, in late rounds that are seldom observed. Initial plant compliance is low but rises monotonically, with probabilities of 0.4, 7.2, 8.9, 16.5 and 17.5 percent over the second through tenth rounds, before levelling off.

[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 to be well out of compliance.

[12]Online Appendix Table B4 gives an example of an extended regulator-plant interaction.

TABLE I

STRUCTURE OF PENALTY STAGE ACTIONS[a]

	Regulatory Action				Plant Action			
Round	Inspect (1)	Warn (2)	Punish (3)	Accept (4)	Ignore (5)	Comply (6)	N (7)	% Left (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	25,217	

[a]The table reports actions taken by the regulatory machine and by the firm in the penalty stage using administrative data. Figure 5 defines actions and their payoffs and Table II maps them to regulatory documents. Each of columns 1–6 gives the probability, within that row, of the party moving at that round when 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.

2.4. *Randomization Balance Check*

Plant characteristics and past regulatory interactions such as inspections, pollution readings, and citations are balanced by treatment assignment. Appendix Table S.IV presents a randomization check. Of 18 baseline measures reported, there is a significant difference between the treatment and the control groups at the 10% 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% of plants had any pollution reading collected in the year prior to the experiment, and 34% of plants had a pollution reading above the limit (fully 85% of those with a reading taken). Many plants (22% in the control group) were cited for violations. More forcefully, 24% of control plants were mandated to install abatement equipment,[13] 7.5% were ordered temporarily closed, 2% had to post a bank guarantee (performance bond), and 1% had utilities cut off.

[13]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)).

3. 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.

3.1. *Regulatory Action*

Table II 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, the second and third columns report the means for control and treatment plants, while the last column reports the coefficient on the inspection treatment dummy from a regression of each outcome on treatment and dummies for 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 were 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, which 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 followup inspections when a violation is found. On net, discretionary inspections were neither crowded-out nor crowded-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 2-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.

TABLE II

REGULATORY INTERACTIONS DURING EXPERIMENT[a]

	Control	Treatment	Difference
Panel A. Inspections by Treatment Status			
Number of 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	

[a]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 2 years. The second and third columns show means with standard deviations given in brackets. The third column shows the coefficient from regressions of each variable on treatment, where each regression includes region fixed effects and a control for the audit sample. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

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 followup directly in Appendix Table S.VII, 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 followup to an inspection is the same for treatment and control plants, conditional on the regulator's own information.

3.2. *Plant Abatement Costs, Pollution, and Compliance*

Table III, panel A presents estimates of treatment effects on plant abatement costs. We use end-line survey descriptions of abatement expenditures to separate abatement costs

TABLE III

END-LINE POLLUTION AND COMPLIANCE ON TREATMENTS[a]

	(1)	(2)	(3)	(4)
	Panel A. Plant-Level Costs			
	Capital Costs		Maintenance Costs	
	(USD × 10^3)	Any (=1)	(USD × 10^3)	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	

[a]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 thousands of U.S. dollars (USD). 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% 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 are given 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 end-line 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 are given in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

for capital and maintenance.[14] Abatement capital expenditures are observable by the regulator and are 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% 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 was incurred (columns 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 III, 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.[15] 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.[16] 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.[17])

[14] 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.

[15] 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.

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

[17] To test the robustness of this compliance effect, Appendix Table S.9 reports placebo checks where compliance is coded to occur at various multiples of the real standard. The effect of inspection treatment on compliance is statistically significant only at the true standard.

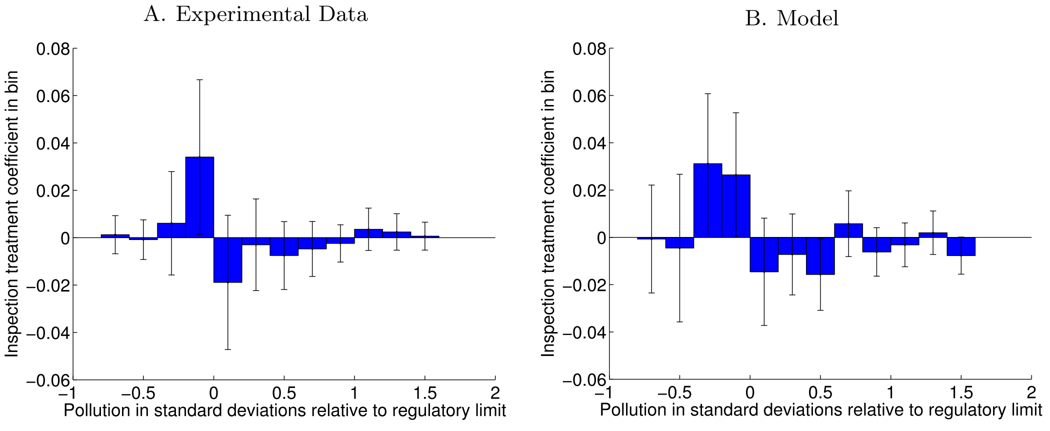

FIGURE 1.—The effect of treatment on pollution distribution. The figures report coefficients on the inspection treatment assignment from regressions of dummies for a pollution reading being in a given bin, relative to the regulatory standard, on inspection treatment, audit treatment, inspection × audit treatment, a dummy for being audit-eligible, and region fixed effects. Part (A) reports coefficients from such regressions on the experimental data and (B) reports coefficients from the same regressions run on model-generated data using the constrained model estimates of Table VI. 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.

Compliance can increase without a large reduction in average pollution if plants near the standard are the most responsive to the inspection treatment. Figure 1(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 III, 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.

3.3. *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 2(A) and (B) demonstrates that the treatment did not appreciably increase the number of plants subject to five or more inspections in a year, which 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% in the control to 13% in the treatment group.

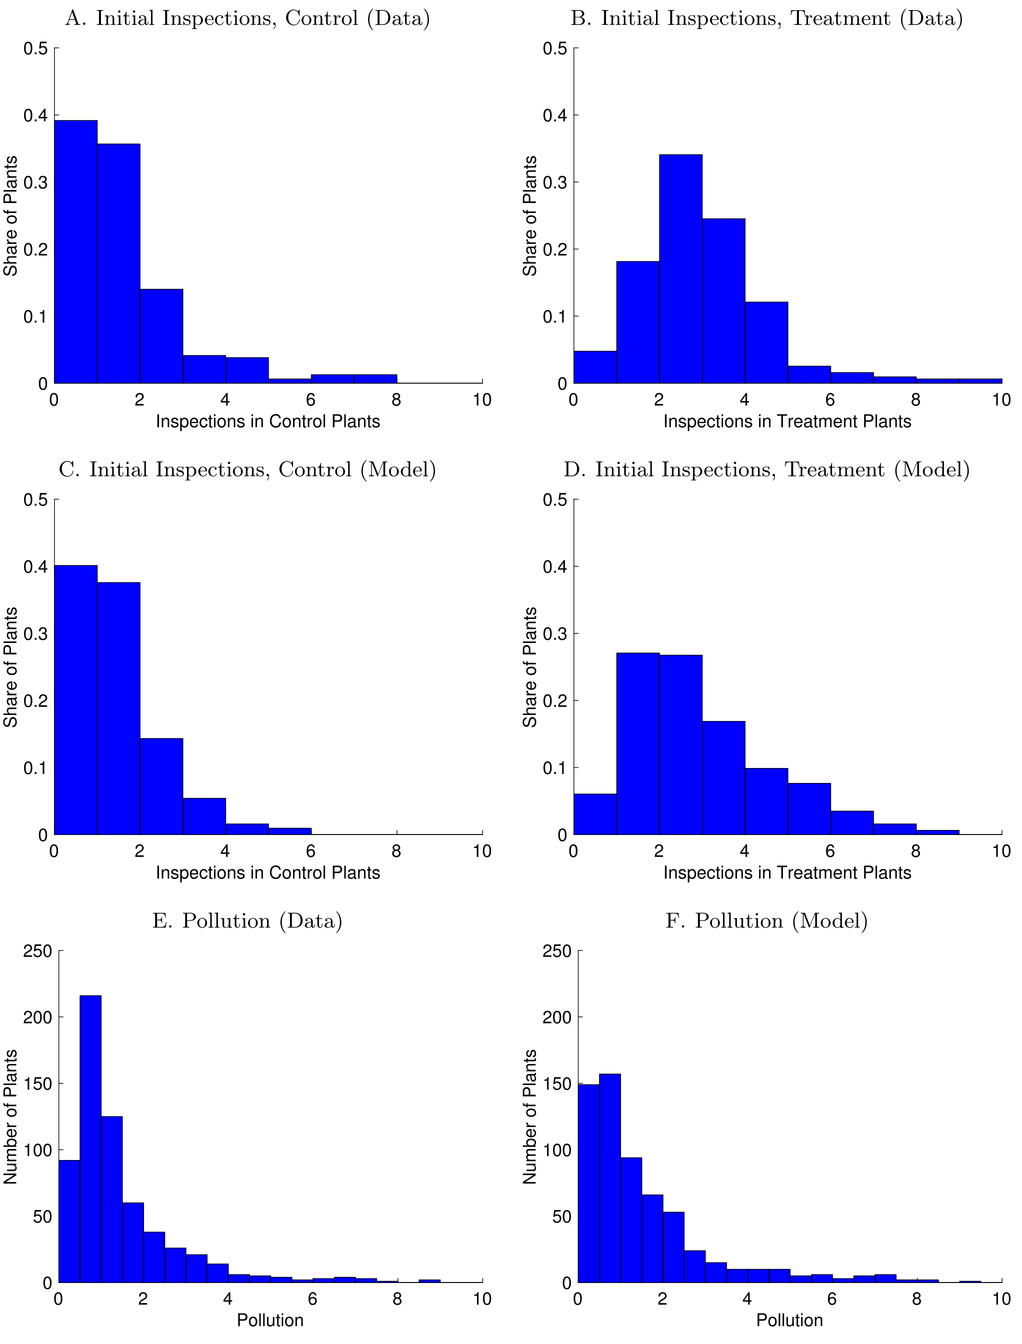

FIGURE 2.—Model fit to inspections and pollution. The figure compares the distributions of inspections and pollution in the model to those in the experimental data. Panels (A)–(D) show the distributions of the annual inspection rate (i.e., inspections per year). Inspections include only initial inspections and not followups. 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.

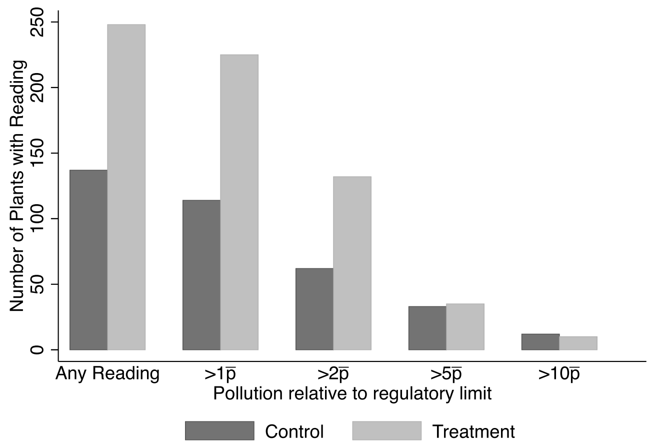

FIGURE 3.—Regulatory targeting of extreme polluters. The figure shows the number of plants with pollution readings either taken or that fall in various bins, relative to the regulatory standard, during the first year of the intervention for the control and treatment groups, respectively. The first pair 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 ($>10\overline{p}$).

Second, despite the additional inspections, Figure 3 reveals that the treatment did not increase the number of extreme violators, that is, plants with pollution readings 5 or 10 times the standard.[18] 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 end-line survey pollution readings, in the control group, predict future regulatory inspections conditional on plant observables and the regulator's own past readings (Appendix Table S.X). Since the regulator did not see the end-line 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.

4. 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

[18]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.

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.

Stage 1. *Targeting*

(i) The regulator chooses an inspection targeting rule to minimize plant pollution subject to a budget of inspections.

(ii) Plants choose whether to run their abatement equipment, given their abatement cost, known level of pollution, and the regulator's targeting and penalty rules.

(iii) The regulator observes a part of plant pollution and inspects plants by applying the targeting rule (i) to this signal, yielding a pollution reading from the inspection.

Stage 2. *Penalty*

(i) The regulator acts as a *regulatory machine*, following exogenous rules for followup and punishment based on pollution measured in inspections and plant actions.

(ii) Plants face a single-agent dynamic problem: they play against the regulatory machine and decide when to comply versus when to 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.

4.1. *Targeting Stage*

4.1.1. *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 is the share of the residual variance in pollution that is due to σ_1, the better is 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.[19] Running abatement equipment reduces pollution proportionally to its latent level, such that $\log P_{jm} = \log \tilde{P}_{jm} + \phi_2 Run$, 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.[20]

4.1.2. *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 5.1 as the expected discounted value to the plant of all regulatory actions in the penalty stage, including followup 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

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

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

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

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

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

[20] Air pollution control devices like filters, electrostatic precipitators, cyclones, and scrubbers are commonly installed in industrial plants in both India and developed countries. The U.S. Environmental Protection Agency (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.

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 $j = 1, \dots, N$ to yield total abatement.[21]

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,

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

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 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 nonparametric targeting rule, and thereby makes it possible to constrain estimates by imposing the optimality of targeting.

4.2. *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 $t = 2, 3, \dots$, 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

[21]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 the plant 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 and (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 is of low marginal environmental value.

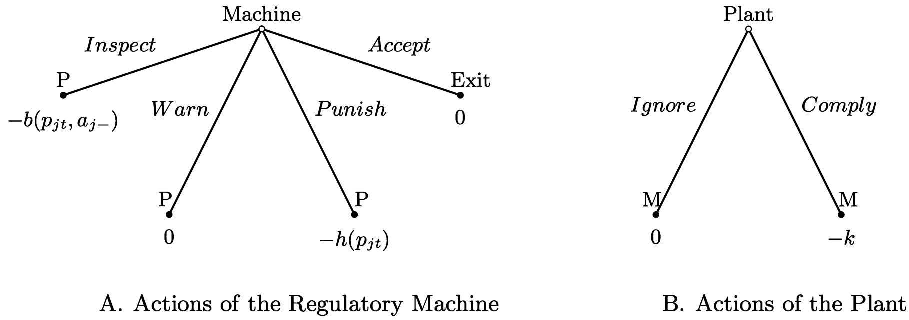

FIGURE 4.—Actions of the regulatory machine and plant at each node. The terminal nodes give the payoffs in each round for the plant. The penalty stage begins with an inspection where the regulatory machine (M) observes p_{j1}. The machine can take four actions. If M inspects, M gets a new signal of pollution 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.

round after the first, the regulatory machine has four actions a_{Rt}: inspect, warn, punish, or accept, which correspond to categorized regulatory data (Section 2.3).

Figure 4 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.[22] We assume that the plant knows the regulatory machine's action probabilities in each possible future state.

4.3. *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

[22]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 support this assumption: the average time between chains, about 5 months, is much larger than the average time between actions within a chain, 2 weeks. Further, recent pollution readings do not change regulatory targeting of inspections (Appendix Table S.X, column 4). Last, the regulator has a short memory: 93% of the time when an action cites a prior inspection it is the most recent prior inspection.

one parameterization of the model. The simulations show that the results of this trade-off vary with the regulator's information (see Appendix Figure S.2 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.

5. ESTIMATION

The model is estimated moving backward. 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.

5.1. *Maximum Likelihood for Penalty Stage*

Building the likelihood for plant actions requires several preliminary steps. (i) We specify that the common state of the game comprises the pollution reading, the last two actions of the regulator and plant, and the game round. (ii) We estimate state transition probabilities using a count estimator. (iii) 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 C of the supplementary material. 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 $b(p_{jt}, a_{j-}) = (1 - \mathbf{1}\{a_{j-} = \mathit{Comply}\}) \times (\nu_1 \mathbf{1}\{\overline{p} < p_{jt} < 2\overline{p}\} + \nu_2 \mathbf{1}\{2\overline{p} \leq p_{jt} < 5\overline{p}\} + \nu_3 \mathbf{1}\{5\overline{p} \leq p_{jt}\})$. This function specifies that inspections are costless for plants that have recently complied, but for plants that have not complied, 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 that 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 backward induction to calculate the value of the penalty stage $V_0(\cdot)$ for each level of pollution at the time of an initial inspection.

5.2. *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 unobserved 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.

5.2.1. *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 and 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 they are 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 6.2.2 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, yields

$$g_1(\phi) = Z_j'(\log P_j - \phi_0 - \phi_1 X_j - \phi_2 \mathit{Run}).$$

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

$$g_2(\lambda, \beta) = \mathbf{1}'\big(\mathbb{E}\big[\mathcal{I}(u_{1j}|X_j, T_j, \lambda, \beta, \rho)\big] - I_j\big),$$

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

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 10% 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 roughly the midpoint 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.[23] 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 6.2.2 considers the robustness of the targeting stage estimates to these assumptions.

5.2.2. *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

[23] 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.

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 nonlinear 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.

5.2.3. *Targeting Stage Objective Function*

We stack the moments to form $g(\theta_T) = [g_1' \; g_2' \dots g_7']$ 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 optimal estimator. Standard errors are calculated to account for the nonlinear constraints and their correlation with the moments (Newey and McFadden (1994)), and to adjust for simulation bias, which is negligible with $S = 5000$.

6. STRUCTURAL ESTIMATES OF REGULATORY COSTS AND TARGETING

6.1. *Penalty Stage*

6.1.1. *Plant and Regulatory Choice Probabilities in the Penalty Stage*

Table IV 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 1–3). 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.

6.1.2. *Revealed Preference Penalty Estimates*

Table V 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 \$3000). 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

TABLE IV

MULTINOMIAL LOGIT MODEL OF ACTION CHOICE CONDITIONAL ON STATE[a]

	Regulatory Machine			Plant
	Inspect	Warn	Punish	Comply
Party to move:	(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

[a]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 S.I 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. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

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.[24] On profits, mean plant annual sales in our end-line survey are \$2.9 million. The

[24]A plant with an extremely high pollution reading has a one-third 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).

TABLE V

ESTIMATES OF PLANT UTILITY PARAMETERS (USD × 10^3)[a]

	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

[a]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 onward; $t = 2$ 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 given in parentheses.

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.

6.1.3. *The Value of Environmental Regulation to Plants*

The value of the penalty stage to a plant summarizes all costs of environmental regulation. Figure 5 shows this value, at a variety of states, as calculated through backward induction given the estimated costs of regulatory penalties from Table V, column 2. At each state, values are divided between expected discounted future abatement costs (light

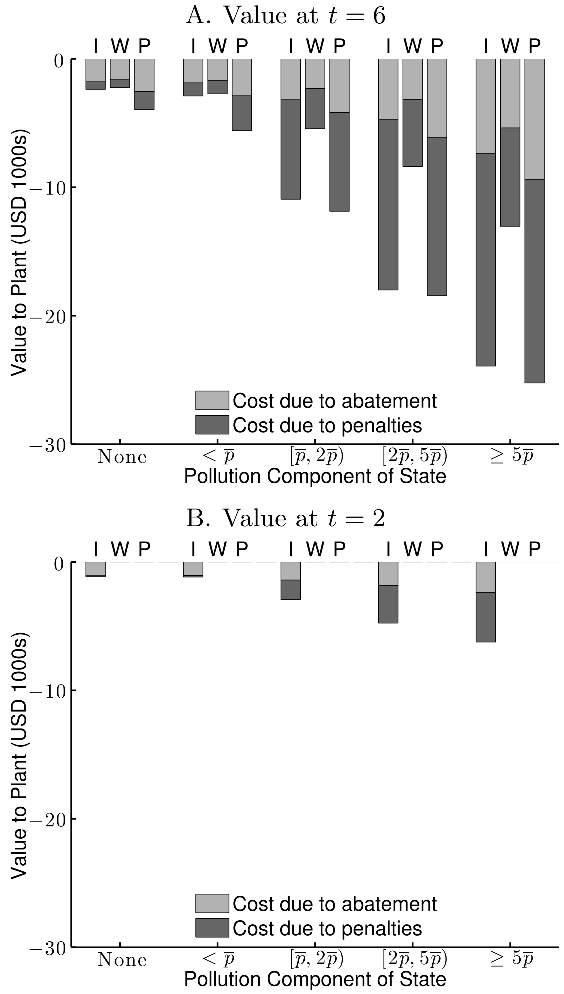

FIGURE 5.—Value of environmental regulation for plants. The figure shows the cost of regulation to plants in thousands of U.S. dollars as measured by the expected discounted value of different states in the penalty stage. Values are divided between expected discounted future abatement costs (light grey) 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 evaluated when it is the plant's turn to move (A) at $t = 6$, and (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 changed if the regulatory machine's lagged action was inspect, warn, or punish, respectively.

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 is shown 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. Figure 5(A) shows the plant value at $t = 6$. The value to the 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.

Figure 5(B) shows the expected discounted value for the plant when it can first act ($t = 2$). The value is shown only for 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 \$1160, if the machine did not take a pollution reading, down to negative \$6240, if the inspection found a pollution reading more than five times the standard, a more than fivefold 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 −\$2131, of which 40% is expected future penalties and 60% is 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 5(B) bars to obtain a smooth $\hat{V}_0(p)$ (Appendix Figure S.1 plots the resulting function). The resulting value function determines plant incentives for preemptory abatement in the targeting stage.

6.2. *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.

6.2.1. *Estimates*

Columns 1 and 2 of Table VI present 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

TABLE VI

ESTIMATES OF TARGETING STAGE PARAMETERS[a]

	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 parameter (λ_1)	−0.395		−0.220	
	(0.003)		(0.066)	
Inspection targeting level parameter (λ_2)	33.022		10.064	
	(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 pollution shock (σ_1)	0.069		0.111	
	(0.003)		(0.022)	
Standard deviation of unobserved pollution shock (σ_2)	1.033		0.864	
	(0.047)		(0.042)	
Mean of log maintenance cost (μ_c)	2.388		1.833	
	(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			

[a]The table reports parameters of the targeting stage of the model. The first two columns 1 and 2 report estimates from the constrained model where the regulator is constrained to target optimally based on observed pollution shocks, and columns 3 and 4 report estimates from the unconstrained model. Within each pair of columns, the first column (columns 1 and 3, respectively) reports the coefficients on the inspection equation and the second column (columns 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 ((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 2-year experiment, and pollution is the maximum (over pollutants) logged standardized end-line 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 5. Standard errors are bootstrapped with $B = 100$ and $S = 5000$.

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 is $\hat{\lambda}_2 = 33$ 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 shock at the 95th percentile.

In columns 3 and 4 of Table VI, 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 VI, 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 6.2.3 argues that the fit of the constrained model is still quite good.

6.2.2. *Sensitivity Analysis and Robustness Checks*

To develop intuition for how features of our data affect parameter estimates, we use the measure introduced by Andrews, Gentzkow, and Shapiro (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 that underlie 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 6 reports the sensitivity of select parameter estimates (across panels) to select moments (row headers within the panel). The length of each bar is the increase (solid

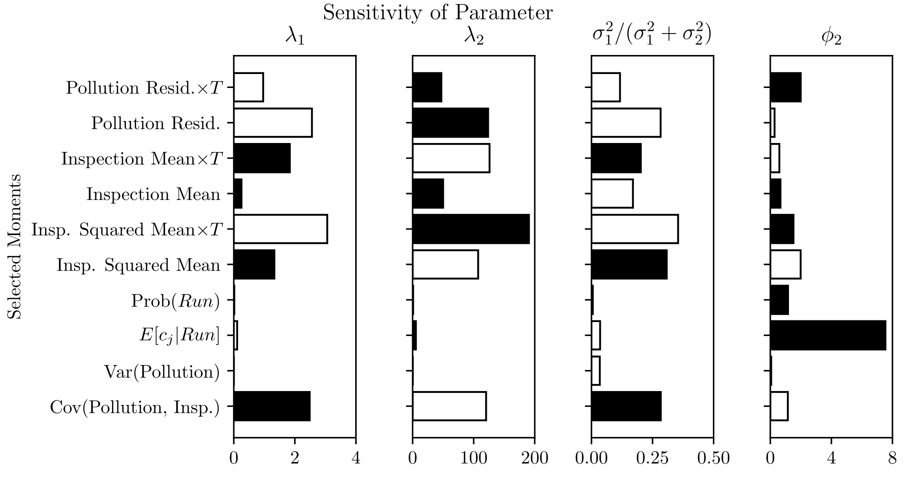

FIGURE 6.—Sensitivity of targeting parameters to moments. The figure shows the Andrews, Gentzkow, and Shapiro (2017) sensitivity measure of selected targeting model parameter estimates with respect to selected moments used to estimate the model. The panels, left to right, show the sensitivity of the parameters or functions of parameters λ_1, λ_2, and $\sigma_1^2/(\sigma_1^2+\sigma_2^2)$, 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 1 standard deviation change in the moment. From the rows, we omit 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 indicate negative sensitivity.

bar) or decrease (hollow bar) in each parameter that would result from increasing the row moment by 1 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 followups.

The differences in inspection moments between the treatment and the control groups play a key role in estimating the targeting equation parameters. The targeting parameter λ_2 is most sensitive, among all moments, to the mean products of the treatment with inspections and inspections squared (Figure 6, panels λ_1 and λ_2; see also Appendix Table S.XV). 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 6, panel $\sigma_1^2/(\sigma_1^2+\sigma_2^2)$). For example, a 1 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 to 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, he/she 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 1 standard deviation, then $\hat{\phi}_2$ would increase from -0.71 to -0.51 (Figure 6, 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 $\mathbb{E}[c_j|\text{run}]$).

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 1 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 S.XIII).

6.2.3. *Model Fit*

There are several ways to assess the model's fit. Figure 7 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.[25] The y-axis reports the number of inspections for a control plant with an 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; 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

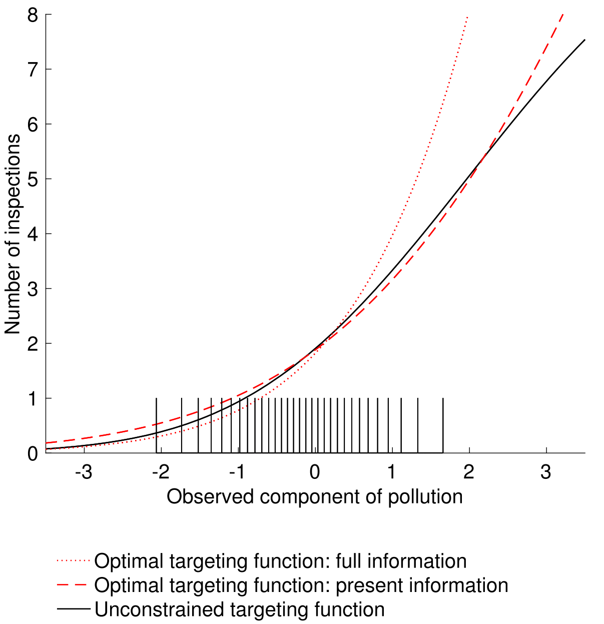

FIGURE 7.—Observed and optimal targeting rules. The figure shows several inspection targeting rules that assign plants an annual rate of inspection as a function of the component of pollution observed by the regulator. Two of the targeting rules are based on the estimates of Table VI: the dashed line gives the estimated optimal targeting function from the constrained estimates in the table, and the solid line gives the estimated targeting function from the unconstrained estimates. The dotted 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.

[25]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.

that is very similar to that 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 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 2(C) and (D) shows the modeled distributions of inspections in the control and treatment groups, respectively, and can be compared to the true distribution in (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 2(E) shows the distribution of pollution in the data and Figure 2(F) shows the model. The model was fitted 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(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(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.

7. 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. The second regime is a targeting regime with discretion, where the regulator solves (3), but where the regulator has only the sparse information in observed pollution. The third regime is 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). The fourth regime is 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 line shows abatement under the uniform rule that requires all plants to be inspected the same

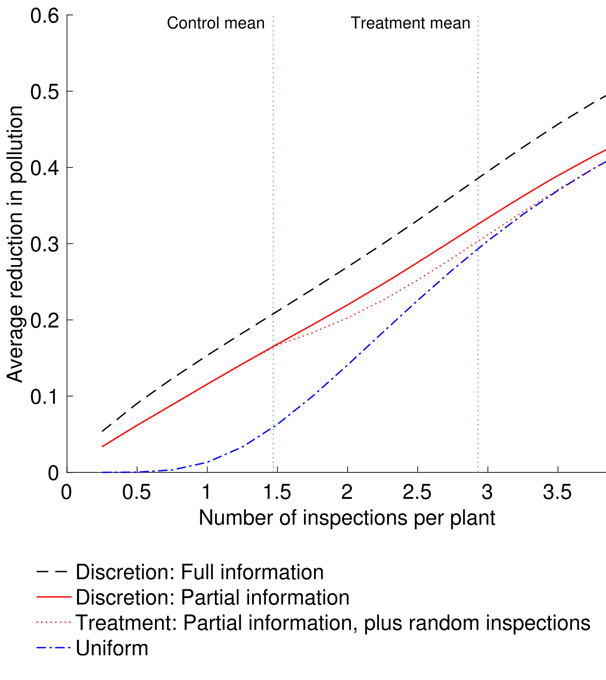

FIGURE 8.—Value of discretion: Abatement by information regime and budget constraint. The figure shows the amount of pollution abatement achieved, in units of regulatory standards of abatement per plant, under different counterfactual inspection regimes. Regimes vary in the information used by the regulator 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 line shows abatement under a minimum threshold rule that requires all plants to be inspected the same number of times. The solid line shows abatement under the optimal targeting rule where targeting is based on the observed component of pollution. The dashed line shows abatement under the same optimal targeting rule where the regulator is assumed to observe all variation in pollution. The dotted 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.

number of times. The solid line shows abatement under the optimal targeting rule where targeting is based on observed pollution only. The dashed 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 standard deviations. 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 preemptive 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 (solid line) as compared to a constant rule (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 line from the left-hand to the right-hand dotted vertical line, increasing average abatement by 0.14 standard deviations. 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 line), we estimate abatement would have been about 15% greater. In other words, doubling inspection while keeping the same level of discretion would have decreased 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 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 regulator has access to a perfectly functioning monitoring technology. The difference in abatement under an optimal targeting regime with full (dashed line) versus estimated (solid line) 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 dashed and solid line increases). Full information in a discretionary regime, at the control budget, is as valuable as doubling inspections in a nondiscretionary regime. It is apparent that full information allows the regulator to target its inspections more precisely and this substantially increases abatement.

Full information can, in principle, be achieved by the use of continuous emissions monitoring systems (CEMS), devices that 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 that CEMS monitoring of particulate matter in Gujarat today costs about $1800(U.S.) per plant-year, on an amortized basis, whereas a single in-person inspection with air pollution sampling costs $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 the plants are left alone. This targeting would grow more valuable if information were improved.

8. 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 30% 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 in 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, in a setting with weak institutional capacity, these devices are in fact a reliable and cost-effective substitute for in-person monitoring.

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., N. KALA, AND A. NYSHADHAM (2016): "Management and Shocks to Worker Productivity," Report, University of Michigan. [2124]

AGHION, P., AND J. TIROLE (1997): "Real and Formal Authority in Organizations," *Journal of Political Economy*, 105, 1–29.[2127]

AGHION, P., R. BURGESS, AND S. REDDING (2008): "The Unequal Effects of Liberalization: Evidence From Dismantling the License Raj in India," *American Economic Review*, 98 (4), 1397–1412.[2127]

ANDREWS, I., M. GENTZKOW, AND J. M. SHAPIRO (2017): "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *The Quarterly Journal of Economics*, 132 (4), 1553–1592.[2126,2152]

BANERJEE, A. V., AND E. DUFLO (2014): "Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program," *The Review of Economic Studies*, 81 (2), 572–607.[2144]

BESLEY, T., AND R. BURGESS (2004): "Can Labor Regulation Hinder Economic Performance? Evidence From India," *Quarterly Journal of Economics*, 119 (1), 91–134.[2127]

BESLEY, T., AND T. PERSSON (2010): "State Capacity, Conflict and Development," *Econometrica*, 78, 1–34.[2127]

BOYER, M., AND J.-J. LAFFONT (1999): "Toward a Political Theory of the Emergence of Environmental Incentive Regulation," *The RAND Journal of Economics*, 30 (1), 137–157.[2127]

BURGESS, R., M. HANSEN, B. OLKEN, P. POTAPOV, AND S. SIEBER (2012): "The Political Economy of Deforestation in the Tropics," *Quarterly Journal of Economics*, 127 (4), 1707–1754.[2127]

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). [2124]

CENTRAL POLLUTION CONTROL BOARD (2013): "Specifications and Guidelines for Continuous Emissions Monitoring Systems (Cems) for pm Measurement," CPCB/e-Publication/2013-14. [2126,2158]

CENTRAL POLLUTION CONTROL BOARD (2014): "Directions… Regarding Self-Monitoring of Compliance," CPCB/B-29016/04/06/PCI-I/5401. [2126,2158]

CHANG, T., J. G. ZIVIN, T. GROSS, AND M. NEIDELL (2016): "The Effect of Pollution on Worker Productivity: Evidence From Call-Center Workers in China," Technical Report, National Bureau of Economic Research. [2124]

CHEN, Y., A. EBENSTEIN, M. GREENSTONE, AND H. LI (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.[2123]

DUFLO, E., M. GREENSTONE, R. PANDE, AND N. RYAN (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.[2124,2128,2135,2159]

——— (2018): "Supplement to 'The Value of Regulatory Discretion: Estimates From Environmental Inspections in India'," *Econometrica Supplemental Material*, 86, https://doi.org/10.3982/ECTA12876. [2128]

EBENSTEIN, A., M. FAN, M. GREENSTONE, G. HE, AND M. ZHOU (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.[2124]
ENVIRONMENTAL PROTECTION AGENCY (2012): "Air Pollution Technology Fact Sheets," Technical Report, Technology Transfer Network: U.S.-Mexico Border Information Center on Air Pollution. [2140]
ESTACHE, A., AND L. WREN-LEWIS (2009): "Toward a Theory of Regulation for Developing Countries: Following Jean-Jaques Laffont's Lead," *Journal of Economic Literature*, 47 (3), 729–770.[2127]
GRAFF-ZIVIN, J., AND M. NEIDELL (2012): "The Impact of Pollution on Worker Productivity," *American Economic Review*, 102 (7), 3652–3673.[2124]
GREENSTONE, M., AND R. HANNA (2014): "Environmental Regulations, Air and Water Pollution and Infant Mortality in India," NBER Working Paper 17210. [2127]
GREENSTONE, M., J. A. LIST, AND C. SYVERSON (2012): "The Effects of Environmental Regulation on the Competitiveness of US Manufacturing," Technical Report, National Bureau of Economic Research. [2127]
GREENSTONE, M., J. NILEKANI, R. PANDE, N. RYAN, A. SUDARSHAN, AND A. SUGATHAN (2014): "Lower Pollution, Longer Lives: Life Expectancy Gains if India Reduced Particulate Matter to Air-Quality Standards," Report, Harvard Kennedy School. [2123,2124]
GUJARAT POLLUTION CONTROL BOARD (2008): "Annual Action Plan 2008-09," Technical Report, Govt. of Gujarat. [2131]
HANNA, R. N., AND P. OLIVA (2010): "The Impact of Inspections on Plant-Level Air Emissions," *The BE Journal of Economic Analysis & Policy*, 10 (1), 19.[2126,2128]
HE, J., H. LIU, AND A. SALVO (2016): "Severe Air Pollution and Labor Productivity: Evidence From Industrial Towns in China," IZA Discussion Paper No. 8916. [2124]
JIA, R. (2014): "Pollution for Promotion," Report, UCSD. [2127]
LAFFONT, J.-J. (1994): *Regulation of Pollution With Asymmetric Information*. Springer. [2127]
——— (2005): *Regulation and Development*. Cambridge University Press. [2127]
LAFFONT, J.-J., AND J. TIROLE (1993): *A Theory of Incentives in Procurement and Regulation*. MIT Press. [2127]
LEAVER, C. (2009): "Bureaucratic Minimal Squawk Behavior: Theory and Evidence From Regulatory Agencies," *American Economic Review*, 99 (3), 572–607.[2124]
MAGAT, W. A., AND W. K. VISCUSI (1990): "Effectiveness of the EPA's Regulatory Enforcement: The Case of Industrial Effluent Standards," *The Journal of Law & Economics*, 33, 331–360.[2126]
MAGNAC, T., AND D. THESMAR (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70 (2), 801–816.[2144]
MINISTRY OF ENVIRONMENT AND FORESTS (1986): "The Environment (Protection) Rules, 1986," *Gazette of India*, 844 (E).[2127]
——— (1999): "Prevention and Control of Pollution (Uniform Consent Procedure) Rules, 1999," Technical Report Part II, Section 3(ii), Schedule IV, *The Gazette of India*. Available at http://icpe.in/pdf/2.pdf. [2128]
NEWEY, W. K., AND D. MCFADDEN (1994): "Large Sample Estimation and Hypothesis Testing," *Handbook of Econometrics*, 4, 2111–2245.[2146,2152]
RASUL, I., AND D. ROGGER (2013): "Management of Bureaucrats and Public Service Delivery: Evidence From the Nigerian Civil Service," Report, UCL. [2127]
RUST, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica: Journal of the Econometric Society*, 55 (5), 999–1033.[2144]
——— (1994): "Advances in Econometrics, Sixth World Congress," in *Chapter Estimation of Dynamic Structural Models, Problems and Prospects: Discrete Decision Processses*. Cambridge University Press.[2144]
RYAN, S. P. (2012): "The Costs of Environmental Regulation in a Concentrated Industry," *Econometrica*, 80 (3), 1019–1061.[2127]
SHIMSHACK, J. P. (2014): "The Economics of Environmental Monitoring and Enforcement," *Annu. Rev. Resour. Econ.*, 6 (1), 339–360.[2126]
STIGLER, G. (1971): "The Theory of Economic Regulation," *Bell Journal of Economics and Management*, 2 (1), 3–21.[2124]
TANAKA, S. (2013): "Environmental Regulations on Air Pollution in China and Their Impact on Infant Mortality," Report, Tufts University. [2127]
THE ECONOMIC TIMES (2014): "PM Narendra Modi Unveils Labour Reforms; Launches Universal Account Number for Employees," Technical Report. Available at https://economictimes.indiatimes.com/news/economy/policy/pm-narendra-modi-unveils-labour-reforms-launches-universal-account-number-for-employees/articleshow/44833365.cms. [2124]

Co-editor Liran Einav handled this manuscript.

Manuscript received 30 September, 2014; final version accepted 15 March, 2018; available online 18 May, 2018.