

NBER WORKING PAPER SERIES

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

Esther Duflo
Michael Greenstone
Rohini Pande
Nicholas Ryan

Working Paper 20590
<http://www.nber.org/papers/w20590>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2014

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 and Kevin Rowe for excellent research assistance and numerous seminar participants for comments. We thank the Sustainability Science Program (SSP), the Harvard Environmental Economics Program, the Center for Energy and Environmental Policy Research (CEEPR), the International Initiative for Impact Evaluation (3ie), the International Growth Centre (IGC) and the National Science Foundation (NSF Award #1066006) for financial support. All views and errors are solely ours. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w20590.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Esther Duflo, Michael Greenstone, Rohini Pande, and Nicholas Ryan. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Value of Regulatory Discretion: Estimates from Environmental Inspections in India
Esther Duflo, Michael Greenstone, Rohini Pande, and Nicholas Ryan
NBER Working Paper No. 20590
October 2014
JEL No. D22,L51,Q56

ABSTRACT

In collaboration with a state environmental regulator in India, we conducted a field experiment to raise the frequency of environmental inspections to the prescribed minimum for a random set of industrial plants. The treatment was successful when judged by process measures, as treatment plants, relative to the control group, were more than twice as likely to be inspected and to be cited for violating pollution standards. Yet the treatment was weaker for more consequential outcomes: the regulator was no more likely to identify extreme polluters (i.e., plants with emissions five times the regulatory standard or more) or to impose costly penalties in the treatment group. In response to the added scrutiny, treatment plants only marginally increased compliance with standards and did not significantly reduce mean pollution emissions. To explain these results and recover the full costs of environmental regulation, we model the regulatory process as a dynamic discrete game where the regulator chooses whether to penalize and plants choose whether to abate to avoid future sanctions. We estimate this model using original data on 10,000 interactions between plants and the regulator. Our estimates imply that the costs of environmental regulation are largely reserved for extremely polluting plants. Applying the cost estimates to the experimental data, we find the average treatment inspection imposes about half the cost on plants that the average control inspection does, because the randomly assigned inspections in the treatment are less likely than normal discretionary inspections to target such extreme polluters.

Esther Duflo
Department of Economics, E17-201B
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
eduflo@mit.edu

Rohini Pande
Kennedy School of Government
Harvard University
79 JFK Street
Cambridge, MA 02138
and NBER
rohini_pande@harvard.edu

Michael Greenstone
University of Chicago
Department of Economics
1126 E. 59th Street
Chicago, IL 60637
and NBER
mgreenst@uchicago.edu

Nicholas Ryan
Harvard Weatherhead Center
K239 1737 Cambridge St.
Cambridge, MA 02138
nicholas.ryan@yale.edu

I Introduction

Regulatory agencies enforce standards using imperfect information on the parties they regulate (Laffont and Tirole, 1993). This imperfection suggests a need for policy to be flexible, to allow regulators to collect and use local information. Yet, flexibility lends discretion, which may lead to enforcement that reflects regulators' personal objectives, rather than the social goals of regulation (Stigler, 1971; Leaver, 2009). This tension between rules and discretion underlies many debates on the design of regulatory agencies (Sitkin and Bies, 1994).

Environmental regulation in emerging economies is an important case in point. Countries like China and India report levels of pollution that exceed the highest ever recorded in rich countries, with dire health consequences (Chen et al., 2013).¹ This pollution problem is severe despite strict *de jure* regulatory standards.

To shed light on regulatory design in practice, we implemented a two-year field experiment which increased the frequency of inspection for a sample of industrial plants in the Indian state of Gujarat, in collaboration with the Gujarat Pollution Control Board (GPCB). The *de jure* environmental standards are strict maximum limits on the concentrations of pollution that plants can emit. There are also regulatory requirements for inspection frequency and national laws that build a criminal framework for punishment, giving regulators broad power to close down highly polluting plants and cut off their utilities. However, *de facto*, we observe that the regulator uses considerable discretion in deciding whom to inspect, whom to pursue and when to apply the costly punishments at its disposal. In the control group, despite widespread noncompliance with pollution standards, 44% of plants are inspected less than once per year during the experiment.

The experiment was conducted in a sample of 960 industrial plants, with half assigned to the treatment group of increased inspection frequency. This "inspection" treatment had two features. First, it provided the resources necessary to bring all treatment plants up to at least the required minimum frequency of inspections. Second, it removed the regulator's discretion by requiring that these added inspections be allocated randomly across *all* treatment plants. For a regulator that goes fishing for violations in a sea of industrial plants, the treatment expanded its fleet, but sent the additional boats to all corners of the ocean, not just its favored grounds. The treatment

¹By one estimate, reducing particulate matter air pollution to India's air quality standards would increase life expectancy for the 660 million people in noncompliant areas by 3.2 years on average (Greenstone et al., 2014).

did not otherwise alter pollution standards, regulatory penalties or the process of regulation. This inspection treatment was cross-randomized with an audit reform experiment in the subset of 473 sample plants that were eligible for environmental audits (Duflo et al., 2013).²

The treatment was successful when judged by process measures, like increased regulatory scrutiny. Actual inspection rates more than doubled for treatment plants relative to control plants, and treatment plants report higher perceived inspection rates, suggesting the additional scrutiny was felt. Additionally, the treatment greatly increased the number of plants found with pollution emissions that exceed the legal standard and the number of regulatory citations for such violations. However, the treatment was weaker for more consequential outcomes. For example, the regulator was no more likely to identify extreme polluters (i.e., plants with emissions five times the regulatory standard or more), and neither the frequency of more costly regulatory penalties nor plants' perception of these penalties differ by plant treatment status. Perhaps most importantly, we measured the effect of the treatment on pollution levels by conducting an endline survey of emissions independent of the regulator. Treatment plants only marginally increased their compliance with pollution standards (four percentage points, or 6%; p-value = 0.087) and did not significantly reduce mean pollution emissions. These pollution results are mixed but consistent, since reductions came mostly from plants close to the regulatory threshold, rather than from large polluters who had the most potential to bring average pollution down.

To understand these results, we follow the regulator's net under the water to see what fish are ensnared and which go free. We obtained access to the regulator's administrative database and identified 9,624 pieces of correspondence between the regulator and the plants in our sample. These provide plant characteristics and document all interactions with sample plants over five years (from two years before the experiment through one year after), including plant inspections, pollution analysis, regulatory notices and penalties, as well as written responses by sample plants, such as documentation of the installation of abatement equipment. We use this data to estimate a dynamic discrete game between the regulator and polluting plants to recover the full costs of environmental regulation (Aguirregabiria and Mira, 2010; Pakes et al., 2007). Using explicit cross-references and dating, we assigned documents to 7,423 separate chains of interactions, each

²In 1996, the High Court of Gujarat ordered GPCB to instate a third-party audit system wherein high-polluting potential plants must provide an annual audit report to GPCB. Duflo et al. (2013) evaluated a potential reform of this audit system and found making third party auditors more accountable to the regulator and less beholden to the plants they audit improves truth-telling and lowers pollution.

of which begins with a new plant inspection and ends with the regulator accepting that the plant, for now, is in compliance. Chains alternate moves between the regulator and plants and go on for as many as nineteen rounds in the data.

The model is necessary to estimate the full costs of regulation because, while we can observe mandated expenditures on pollution abatement equipment, the costs of regulatory penalties (e.g., temporary plant closures) in terms of lost profits are not measured, but only implicitly revealed through plant behavior. In the game, the regulator observes a measure of plant pollution on its inspection and may choose, at any later round, to re-inspect a plant, warn it for past violations, punish it with a costly penalty or accept it as compliant and end the game. The plant can comply by installing abatement equipment or ignore the regulator altogether. Plants are trying to minimize regulatory costs (maximize profits) by trading-off a known cost for abatement equipment today against the risk of regulatory penalties for noncompliance, such as temporary closure. Using this structure, we estimate conditional choice probabilities for regulatory actions depending on pollution and the past actions of both parties. We use these choice probabilities and state transition probabilities to construct the plant's value function for each action at each possible state, and estimate what regulatory penalties cost to plants by maximizing the pseudo-likelihood of observing the actions taken in the data.

The structural estimates show that investments in pollution abatement, which are a typical measure of regulation's costs, account for less than two-thirds the total cost of regulation. The balance is due to regulatory penalties that can be substantial: 3% of plants were closed multiple times during our two-year experiment and one plant was closed on five separate occasions. We find that the regulator largely reserves major penalties for plants with pollution readings that exceed the *de jure* standard by a factor of five times or more. In turn, plants are most likely to comply when such a high reading is cited against them. To provide some idea of the magnitude of the costs of regulation, the average expected discounted cost of regulation for an extreme polluter in a later round of the game is US\$ 23,000, with 67% of this cost due to regulatory penalties and the remainder to pollution abatement costs. This cost is equivalent to about one month of mean plant profits, assuming a 10% profit margin, or two months of mean plant electricity bills.

We use the full costs of regulation as estimated from the structural model and the experimental variation in inspections to assess the value of regulatory discretion. We find that the average

inspection in the treatment, which is mostly assigned randomly, leads to about US\$ 1,340 of regulatory costs.³ This is roughly half of the US\$ 2,650 in regulatory cost due to the average inspection in the control where inspections are mostly discretionary. We interpret regulatory costs per inspection as the value of an inspection consistent with the regulator’s own objectives and preferences, but do not take it as a measure of social welfare.

The combination of a large-scale field experiment and structural estimation allows us to make several contributions to the economic literatures on regulatory design and environmental regulation. First, we believe this paper provides the first experimental evidence of the effect of inspection frequency on compliance with environmental regulation.⁴ Second, the structural estimates allow us to estimate the full costs of environmental regulation and separate them into mandated investment in abatement equipment and the regulatory penalties needed to force that investment. We show that these penalties are a substantial portion of the total cost of environmental regulation, raising the social cost of abatement through regulation relative to the private cost of abatement to plants. Third, marrying the experimental variation in inspection frequency with the estimates from the structural model allows us to value regulatory discretion. With the rigid penalty structure in this setting, the data suggest that regulatory discretion is valuable because it allows the regulator to target inspections at extreme polluters.⁵

The rest of the paper runs as follows. Section II describes environmental regulation in India, the experimental design and data collection. Section III gives reduced-form results on the effects

³We say treatment inspections are “mostly” random because even in the treatment there are non-random follow-up inspections.

⁴Past studies of regulatory inspections in the United States show that inspections reduce pollution significantly (Hanna and Oliva, 2010; Magat and Viscusi, 1990), but at the cost of lower manufacturing productivity (Greenstone et al., 2012). All these studies rely on observational data wherein dirtier plants are more likely to get inspections (Hanna and Oliva, 2010). 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. Few studies of the productivity impacts of environmental regulation exist, though negative productivity effects of rigid industrial regulation have been documented for India (Aghion et al., 2008; Besley and Burgess, 2004).

⁵Environmental regulation is a classic setting for incentive regulation (Laffont and Tirole, 1993; Laffont, 1994; Boyer and Laffont, 1999). Limited regulatory capacity and limited commitment will typically cause the optimal regulatory policy in emerging economies to differ dramatically from settings with complete markets (Laffont, 2005; Estache and Wren-Lewis, 2009). In addition, in these settings, the regulator may have especially poor information on firm cost structure as firms are numerous and monitoring spread thin (Dufo et al., 2013). 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 an emerging 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, that discretion in environmental regulation may be misused by bureaucrats and politicians (Burgess et al., 2012; Jia, 2014).

of additional inspections on regulatory actions and plant pollution behavior. Section IV presents a dynamic discrete game model of the interactions between the regulator and polluting plants. Section V presents the model estimates on the costs of regulatory penalties for treatment and control plants and uses them to value discretion in regulatory inspections. Section VI concludes.

II Context, Experimental Design, Data, and Summary Statistics

This section describes the Indian regulatory framework, the experimental design and data, and presents summary statistics on the sample plants and their interactions with the regulator.

A Regulation of Industrial Pollution in India

India remains severely polluted despite powerful regulatory agencies and stringent environmental regulations (Greenstone and Hanna, 2014). Regulations are command-and-control in nature: plants must meet maximum allowable concentration standards for pollution emissions in air and water. 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 conduct basically all enforcement of environmental regulations.

We describe the enforcement practices of our partner regulator, the Gujarat Pollution Control Board (GPCB), which are largely common with other Indian states. GPCB has 19 regional offices, each with several inspection teams and a Regional Officer in charge of assigning inspections. Plant inspections are undertaken by full-time engineer and scientist employees. During an inspection, the team observes the state of the plant and its environmental management and often, but not always, collects pollution samples for laboratory analysis. Findings are summarized in an inspection report which enters the plant's file, and the laboratory analysis report, once completed, is added to the same file. Officers in the region and at the head office review inspection results and analysis reports, covering qualitative information, such as on waste management practices and plant housekeeping, the plant's history of violations and especially pollution concentrations for air and water pollutants. The officers, not the field staff, decide whether to take any action, including whether to inspect the plant again, warning the plant or imposing penalties (on which see below), or to leave the plant alone.

A first plant inspection may occur for several reasons. First, 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.⁶ In administrative baseline data, routine inspections make up 35% of inspections. Second, inspections occur when plants apply for an environmental consent, a license to operate a certain type of manufacturing at a certain scale, and for renewal of this consent every five years; 30% of GPCB inspections fall in this category. Third, the inspection may be to follow-up on a violation found in a prior inspection, or to verify that the plant has taken a specific action as it claimed (24% of inspections). The balance of 11% of inspections are in response to public complaints and other miscellaneous reasons.

The staff strength to attend to these competing priorities is, however, limited. Bhushan et al. (2009) note that the approved number of employees for many Boards has declined even as the number of plants the Boards were to regulate doubled or trebled. Of nominally approved staff positions, many Boards left a quarter or more unfilled due to financial or bureaucratic difficulties in hiring. As a result, inspections occur less often than prescribed. Data from the year before the experiment shows that inspection rates for red-category plants of small-, medium- and large-scale occur roughly half as often as prescribed (once in 187, 224 and 534 days, respectively). Finally, a subset of red-category plants are also required to submit annual environmental audit reports to the regulator. Auditors’ conflict of interest has historically rendered these audits inaccurate, such that inspections remain the key channel through which the regulator monitors plants (Duflo et al., 2013).

While inspections are infrequent, plants found in violation of pollution standards can be harshly penalized. National environmental regulations empower the regulator to issue written directions for polluting plants to take abatement action or for the stoppage of water, electricity or other services to those plants (Ministry of Environment and Forests, 1986). The regulator applies these penalties often, including mandating that a plant install abatement equipment, mandating that a plant post a bond against future performance and ordering that a plant’s water and electricity be disconnected, closing the plant down. Utility disconnections can remain

⁶The GPCB follows a government classification for plants based on their reported scale of capital investment, with small-scale being investment less than INR 50m (US\$ 1m), medium INR 50m to 100m (US\$ 1m to 2m) and large above 100m (US\$ 2m). Throughout the paper, we use an exchange rate of US\$ 1 = INR 50, a round figure about equal to the exchange rate at the end of 2011 and slightly stronger (for the dollar) than the average of INR 46 over the experiment. The prescribed inspection rates are roughly comparable to those applied to large plants, by air pollution potential, in the United States (Hanna and Oliva, 2010).

in force until the plant has shown progress towards meeting environmental standards, often, again, by installing abatement equipment. The duration of closure depends on the offense but is typically several weeks; in our data the median duration of closure is 24 days.⁷ 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 to legal standards is burdensome and there are long delays in prosecution.

B Experimental Design

We worked with GPCB to design an experiment wherein the Board, between 2009 (quarter 3) and 2011 (quarter 2), raised inspection frequency for a randomly selected subset of plants out of a sample of 960 highly-polluting plants. The sample was drawn as follows. First, we identified the population of 3,455 red-category (i.e., high pollution potential) small- and medium-scale plants in three regions of Gujarat (Ahmedabad, Surat and Valsad). This population, which makes up roughly 15% of the more than 20,000 regulated plants in Gujarat, is of special interest as it is highly polluting but hard to regulate. From this population, we selected all 473 audit eligible plants in Ahmedabad and Surat.⁸ Next, we randomly selected 488 other plants from the remaining non-audit-eligible population.

Inspection treatment assignment was randomized within region by audit-treatment-status strata (treatment, control and non-eligible). The treatment was thus cross-randomized and implemented concurrently with the audit reform treatment studied by Duflo et al. (2013). A total of 481 plants were assigned to the inspection treatment, and each treatment plant was guaranteed at least one annual inspection. Specifically, in the first quarter the plant was assigned an initial inspection, after which it was randomly assigned on a quarterly basis to be inspected again with probability 0.66, conditional on strata. After four quarters this cycle started over. Thus the treatment assigned extra variation in inspection frequency within the treatment group of plants, up to a maximum of four inspections per year.

Treatment inspections were conducted by newly created regional GPCB teams. Each team

⁷This estimate is from comparing the dates of closure orders and the revocation of closure orders for the same plant in regulatory files. This comparison probably somewhat overstates closure duration, because closures are carried out by an order copied to the electric utility, which executes these orders at a lag.

⁸Due to the requirements for audit eligibility, which were originally targeted at plants in Ahmedabad, few if any plants in Valsad are audit-eligible.

consisted of an environmental engineer and scientist. In order to increase inspections without overburdening GPCB staff, one member of each team was a recently retired GPCB scientist rehired for the experiment, and one an active environmental engineer.⁹ While this design does draw on active staff, GPCB judged that scientists, who are the only staff able to take certain pollution samples, were the main constraint in conducting additional inspections. Each morning, the inspection teams thus created were randomly assigned a list of plants from the treatment group at which to conduct a “routine” inspection that day. This mimicked GPCB’s practice of assigning teams to plants each day, except that for these teams the plant assignment, rather than being based on the RO’s discretion, was random. The additional teams were used only for these routine inspections and not other GPCB work or follow-ups. Regional officers could choose to shift their other teams away from routine inspections in treatment plants if they wished (though, below, we find little evidence of displacement). Inspection reports from the additional teams were treated like any other inspection report: the reports entered the same database, had samples analyzed by the same GPCB labs and had the same officials at the GPCB deciding whether or not to take action against plants.

Towards the end of the inspection treatment, in the month prior to the endline survey, we also assigned, randomly and independently of the other treatments, some plants to receive a letter from GPCB reminding them of their obligations to meet emissions limits. This letter reiterated the terms of plants’ environmental consent, which in principle they already knew, but it may have also served to increase the salience of regulatory compliance.

C Channels of Influence

How should we expect the experimental manipulation to affect plant and regulator behavior?

On the plants’ side, noncompliant plants facing a higher inspection rate may change their investments in abatement equipment and/or their emission levels. Plants may respond in this way to explicit penalties from the regulator, following inspections, or as part of a deterrence effect, from plants updating their expected penalties from remaining noncompliant. Treatment plants were not formally notified of the change in inspection frequency, but they would have inferred it from their experience. Plants that additionally received a letter from the regulator may have felt

⁹GPCB observes a mandatory retirement age so many retired staff are willing to work. GPCB identified recently retired but interested staff in each region.

under greater regulatory scrutiny. During the experiment, we do observe a significant difference in perceived inspection probability between treatment and control plants, which is necessary for a deterrence effect. Whether, in practice, plant behavior changes will depend on a plant's initial compliance status, expected penalties and the cost of abatement, a trade-off explored in detail in the context of the model.

On the regulator's side, the treatment had two features: inspection resources increased and were allocated across a representative group of plants. Given that pollution varies over time, this increase in inspections generates more opportunities to catch potential violators that are already being visited, as well as chances to catch new violators that are seldom visited. The number of new violations detected and penalties imposed will depend on the extent of regulatory targeting in the status quo and the regulatory penalty function, respectively, where the penalty function relates pollution violations to penalties. If the regulator was already inspecting a random sample of plants, and is now inspecting more of them and applying the same penalty function, violations and penalties would scale linearly with additional inspections. If the regulator was targeting more polluting plants in its inspections and penalties, then marginal random inspections may yield fewer violations and penalties than average inspections, whereas if the regulator was avoiding the heaviest polluters, marginal inspections may yield more.

D Data Collection

The analysis uses two sources of data, an endline survey and GPCB administrative records. The endline plant survey occurred between April and July, 2011 (the treatment increase in inspection frequency ran through the end of June), and was conducted by independent agencies, mainly engineering departments of local universities. The survey collected pollution readings and data on recent investments in pollution abatement equipment and other aspects of plant operations. The GPCB issued letters that required plants to cooperate with the surveyors and also stated that the results would not be used for regulatory matters. With this regulatory help, attrition was low; 12.9% of sample plants closed during the study and only 4.7% attrited for other reasons. Attrition did not differ by treatment status and nor did attrition on account of plant closure in particular (See Online Appendix Tables B1 and B2).

The second source of data, GPCB records of its own interactions with sample plants, allows

us to examine the process of regulatory enforcement. Regulation, in practice, does not move directly from observation to penalty but is mediated by a formal, sometimes lengthy dialogue between the regulator and regulated plants. The records consist of internal regulatory documents, letters and actions sent to plants and the written responses of plants to these actions. Table 1 presents an overview of the types of documents in the data, organized by the party and purpose or action to which they correspond. These categories presage the discrete actions in our model of regulation. Appendix A describes the data cleaning and collation process.

On the regulatory side, the main documents are internal inspection reports, warnings and orders to plants. *Inspection* reports describe GPCB’s findings on visiting a plant and accompanying analysis reports record pollution readings for any samples collected in a visit. *Warnings*, which may take several degrees of severity, urge the plant to clean up, cite it for pollution violations or threaten that the plant will be closed if it does not take some action to improve. *Punishments* consist of requiring plants to post bank guarantees or orders that the plant be closed down, possibly including a letter directly to the utility ordering that water or electricity be disconnected.

On the plant side, the main documents are plant responses to the regulator’s warnings and actions. We designate with the plant action *Comply* documents that record the installation of abatement equipment, typically with a certification or invoice from the vendor that did the work. The alternative to compliance is for the plant to *Ignore* the regulator, which is often inferred from the absence of any response. If the plant writes to the regulator to protest but does not demonstrate evidence of compliance, we also class this as *Ignore*, since protests are essentially cheap talk and involve no expenditures.¹⁰

We group GPCB and plant documents into chains of related actions, where each action in a chain references or logically follows the prior action. We observe 25,217 actions, of which 9,624 are documented and 15,593 imputed (*Ignore* is imputed 8,382 and *Accept* 7,211 times, as the complements to observable actions), which are grouped into 7,423 chains of related actions. All chains start with an inspection and end when GPCB *Accepts*. In between, the plant and

¹⁰For example, a plant where GPCB found high air pollution readings claimed in correspondence: “At the time of visit our chilling plant accidentally failed to proper working, so chilling system of scrubber was not effective by simple water. Same time batch was under reaction and we were unable to stop our reaction at that time. Now it is working properly.” That is, a piece of air pollution control equipment failed, causing pollution to be higher than normal during the visit. These types of explanations are common when plants are found well out of compliance.

regulator alternate moves at even- and odd-numbered rounds, respectively. (For completeness and to maintain the game structure, we allow the plant to move in the second round, though in the overwhelming majority of cases the regulator will take an action after an initial inspection without the plant having a chance to respond.) Of the 5,771 actions in chains with more than three rounds, 87% are linked based on the dating of the document relative to other documents concerning the same plant, and 13% (16% of non-imputed actions) are linked by an explicit reference to another document, for example when the regulator cites a prior inspection in taking action against a plant. (For reference, the Online Appendix Table B3 gives an example of a longer chain, showing the actions taken by each party at each turn and the underlying document for each action.)

E Status Quo Regulation and Randomization Check

Table 2 presents sample means and a randomization check for the inspection treatment. Each row considers a separate plant characteristic; columns 1 and 2 report the means for treatment and control plants, respectively, using administrative data for the year prior to the experiment. Column 3 reports the coefficient β on the inspection treatment dummy T_{ir} from the following regression, where for each outcome Y_{ir} for plant i in region r ,

$$Y_{ir} = \alpha_r + AuditSample_{ir} + \beta T_{ir} + \epsilon_{ir} \quad (1)$$

where α_r are region effects and $AuditSample_{ir}$ is a dummy for a plant belonging to the audit sample (i.e., being audit eligible, rather than being assigned to the audit treatment). Standard deviations are in brackets and standard errors for the column 3 coefficient in parentheses.

Panel A describes plant characteristics, which are well-balanced by treatment status. Forty-five percent of sample plants are in the textiles sector and about 15% are chemical plants (mainly upstream dye or dye intermediate producers). The average plant generates nearly 200,000 liters of waste effluent per day. About half of sample plants have air emissions from a boiler and many are burning dirty fuels such as diesel or lignite in these boilers.

Panel B reports recent plant-regulator interactions for the year prior to the study. The regulator cites plants for violating pollution standards often and follows through on those citations with costly penalties. The average plant was inspected about 1.3 times per year from January

2009 to the start of the experiment (July) in both the treatment and control groups. Despite limited contact with the regulator, many plants (15% and 18% in treatment and control) were cited for violations. Moreover, regulatory action did not stop with citations; about 20% of plants were mandated to install abatement equipment and 6% had their utilities disconnected. Across both panels we observe no statistically significant differences between treatment and control plants.

Figure 1, Panel A shows for control plants the distribution of the annual rate of inspection, rounded to the nearest inspection, over the roughly two-year course of the experiment. We observe a very right-skewed distribution where a large number of plants have fewer than 2 inspections per year, with 44% of plants inspected less than once per year. Then, above about four inspections per year, the distribution declines sharply, though some plants are visited nine times a year. This disparity is striking, given that all sample plants were selected for high pollution potential.

III Results: Experimental Estimates

This section reports on how experimental variation in inspection frequency affects regulatory actions and compliance.

A Regulatory Action

Equation (1) gives the basic estimating equation. Some specifications, noted below, also include regressors for the audit treatment and the interactions of the audit and inspection treatments. Each row in Table 3 considers a different regulatory outcome, and as before columns 1 and 2 report the means for control and treatment plants and column 3 the coefficient on the inspection treatment dummy. Panel A shows the rates of inspection in the treatment and control groups. Control plants were inspected an average of 1.40 times per year over the course of the experiment. Treatment plants were assigned to be inspected 2.12 more times per year and actually inspected an additional 1.84 times per year, more than doubling the rate of inspection, to 3.24 times per year. Each assigned inspection yields 0.9 more actual inspections implying no significant net substitution of GPCB's own inspections away from treatment. This is not a mechanical implication of the experimental design since inspection teams could well have induced

crowd-in and/or crowd-out of regular GPCB actions.¹¹

Panel B reports perceived inspection frequency. Though not officially told about increased inspection frequency, treatment plants recognized the change. When asked to recall how many inspections they received in a given year, both treatment and control plants overstated the number of inspections they got. Treatment plants, though, recalled being inspected a significant 0.71 times more than control plants in 2010, which is correct in sign but understates by 58% the actual difference in inspection rates. There was no difference in perceived inspections for years prior to the experiment (not shown).

It is instructive to compare the distribution of annual inspection rates in the treatment group in Figure 1, Panel B with the control group's distribution in Panel A. The left side of the treatment distribution is shifted to the right; the mass bunched against zero in the control is instead in the range from 2-3 inspections per year. Only 8% of plants in the treatment have less than one inspection per year during the experiment, a reduction of 36 percentage points from the control. The treatment shifts mass up the distribution, but the skinny right tails of the two distributions still look similar. Plants with this high rate of inspection are typically severe violators and, regardless of the treatment, are apparently in continual contact with the regulator, which inspects them repeatedly to gather additional evidence, to check whether they have taken measures to reduce pollution emissions, or to verify that a punishment has been imposed.

The treatment-induced increase in inspections led to more detected pollution violations and regulatory citations against plants, but no more of the really costly regulatory penalties. Table 3, Panel C examines the number of regulatory actions against sample plants. Moving down the table, the regulatory actions are ordered by increasing severity, from pollution readings, citations and warnings through to actions like utility disconnections that have a direct cost to plants. Treatment plants are a significant 0.21 more likely to have a pollution reading collected over the nearly two-year treatment, on a meager 0.38 base in the control. These readings lead directly to more treatment plants being found in violation of a standard (0.22) and a greater number of citations (0.20) for these violations in the treatment. Given the treatment duration of about two years, this means about one in 10 treatment plants receives an additional citation each year,

¹¹We test for crowd-out (the regulator shifting its own routine inspections away from the treatment group), and crowd-in (treatment inspections generating GPCB revisits to treatment plants). Online Appendix, Table B4 shows that these effects are small and statistically insignificant; hence the result in Table 3 that the net inspections were nearly equal to inspections assigned by the treatment.

more than doubling the citation rate in the control.

Moving down to more severe regulatory actions, the differences between treatment and control plants begin to taper off. More citations lead to a statistically significant increase of 0.08 closure warnings, wherein the regulator formally threatens a plant with closure unless some action is taken. However, closure directions are only insignificantly higher in treatment (coefficient 0.041, standard error 0.033). Other costly actions, such as the mandated installation of equipment or disconnection of utilities, are more common in the treatment but by very small and statistically insignificant amounts.

These costly penalties were not spared because treatment plants increased pollution abatement efforts. Table 3, Panel D shows that 61% of plants in the control installed some piece of abatement equipment over the experiment, with the average control plant installing 1.33 pieces of equipment and spending somewhat under US\$ 22,000 to do so. Treatment and control plants do not differ significantly on any of these measures of pollution abatement or numerous others we collected to measure abatement expenditures during the experiment.

We probe the finding that treatment plants were not penalized more, despite getting more inspections and citations and not substantially cleaning up. We observed in Duflo et al. (2013) (Figure V, pg. 1538) that the GPCB reserves its harshest penalties mainly for plants that exceed pollution standards by a factor of five or more. Figure 2 shows the number of plants in the treatment and control groups with at least one pollution reading taken and that exceed the relevant pollutant-specific standard by varying amounts (during the first year of the experiment, in order to minimize changes due to the treatment and focus on regulatory selection of plants). In the treatment group, the regulator takes many more pollution readings and finds many more plants, roughly twice as many, with pollution readings above the standard and above twice the standard. However, the striking finding is that the additional inspections and pollution readings in the treatment fail to uncover any additional plants with pollution readings above $5\bar{p}$, above which penalties become much more severe, or above $10\bar{p}$. There are 35 (10) plants in the treatment group with a pollution reading that is more than $5\bar{p}$ ($10\bar{p}$), compared to 33 (12) in the control group. As is evident in the figure, these rates are practically identical, and the null hypotheses that the probabilities of detection of plants with readings $> 5\bar{p}$ and $> 10\bar{p}$ do not differ by treatment status cannot be rejected (p-values: 0.80 and 0.67, respectively). This

suggests that the regulator is largely aware of which plants are violating its *de facto* standard, even with the limited resources available for inspections in the status quo.

B Pollution and Compliance

In Table 4 we report 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. All specifications include region fixed effects and an indicator for whether a plant is audit-eligible.

The treatment is estimated to have reduced plant pollution concentrations (column 1) by a statistically insignificant amount, with a coefficient of -0.105 standard deviations (standard error 0.084 standard deviations, p-value: 0.21). This effect is about half the size of the statistically significant -0.187 standard deviation reduction in pollution due to the audit treatment.¹² The audit-by-inspection interaction is large and positive, offsetting the reductions in pollution from the main effects of audits and inspections. The magnitude of the interaction is puzzling. One explanation is that the regulator is most focused on reducing emissions at plants that greatly exceed the regulatory standard and there may not be enough of these plants in both the audit and inspection treatment groups. An alternative explanation is that audits fail to produce any additional information beyond what is uncovered in an inspection; indeed, the null that the sum of the interaction and audit coefficients is zero cannot be rejected (p-value: 0.41).

Column 2 repeats the column 1 specification, except that the outcome variable is compliance, measured as an indicator for whether a plant-by-pollutant reading is below the pollutant-specific standard. The inspection treatment marginally increased compliance with pollution standards: treatment plants are 3.6 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.¹³ The audit treatment had a positive but smaller and statistically insignificant effect on compliance.

The Online Appendix tests the robustness of the effect of the inspection treatment on compliance. Online Appendix Table B5 runs placebo checks where compliance is coded to occur

¹²The audit treatment effect on pollution reported in Duflo et al. (2013) was estimated in the inspection control group only and slightly larger.

¹³Note that this rate of compliance is at the plant-by-pollutant level, but multiple pollutants are observed; the rate of compliance is a scant 10% when measured as the share of plants compliant on all pollution readings taken.

at various multiples of the real standard ($2\bar{p}$, $5\bar{p}$, $10\bar{p}$), and finds that the effect of inspection treatment on compliance is only statistically significant at the true standard.¹⁴

Compliance may increase without a large reduction in average pollution if plants near the standard were the most likely to respond to the inspection treatment. Figure 3 tests this proposition by plotting 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 4, column 2 but with finer bins than just a single dummy for compliance). Note that these pollution readings, from the endline survey, are not used for regulation and were collected by independent teams without knowledge of treatment status. The inspection treatment is estimated to reduce pollution readings just above the standard, in the range of $(0.0, 0.2]$ standard deviations, more than in any other bin, though this decrease is not statistically significant (p-value: 0.17). However, the treatment does significantly increase the number of readings just below the standard in $[-0.2, 0.0]$. It is evident that the treatment and its increased regulatory scrutiny shifted some plants that were modestly out of compliance with the *de jure* standard into compliance.

We note that the finding here that inspections change compliance but not average pollution is the reverse of our finding that more truthful environmental audits induced a reduction in pollution from highly polluting plants only, with no effect on compliance (Duflo et al., 2013). There are several reasons why the audit and inspection interventions may have differed in this way. Audits already focus on a set of audit-eligible firms that have especially high pollution potential, whereas inspections increased frequency for a broader sample of plants across the board. Additionally, audits may have offered new information on what plants were extreme polluters, since these plants were frequently inspected by the government but had never been rigorously audited by a private firm. By contrast, randomly allocated inspections may not have affected scrutiny for extreme plants, both because the inspections are less likely to be assigned to extreme plants and because any added inspections of extreme plants through the same mechanism give the regulator less new information.

¹⁴Online Appendix Table B6 tests whether the letter treatment, of randomly sending letters reminding firms of their compliance obligations, had any effect on emissions and compliance. We cannot reject that the letter treatment had no effect on either treatment or compliance, and the effect of inspections on compliance is somewhat smaller and statistically insignificant in specifications that include the letter treatment and the interaction between the letter and inspection treatments.

These findings raise some tension with respect to the regulatory status quo. The treatment doubled inspection rates and led to much higher rates of detected violations, citations and warnings. However, treatment plants were no more likely to be subjected to costly regulatory penalties, nor did they install more pollution abatement equipment. Further, the treatment caused only modest changes in pollution emissions, for plants already close to compliance. These mixed findings seem surprising because the regulator is powerful and can act forcefully; for example in the status quo, many plants *are* punished through mandated installations of pollution abatement equipment and forced closures. The next section lays out and estimates a dynamic model of regulator and plant behavior with the broad aim of understanding the source of these results.

IV A Dynamic Model of Plant and Regulator Behavior

Inspections begin a process that may result in the threat or use of costly regulatory penalties and abatement. We model this process to answer two specific questions. First, what are the full costs of environmental regulation for a plant, measured not only by the costs of pollution abatement but also by the monetary value of regulatory penalties, like plant closures and utility disconnections, in terms of lost profits that must be inferred from plants' revealed behavior? The former cost can be directly observed, while we use the model to estimate the latter. Second, what is the value of regulatory discretion in choosing where to assign inspections, relative to randomly assigned inspections?

We estimate a dynamic discrete choice structural model. At its core, the model involves alternating rounds where, as described above, the regulator can choose among the four potential actions of inspect, warn, punish, or accept and the plant at each round chooses whether to ignore the regulator or comply. Table 5 summarizes the structure of the chained interactions between the regulator and plants across rounds. Columns 1 through 7 give the frequency of actions of the regulator or plant in that round, from regulatory records, and column 8 gives the total number of observations in that round. All chains begin with a regulatory inspection. The players then alternate moves in a chain until the regulator decides to *Accept* the plant's compliance for the time being, which terminates the chain. We pool treatment and control plants to estimate the model; as discussed in Section II, treatment and control plants are treated identically in the

process of regulation, conditional on an initial inspection.

The regulator accepts many plants' compliance at an early stage, as the rapidly descending numbers of observations in column 8 show: fully 87% of chains end after a single inspection, with the regulator accepting in the third round. There are still over nine hundred chains that continue beyond that stage, however, and a handful that go on for a dozen rounds or more as the regulator revisits violating plants and enforces actions against them. If the regulator does not immediately accept the state of the plant, it is initially more likely to issue a warning: in round three, 9.5% of actions are warnings against 2.2% punishments. Thereafter, the regulator is increasingly more likely to punish. The probability of punishment conditional on reaching a given round rises monotonically with each round from 2.2% in round 3 to 18.1% in round 9 before turning downwards, in late rounds that are seldom observed. On their side, plants are unlikely to comply at first, but grow monotonically more likely to comply, with probabilities of 0.4, 7.2, 8.9, 16.5 and 17.5 percent over the second through tenth rounds, before their compliance probability levels off. The shares of compliance and punishments are endogenous outcomes for plants that have generally chosen not to comply up to a given round. We would expect these plants to have higher abatement costs than those that comply at an early stage.

Thus status quo regulation implies that most fish slip quickly through the net, but a few big ones are ensnared and thrash about. Regulation is a complex and sometimes protracted multi-stage interaction between the regulator and polluting plants. The remainder of this section models this process as a dynamic game and presents the steps involved in the estimation of its parameters.

A Key Features of the Dynamic Game

1 . Actions and Payoffs in a Round

We now outline the basics of the dynamic game played by the plant and the regulator. The plant's objective is to minimize regulatory costs, the sum of pollution abatement costs and regulatory penalties (or, in other words, to maximize profits). The regulator's objective may involve private and social goals. We do not specify this objective, but instead estimate an empirical model for regulatory action. The main objects of interest are the parameters of the regulatory penalty function, in terms of dollars imposed on regulated plants, which are revealed

by the plants decisions about when to abate and when to risk penalties.

A game (set of chained interactions) starts with an initial inspection, whether assigned in the experiment or by the regulator. We let j index a plant, R refer to the regulator, and i index either agent. Each game round is indexed by $t = 1, 2, \dots$ and the regulator and plant alternate moves. The regulator observes in round t a maximum pollution reading across pollutants of p_{jt} . (If no inspection in round t itself, then p_{jt} is recalled as the maximum pollution reading in the most recent inspection.) Each plant j has a type κ_j , not observed by the regulator R , that represents its idiosyncratic cost of reducing pollution through the installation of pollution abatement equipment, which we assume is normally distributed.

Post inspection the regulator has a choice of four actions a_{Rt} in any round t : *Inspect*, *Warn*, *Punish* or *Accept*. Figure 4 shows these actions and their payoffs within a round for the plant, leaving the regulator's payoffs unspecified. To *Inspect* is to revisit the plant and gather another pollution reading and to *Warn* is to caution the plant that it is at risk of some regulatory action. These actions are costless to the plant but continue the game and obligate the plant to respond. The only action with a non-zero current payoff for the plant is *Punish*,

$$\pi_j(a_{Rt} = \textit{Punish}|s_t) = -h(p_{jt})$$

This imposes a cost $h(p_{jt})$ that depends on observed pollution, after which the game also continues. In practice this punishment is temporarily closing the plant down, with higher pollution potentially leading to more costly penalties due to longer closures. Lastly the regulator may *Accept* that the plant is compliant, which costs the plant nothing and ends the game. The regulator is the only player who can end the game and thus moves after any plant move.

The plant acts, with the goal of minimizing costs, after the regulator chooses any action other than *Accept*. Complying costs the plant

$$\pi_j(a_{jt} = \textit{Comply}|s_t) = -c(\gamma_j, \kappa_j)$$

to install abatement equipment, where γ_j are observable plant characteristics and κ_j an unobservable, plant-specific cost shock.¹⁵ The plant may instead *Ignore* the inspection, warning or

¹⁵Because abatement capital is observable it is used to demonstrate compliance even when an initial pollutant violation could be remedied by operational changes.

punishment, which costs nothing.

2 . State Transitions

An action today gives immediate payoffs in the round and may affect future actions and payoffs. For example, a high reading p_{jt} when the regulator moves may imply a high likelihood of future penalty and, therefore, make the plant more likely to comply today. Such effects are channeled through the state of the game, a vector s_t containing information on the history of actions and pollution readings. The state transition matrix $f(s_{t+1}|s_t, a_{it})$ gives the probability of transitioning to state s_{t+1} conditional on the present state and own action. The transition matrix is known to each player through observing play in equilibrium.

3 . Value Functions

Letting $V(s_t)$ represent a plant's value of state s_t , and $v(a_{it}|s_t)$ the choice specific utility of taking action a_{it} , the total value of each of the plant's actions when moving is given by

$$v_j(a_{jt}|s_t) = \pi_j(a_{jt}|s_t) + \beta \sum_{s_{t+1}} f(s_{t+1}|a_{jt}, s_t) \sum_{a_{R,t+1}} Pr(a_{R,t+1}|s_{t+1}) \times \left\{ \pi_j(a_{R,t+1}|s_{t+1}) + \beta \sum_{s_{t+2}} f(s_{t+2}|a_{R,t+1}, s_{t+1}) V(s_{t+2}) \right\} \quad (2)$$

The plant discounts the value of future rounds by β . The transition $f(s_{t+2}|s_{t+1})$, from the plant's point of view, contains both the regulator's action and any other change in the state before the plant moves again.

For the plant, the value of a given state equals the value of taking the action with the highest payoff:

$$V_j(s_t) = \max_{a \in A_P} v_j(a_{jt}|s_t). \quad (3)$$

In determining its move today, the plant takes into account today's payoffs and the value of future states that are likely to follow from that move.

As a repeated game with private information, this model has many possible equilibria.¹⁶

¹⁶ A fully specified equilibrium includes complete contingent strategies for the players and beliefs of the regulator

We follow the dynamic games literature in not specifying what equilibria are played, but rather making several assumptions about equilibrium play that allow us to estimate the model parameters (Aguirregabiria and Mira, 2010; Pakes et al., 2007). Specifically, we assume that all plants and the regulator play the same equilibrium in all interactions, that they have correct beliefs about their opponents' actions along the equilibrium path, and that players choose strategies to maximize their expected payoffs subject to their beliefs.

B Estimation Framework and Static Estimation of Abatement Cost and Action Choice

The goal is to develop measures of the full costs of environmental regulation by understanding how plants trade-off the costs associated with the mandated installation of pollution abatement equipment against the costs from regulatory penalties like plant closures. The abatement costs can be measured directly in our data, whereas the parameters that determine regulatory penalties are estimated by maximizing the likelihood of observing the actions the plant chose in the data. Since the actions of the plant at each round depend on the value of future states, calculating the likelihood involves several steps to build the plant's dynamic value in each round by backwards induction over future costs and regulatory penalties. This subsection walks through these steps.

1 . Abatement Costs

The first step is to specify the static cost to the plant of installing abatement equipment today. Investment in abatement equipment, a fixed cost, is taken as the cost of the plant action *Comply*, because capital equipment investment is what is documented by plants, observed by the regulator and used in their judgment as demonstrating compliance. We define pollution abatement cost as consisting of an observable and unobservable component:

$$\ln c(\kappa_i, \gamma_i) = \gamma_i' \theta_c + \kappa_i \quad (4)$$

where γ_i is a vector of observable plant characteristics, θ_c is a coefficient vector for the effects of those characteristics on abatement costs, and κ_i is a cost shock.

The parameters γ_i are estimated from a linear regression of log abatement cost investments, as to the distribution of idiosyncratic compliance costs among plants surviving after each history of actions.

measured in the endline survey, on a vector of plant characteristics θ_c .¹⁷ Table 6 presents estimates of equation 4 for plants that invested in abatement equipment during the experiment (mid-2009 to the endline survey in mid-2011).¹⁸ Looking at columns 1 and 2, plants that are medium-scale (bigger than the omitted category of small plants) and textile plants (which tend to be larger and have greater air and water emissions) have economically and statistically significantly higher abatement costs. In column 3, we replace the textile dummy with indicators for using coal as a fuel and for having a high volume of wastewater, which are associated with higher air and water abatement costs, by 0.772 and 0.473 log points, respectively. Column 4 shows that the textile dummy has no explanatory power once these more precise characteristics are included. Observing this, we use the parsimonious model in column 3 as our specification for predicting plant abatement costs.

We assume the cost shock κ_i is distributed normally conditional on γ_i , i.e. $\kappa_i|\gamma_i \sim \mathcal{N}(0, \sigma_\kappa^2)$. An empirical concern is that plants with very high unobserved abatement cost shocks may not make abatement investments, even under pressure, and in this case we would not observe what their abatement expenditures would have been had they invested. This concern is somewhat mitigated by our observing a large number of equipment installations over the course of the experiment, i.e. 717 abatement investments among 791 plants completing the endline survey.

Nevertheless, we estimate the standard deviation σ_κ of this distribution via a Heckman selection model (Heckman, 1979). Specifically, we use the same covariates and functional forms from column 3 to estimate the selection equation where the dependent variable is an indicator for the plant making some investment in pollution abatement equipment. The second-stage is then the column 3 version of equation (4), except that it also includes as a regressor the inverse Mill's ratio that is determined in the selection equation. The selection model is therefore identified from the assumption of joint normality of the selection and second-stage error terms. We recover the standard deviation of idiosyncratic log abatement costs by estimating these two stages by maximum likelihood.

In the dynamic estimation we will consider specifications which (i) assign the average abatement cost to all plants, (ii) allow costs to vary with observables, based on the column 3 estimates,

¹⁷This vector includes a constant, indicators for plant size, whether a plant uses coal or lignite as a fuel and whether a plant has more than 100,000 liters a day of wastewater. These last two variables are respectively associated with higher air pollution emissions and greater regulatory stringency for water pollution.

¹⁸The mean value of an equipment investment, not shown in the table, is US\$ 17,030

and (iii) allow costs to vary with both observables and unobservables, based on the column 5 estimates. The factors in the cost model are associated with not only higher pollution and abatement costs, but also greater plant size and revenues. Because penalties are imposed in practice by plant closure we expect plants with higher revenues face higher penalties, which implies that abatement costs will be correlated with differences between actual plant penalties and the average penalty function $h(p_{jt})$. When using plant cost estimates based on observables in the dynamic estimation, we therefore scale predicted cost by predicted plant revenue, predicted using the same set of factors as in the cost model (See Online Appendix, Table B7). With this scaling the cost estimates measure whether plants have higher or lower abatement cost per unit revenue rather than in absolute terms. This is equivalent to assuming the penalty function is proportional to predicted revenue.

2 . States

The next step is to set the state vector, the variables that the players observe and act upon. We specify the common state of the game as comprised of the pollution reading, the last two actions of the players and the game round:

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

where the subscripts in a_{i-} reference the prior action of each player. If the regulator is to move at turn t , then a_{R-} will be the regulator's prior action at $t - 2$; if the plant is to move at t then a_{R-} will have been taken at $t - 1$.¹⁹ We specify the round as entering with several dummies rather than continuously to allow the regulator to flexibly respond to the selection of plants that may occur across rounds. The plant, in addition to this common state, knows its cost of abating pollution. The state for the plant is thus $s_{jt} = s_t \cup c(\kappa_j, \gamma_j)$.

3 . State Transitions

The plant's value depends on the current state and the future states it is likely to encounter, and therefore on transitions between states. The state transition after the plant moves is wholly

¹⁹In principle the whole history of player actions could enter the state. We investigated more complex states but further lags of actions did not help predict player's actions beyond the simple state as specified above.

deterministic, because the plant affects only how its own action is recorded in the state: if it chooses to *Comply* today, then $a_{j-} = \textit{Comply}$ tomorrow. The transition after the regulator moves has a deterministic part, for the regulator’s action, and a stochastic part. The transition of the pollution state is stochastic, since the pollution readings the regulator takes vary with the sampling and operating conditions of the plant and due to noise in the pollution tests themselves.

We use a simple count estimator for the pollution state transition when the regulator moves.

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

The pollution state may transition if $a_{Rt} = \textit{Inspect}$ but otherwise remains the same. We restrict the pollution transition to depend on past pollution and the regulator’s move, but not the plant’s past moves, because the count estimator may be biased for low-probability events in finite samples, so that conditioning on more past actions will leave many cells empty (e.g., the probability of pollution transitioning from above 5 times the standard to between 1 and 2 times given that the plant complied and the regulator inspected).²⁰

4 . Conditional Action Probabilities

The heart of the model consists of players’ decisions, given the observed state. We use a multinomial logit model to estimate the conditional action probabilities, where:

$$Pr(a_{it} = a|s_{it}) = \frac{\exp(\theta_a \mathbf{q}(s_{it}))}{\sum_{a'} \exp(\theta_{a'} \mathbf{q}(s_{it}))}$$

θ_a is a vector of coefficients for each action and $\mathbf{q}(s_{it})$ is a vector of state values. In particular, we specify $\mathbf{q}(s_{it})$ to include dummies for the possible most recent actions, categorical bins for the observed pollution level p_{jt} , and dummies for the stage of the game.

Table 7 presents estimates of multinomial logit coefficients for the conditional action probabilities for both the regulator and plants. The first three columns give the estimated coefficients for the regulator’s actions, at its turns to move, and the last column the coefficients for the plant’s action *Comply*. We omit the costless actions *Accept* from the regulator’s actions and

²⁰Nonetheless, we find the count estimator preferable to smooth alternatives, such as an ordered logit model, because it is simple and because some state transition patterns appear irregular in ways that would contradict common modeling assumptions. For example, the pollution state transitions from the highest level into compliance more often than into the next-highest level.

Ignore from the plant's, so coefficients are relative to those actions. Each row represents the effect of a different component of the state on the column action choices of the players.

The regulator responds strongly to high pollution levels. The regulator is significantly more likely to *Warn* or *Punish* the plant if pollution is slightly above the standard ($1-2\bar{p}$), well above the standard ($2-5\bar{p}$) or far above the standard ($> 5\bar{p}$). The coefficients that predict *Punish* as an action are monotonically and steeply increasing in pollution (column 3).

The regulator and plant's past actions also matter for the regulator's current actions. The regulator is much *less* likely to *Warn* or *Punish* if it has warned before. On the regulator's response to the plant, consider a state in the fifth round where the lagged action of the plant is *Ignore*, of the regulator *Inspect*, and no pollution reading was taken. If the plant previously chose *Comply* instead of *Ignore*, it would cut the probability of *Punish* by 5.4 percentage points, or about a third. The plant complying has a greater effect on overall payoffs than only this immediate decrease in punishments, because it also makes the regulator less likely to *Inspect* (by 10 pp) or *Warn* (by 8 pp). Thus plant compliance drops the likelihood of any inside action, that continues the game, and raises the probability the regulator just *Accepts* to end the game.

Figure 5 gives a sense of the magnitude of how the regulator's actions depend on pollution and how the plant responds to these actions. The figure plots the predicted probabilities of plant compliance and regulatory punishment, in pairs of bars for the five different levels of pollution in the model: null (no reading), $[0, \bar{p})$, $[\bar{p}, 2\bar{p})$, $[2\bar{p}, 5\bar{p})$ and above $5\bar{p}$. As pollution increases, the probability that the regulator penalizes and the probability the plant complies both increase. The predicted baseline probability of *Punish*, in the state described above, is 14.5%, whereas if the pollution reading goes from null to above $5\bar{p}$ this probability more than doubles, to 31%. For the same change, the plant's probability of compliance rises 6 percentage points on a base of 4.2%. Moreover, if the regulator had recently *Punished* the plant, this raises the probability of compliance by 9.6 pp, more than tripling the baseline probability. The marginal effects of moving from null pollution to either $[2\bar{p}, 5\bar{p})$ or $5\bar{p}$ on compliance and punishment are statistically significant at the five-percent level for the plant and the regulator, respectively.

Overall, the estimated action probabilities suggest that the regulator and the plant respond to past actions and pollution readings in a manner consistent with our equilibrium model of their dynamic interactions. The regulator pursues plants with pollution violations and penalizes

plants for extreme pollution readings, and plants clearly respond to states that foretell future regulatory action.

5 . Penalty Cost Function

The plant's payoffs depend on abatement costs and the penalties associated with various levels of pollution. We parameterize the pollution penalty function as

$$h(p) = \tau_0 \mathbf{1}\{\bar{p} < p < 2\bar{p}\} + \tau_1 \mathbf{1}\{2\bar{p} \leq p < 5\bar{p}\} + \tau_2 \mathbf{1}\{5\bar{p} \leq p\},$$

where \bar{p} is the legally mandated pollution threshold. Recall that $p = p_{jt}$ is the maximum pollution reading observed in the regulator's most recent inspection of the plant, or the maximum prior reading in the chain if no reading was taken on a given inspection. This functional form allows that the regulator may not only punish high polluters with a higher probability but also levy different penalties conditional on punishment. The complete vector of plant cost- and payoff-parameters is then $\theta_{uP} = \{\theta_c, \kappa_j, \tau\}$ with $\tau = \{\tau_0, \tau_1, \tau_2\}$.

6 . Identification and Calculation of Likelihood

To recap, the preceding steps provide the critical ingredients for building the likelihood. The vector θ_c allows for predicting each plant's abatement costs based on observable characteristics, σ_κ allows for the incorporation of unobserved plant abatement costs, the count estimator yields state transition probabilities, and the vector θ_a for all actions $a \in A_R$, gives predicted conditional action probabilities for the regulator. The remaining unknowns in the dynamic game are the penalty-cost vector τ and the variance σ_a of the plant action shock.

Identification of the penalty costs in the dynamic discrete choice model, given the assumptions above that a single equilibrium is played and players have rational expectations, is analogous to identification in a static discrete choice model (Rust, 1994; Aguirregabiria and Mira, 2007). Two normalizations are required to identify the model parameters. For the first normalization, we have set the payoff from *Ignore* to zero for the plant. The second normalization is applied in the penalty function: we omit penalty parameters for states when plants have no pollution reading or a pollution reading less than the standard. Nonetheless, plants are occasionally penalized at these states (0.34 and 0.51% of the time, respectively). We thus assume that the penalty actually

imposed when pollution is below the standard is zero; in effect the regulator has made a mistake and this is not costly to the plant.²¹ Note that, with two level normalizations, the variance of the plant action shock is a free parameter to be estimated.

The estimation of the dynamic game finds estimates of the parameter vector τ that maximize the likelihood of the sequences of actions in the data being taken. For each iteration of the maximization routine, model parameters and the draw κ_i for idiosyncratic costs are taken as fixed. The main difficulty in estimation is to calculate the values of different actions for the players given these parameters. For the empirical specification, we specify shocks to the utility of each action. The action-specific value of a_{it} then becomes

$$v_i(a_{it}, s_t) = \pi_i(a_{it}, s_t) + e_i(a_{it}, s_t).$$

We assume these shocks e_i are distributed identically and independently across actions with a type-I extreme value distribution.²²

We calculate action-specific values for the plant using backwards induction. We assume the game is finite and that the regulator will always accept in period $T = 35$, which is well beyond the ultimate round of $t = 19$ actually observed in the data.²³ In the finite game, we infer action-specific values using the state transitions and choice probabilities, starting at the final round.²⁴ For the estimation, we restrict the sample to all plant actions taken in round $t = 4$ and after, omitting $t = 2$ on the grounds that we believe the plant often does not have a chance to respond to the regulator in $t = 2$ before another regulatory action is observed (See discussion in Section II). We use a discount factor of 0.991 between rounds, which has been calibrated given the average round length to match the annual returns on capital for Indian firms found in Banerjee and Duflo (2014).

With the above specification of shocks and backward induction procedure, we have solved

²¹If one believed that this were costly, then one could interpret the penalty coefficients as the values of being penalized for a given pollution reading relative to a null or low pollution reading.

²²This assumption, while not needed for identification, is very common in the literature on dynamic structural estimation because it leads to closed-form solutions for action probabilities (Rust, 1987).

²³Given that the probability of regulator acceptance in any given round acts like a discount factor, this assumption on the game length is conservative (varying it somewhat did not significantly affect plant values in earlier rounds).

²⁴When $t = T$, then $v_j(a_{Rt}|s_t) = 0$. At $t = T - 1$, the plant's value equals its one-period profit plus an action-specific shock, $v_j(a_{jt}|s_t) = \pi_j(a_{jt}|s_t) + e_j(a_{jt}, s_t)$. The regulator always moves as estimated in the data. At moves $t = T - 3$ and all earlier moves of the plant, the plants action-specific value is found with the empirical analogue to equation (2), where the plant's profit in a given round depends on the parameters θ_{uP} .

for the values of each state for the plant, conditional on a given set of parameters. Estimation involves selecting parameters to maximize the probability of the choices observed in the data. The likelihood for the plant is

$$\mathcal{L}_j(\theta_j) = \prod_n \sum_d \omega(\theta_{jd}) \prod_{t=1}^{t=T_{jn}} Pr(a_{jnt} | s_{jnt}, \theta_{jd}).$$

Where ω are weights representing the distribution of κ . We integrate over this likelihood using Gauss-Hermite quadrature. We keep the distribution of weights fixed according to our preliminary estimation of σ_κ , the standard deviation of idiosyncratic cost shocks. Below, we describe the results from this estimation procedure.

V Results: The Full Costs of Regulation and the Value of Discretion

We begin by describing the estimates of the plant payoff parameters. Next, we use these estimates to obtain the full costs of environmental regulation for a plant (i.e., the sum of pollution abatement expenditures and regulatory penalties, and examine how this cost varies with a plant's initial level of pollution. Finally, to place a value on regulatory discretion, we exploit the experimental variation in inspection frequency to compare the costs imposed on plants by randomly assigned inspections in the treatment and inspections assigned at the regulator's discretion in the control.

A Estimates of Penalty Parameters

Table 8 presents estimates for the regulatory penalty function, $h(p)$, and the standard deviation of the unobserved action shock in the plant's action choice. These estimates provide a monetary estimate of the value of the GPCB's mandated plant closings, utility disconnections, and other penalties that would not be observable without the estimation of the dynamic game. The three columns of the table give the penalty parameters under different abatement cost assumptions: when all plants have the mean abatement cost (column 1), using the cost model based on observed characteristics (column 2) and both using the cost model and integrating over the distribution of unobserved cost shocks (column 3). Units of all parameters are thousands of

US dollars.

Inference is by clustered bootstrap over 100 bootstrap samples with replacement, with clusters at the plant-by-chain level (i.e., a series of regulatory interactions is drawn together).²⁵ We report the standard deviation and 90% confidence intervals of the coefficient estimate across replications.

In column 1, we find that a regulatory penalty imposed when observed pollution is slightly above $(1-2\bar{p})$ the standard costs US\$ 40,000. The point estimate for the penalty for pollution well above the standard $(2-5\bar{p})$ is slightly higher (US\$ 54,000), and for pollution far above the standard $(> 5\bar{p})$ a bit smaller (US\$ 41,000). In all cases, the estimates are significantly different from zero but, because the estimates are fairly imprecise, we cannot reject that the penalty function is flat with respect to pollution.

The model coefficients that adjust for predicted abatement cost per unit revenue are lower across the board, by about 40% for the highest two pollution categories. This makes sense: when all plants have the average cost, as in column 1, the only way to force a plant to comply is through high penalties. But when costs vary, as in column 2, some plants comply because their costs are lower than average, and the model can, therefore, match compliance rates with lower penalties.

Finally, column 3 allows for both observed cost differences and unobserved, log-normally distributed random cost shocks. The estimates from this model are US\$ 29,000 - US\$ 34,000 for *all* pollution levels, flat with respect to pollution, and more precisely estimated. The coefficients are again lower than those based on average cost. Allowing some plants to have idiosyncratically low costs, the model can rationalize the same rates of compliance with lower penalties, since compliant plants may just have had low cost draws. The estimated variance of the action shocks drops sharply in this specification, from over US\$ 5,000 to US\$ 2,600, because part of the variance in plant actions is attributable to idiosyncratic and persistent cost draws.

Overall, the penalty estimates suggest that penalties are flat in pollution. In light of the finding in Figure 5 that the probability of punishment is steeply increasing in observed pollution, it is evident that expected value of regulatory penalties are higher for heavy polluters due to the higher chance of being penalized at high emissions levels. The baseline penalty estimates, using average abatement cost for all plants, amount to around US\$ 40,000 per plant. If we allow

²⁵Note that on each bootstrap replication, we re-estimate not only the dynamic parameters but also the conditional choice probabilities, state transitions and discount factor that enter the dynamic estimation.

for cost-based selection into compliance, then the estimates with observed and unobserved cost differences are lower across the board, but with confidence intervals that include the baseline estimates. For simplicity, the next section reports results on the full costs of environmental regulation when using the column 1 estimates that are based on an average plant’s costs. However, the results scale about linearly with penalty cost, so it is straightforward to adjust values for the lower penalties estimated in models with heterogeneous costs.

Does the scale of estimated penalties accord with economic intuition? In column 2, the highest penalty estimated, τ_1 , is somewhat over three times the mean value of equipment abatement cost. This ratio of penalties to costs seems reasonable given that penalties must meet or exceed costs to induce abatement, and that penalties are applied infrequently, even for violating plants. Recall from Figure 5 that a plant with an extremely high pollution reading has about a 1/3 chance of being punished; this implies that the expected value of penalties at this state is about equal to the average abatement cost. Bringing in data from our endline survey, the mean plant annual sales are US\$ 2,891,500. The typical GPCB penalty for severe pollution is to close a plant by ordering its utilities disconnected; these disconnections have a median duration of 24.5 days. If we assume a 10% profit margin for plants and that profit loss is proportional to closure (i.e., there is no substitution across periods), then this implies a profit loss of US\$ 20 thousand, which lies below the estimates from the column 1 model and about in line with the average estimate across pollution levels of the model in columns 2. The penalty estimates thus are reasonable, whether approached from the cost data used to estimate the model or from independent data on plant revenues and penalty duration.

B The Full Costs of Environmental Regulation

The value of the game to a plant summarizes the costs of environmental regulation. Figure 6 shows this value, at a variety of states, as calculated through backward induction given the estimated costs of regulatory penalties and the mean cost of abatement. At each state, values are divided between expected discounted future abatement costs (light grey) and expected discounted future regulatory penalties (dark grey). The figure shows three different dimensions of the state: the time dimension is shown across panels, the pollution dimension is shown across clusters of bars within a panel, and the dimension of regulatory action across bars within a cluster.

Panel A shows the expected discounted value for the plant when it can first act in $t = 2$. The value is shown for the regulator's lagged action *Inspect*, since this is the only action the regulator can take in $t = 1$ by construction. There is a more than five-fold difference in value depending on the pollution state; the value ranges from negative US\$ 1,160, if the regulator did not take a pollution reading, down to negative US\$ 6,240 if the regulator found a pollution reading more than five times the standard. Moreover, the composition of this value also changes. For pollution readings below the standard, only 9% of the value is due to expected penalties, because penalties are very unlikely. For higher pollution values this rises steadily until penalties account for 61% of the value for extreme polluters with $p > 5\bar{p}$. It is apparent that a measure of the costs of environmental regulation that does not account for the monetary value of regulatory penalties would be greatly understated, and differentially understated for more polluting plants.

States later in the game, when the regulator is more likely to act, have sharply lower valuations. Figure 6, Panel B shows the plant value at $t = 6$. Punishment is now much more likely as is the prospect of the plant being forced into compliance and the installation of costly abatement equipment. Within each pollution state, the regulator having punished in the last round is associated with lower valuations, presumably because it reveals plants that are unobservably dirty in a way not captured by pollution readings, and thus more likely to be punished again (note these values are forward-looking and so do not include the value of any punishment already incurred). By contrast, if the regulator has recently *Warned* the plant, then this is associated with a higher (less negative) value, indicating the regulator is less likely to take severe action in the future. Again, the share of the value due to penalty costs is sharply increasing in pollution. A plant in the highest pollution state, for example (rightmost cluster, weighting bars by their relative probabilities), expects 67% of its US\$ -23,248 future value to be due to penalties. The total value to the plant, given the penalty estimates, equates to about half of the penalty conditional on certain punishment.

What can these estimates tell us about the experimental results? First, despite a flat value of the penalty function in pollution, the higher probability of penalty and of continuation in the game gives starkly more negative values of regulation for plants with high pollution levels. Hence, uncovering a large number of minor violators would not be expected to induce abatement, because these minor violators do not expect to incur penalties. Second, the kinds of citations

actually issued in response to the treatment, which are classed as warnings, are associated with higher (less negative) values and lower abatement cost. These estimates thus reconcile the fact that the experiment generated additional violations, or warnings, but induced abatement only near the regulatory standard—warning carries little weight and in fact may signal that plants do not need to comply.

C Interpreting Regulatory Costs as the Value of Discretion

This section combines the full costs of regulation, as estimated from the model, with experimental variation in inspections to compare the costs of regulation for treatment and control plants. We discuss the costs imposed on plants in these different groups and then unpack under what assumptions differences in the cost of regulation to plants measure the value of regulatory discretion.

Table 9 presents estimates of regulatory costs by treatment status. Total costs include the regulatory penalty costs estimated from the dynamic model of Table 8 and abatement costs. These estimated costs are then applied to the pollution states and penalty and abatement actions observed in the data for the treatment and control groups. The costs are summed up across all plants and divided by the total number of regulatory inspections to present costs per inspection for each group. Columns 1 and 2 show the treatment and control groups and column 3 reports the difference between the ratios and the associated standard error in parentheses.

It is evident that inspections in the treatment group, which are far more likely to be randomly assigned²⁶, are less likely to induce costs for plants. The total costs of environmental regulation per treatment plant inspection are about US\$ 1,340, compared to roughly US\$ 2,650 for control plant inspections. Both cost components are higher per inspection in the control—plants are penalized more often, at a value of about US\$ 960 per inspection (versus US\$ 620 in the treatment) and install US\$ 1,690 of equipment per inspection as opposed to US\$ 710 in the treatment.

In principle, the costs imposed by control inspections may be higher because the regulator targets more of these inspections or because there are diminishing marginal returns to additional

²⁶GPCB carried on its own routine inspections, which are meant to be more or less random but are assigned at officials' discretion, in both groups. Thus the random share of inspections in the treatment includes the 67% of truly random inspections assigned by the treatment, plus any routine inspections, whereas the random share of inspections in the control is the routine inspections alone. Before the experiment 35% of inspections in the sample were routine.

inspections in the treatment. These are not exclusive, because as shown in the discussion of the status quo, some plants are targeted and visited often, which may lead to diminishing returns to inspections for those plants. One observation in favor of a targeting view is that the ratio of abatement costs to penalty costs is higher in the control: penalties and abatement cost per inspection are about the same in the treatment, but the abatement cost is nearly twice as high as the penalty cost in the control, suggesting that the regulator is able to uncover chances to abate and compel plants to take them.

Under what assumptions does this difference in costs imposed measure the value of regulatory discretion or targeting? The penalty cost estimates are estimated from observed abatement investments and the behavior of the regulator and plants in the data, as far as when to penalize and when to abate. We believe that it is valid to interpret total regulatory costs per inspection as the value of an inspection *according to* the regulator's own objectives or preferences. There are two important caveats to this interpretation.

First, some kinds of rent-seeking would cause our model to be mis-specified and the penalty cost estimates to differ from the objectives of the regulator. The penalty estimates are robust in the sense that they will include side payments, if such payments occur when the regulator chooses to punish the plant, e.g., if a punishment is applied and then removed when a payment is made. This robustness is due to penalty costs being identified from plants' willingness to make abatement investments to avoid the punishment stage, and all costs therein, whether they are transfer payments or reductions in profits from closure. If costs were being imposed at other nodes of the game tree, such as plants having to pay an inspection team whenever they show up, these costs would be a negative payoff to environmental regulation omitted by our model. The effect such costs would have on our estimates of regulatory penalties is hard to forecast, as it would depend on precisely in which states the costs occurred.

Second, we do not take the difference in costs per inspection as a measure of social benefits or welfare. Such an interpretation would require that the regulator imposes costs on plants equal to the change these costs induce in the external social damages of plant pollution. A full cost-benefit analysis of the value of inspections and penalties would therefore require estimating monetary damages of pollution in Gujarat, elasticities of plant pollution with respect to penalties and any threat or deterrence effects from these penalties. It would be heroic to expect even a regulator

focused only on social costs to optimally target and penalize plants. Moreover, any kind of rent-seeking would break the connection between the costs of regulation and social welfare or damages. To put it another way, the penalty function we observe looks rational given resource constraints, but we have no evidence that it is optimal.

VI Conclusion

This paper reports the results of an experiment to increase the frequency of inspection of polluting plants in collaboration with the environmental regulator in Gujarat, India. Using detailed data on regulatory interactions with plants and independent pollution readings, we are able to follow this increase in regulatory scrutiny through to the most important regulatory outcomes and assign a comprehensive value to regulatory inspections.

The treatment was successful when judged by process measures, since treatment plants were more than twice as likely as control plants to be inspected and to be cited for violating pollution standards. However, the treatment was much less successful when judged by more consequential outcomes; the regulator was no more likely to identify extreme polluters or to impose costly penalties in the treatment group. Further, treatment plants only marginally increased compliance with standards and did not significantly reduce mean pollution emissions.

To explain these results and recover the full costs of environmental regulation, we modeled the regulatory process as a dynamic discrete game where the regulator chooses whether to penalize and plants choose whether to abate to avoid future sanctions. Using data we assembled on 10,000 interactions between plants and the regulator, the model produces estimates that indicate that the costs of environmental regulation are mainly reserved for extreme polluters. By combining the experimental and model results, we find that the average treatment inspection imposes about half the costs on plants as the average control inspection, because these randomly assigned inspections are less effective at identifying extreme polluter than the regulator's discretionary inspections in the control group. Overall, the data suggest that regulatory discretion has considerable value in this setting.

Considering lessons for the state of environmental regulation in India, we see the fisherman's net as half-full, mainly with big fish. The other half may be left empty because command-and-control regulation in the form studied here is clearly very costly. The estimates imply that a

large part of the cost of environmental regulation for plants is due not to abatement investments but to regulatory penalties, applied mainly through plants being closed down. Since the threat of closure is necessary to induce compliance, every closure helps the regulator to establish a reputation, but each instance of closure also creates deadweight loss, in the form of foregone plant production and profits. By our estimates, even in the control, inducing every dollar of actual abatement investment requires 57 cents of deadweight penalties. A full cost-benefit analysis of regulation is beyond the scope of this paper, but the large share of deadweight loss in the cost to plants would burden the cost side of the ledger. From the perspective of a social planner taking the technology of regulation as given, these losses in the regulatory process increase the costs of pollution abatement and may rationalize, to some extent, why for many plants abatement remains low and pollution high.

Yet, in the full half of the net—where the big fish are caught up—the regulator appears to be applying a rational penalty function, albeit one where the *de facto* standard lies much above the *de jure* standard. Moreover, the regulator is able to effectively target plants which are high polluters and have greater scope for improvement. We do not claim this targeting is optimal, but targeting in something like this manner appears necessary to address high pollution levels given tight resource constraints, poor information and a rigid penalty structure. Future research might consider more fundamental reforms of the command-and-control status quo, for example changes to the technology of monitoring or to the nature of incentives for compliance.

References

- Aghion, P., Burgess, R., and Redding, S. (2008). The unequal effects of liberalization: Evidence from dismantling the license Raj in India. *American Economic Review*, 98(4):1397–1412.
- Aghion, P. and Tirole, J. (1997). Real and formal authority in organizations. *Journal of Political Economy*, 105:1–29.
- Aguirregabiria, V. and Mira, P. (2007). Sequential estimation of dynamic discrete games. *Econometrica*, 75(1):1–53.
- Aguirregabiria, V. and Mira, P. (2010). Dynamic discrete choice structural models: A survey. *Journal of Econometrics*, 156(1):38–67.

- Banerjee, A. V. and Duflo, E. (2014). Do firms want to borrow more? Testing credit constraints using a directed lending program. *The Review of Economic Studies*, 81(2):572–607.
- Besley, T. and Burgess, R. (2004). Can labor regulation hinder economic performance? Evidence from India. *Quarterly Journal of Economics*, 119(1):91–134.
- Besley, T. and Persson, T. (2010). State capacity, conflict and development. *Econometrica*, 78:1–34.
- Bhushan, C., Yadav, N. K., Roy, A., , and Banerjee, S. (2009). Turnaround: Reform agenda for India’s environmental regulators. *Centre for Science and Environment*.
- Boyer, M. and Laffont, J.-J. (1999). Toward a political theory of the emergence of environmental incentive regulation. *The RAND Journal of Economics*, pages 137–157.
- Burgess, R., Hansen, M., Olken, B., Potapov, P., and Sieber., S. (2012). The political economy of deforestation in the tropics. *Quarterly Journal of Economics*, 127(4):1707–1754.
- Chen, Y., Ebenstein, A., Greenstone, M., and Li, H. (2013). Evidence on the impact of sustained exposure to air pollution on life expectancy from China’s Huai River policy. *Proceedings of the National Academy of Sciences*, 110(32):12936–12941.
- Duflo, E., Greenstone, M., Pande, R., and Ryan, N. (2013). Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India. *The Quarterly Journal of Economics*, 128(4):1499–1545.
- Estache, A. and Wren-Lewis, L. (2009). Toward a theory of regulation for developing countries: Following Jean-Jaques Laffont’s lead. *Journal of Economic Literature*, 47(3):729–770.
- Greenstone, M. and Hanna, R. (2014). Environmental regulations, air and water pollution and infant mortality in India. *NBER Working Paper 17210*.
- Greenstone, M., List, J. A., and Syverson, C. (2012). The effects of environmental regulation on the competitiveness of us manufacturing. Technical report, National Bureau of Economic Research.

- Greenstone, M., Nilekani, J., Pande, R., Ryan, N., Sudarshan, A., and Sugathan, A. (2014). Lower pollution, longer lives: Life expectancy gains if India reduced particulate matter to air-quality standards. *Mimeo, Harvard Kennedy School*.
- Hanna, R. N. and Oliva, P. (2010). The impact of inspections on plant-level air emissions. *The BE Journal of Economic Analysis & Policy*, 10(1).
- Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica*, 47(1):153–161.
- Jia, R. (2014). Pollution for promotion. *mimeo*.
- Laffont, J.-J. (1994). *Regulation of pollution with asymmetric information*. Springer.
- Laffont, J.-J. (2005). *Regulation and Development*. Cambridge University Press.
- Laffont, J.-J. and Tirole, J. (1993). *A Theory of Incentives in Procurement and Regulation*. MIT Press.
- Leaver, C. (2009). Bureaucratic minimal squawk behavior: Theory and evidence from regulatory agencies. *American Economic Review*, 99(3):572–607.
- Magat, W. A. and Viscusi, W. K. (1990). Effectiveness of the epa’s regulatory enforcement: The case of industrial effluent standards. *JL & Econ.*, 33:331.
- Ministry of Environment and Forests (1986). The environment (protection) rules, 1986. *Gazette of India*, 844(E).
- Pakes, A., Ostrovsky, M., and Berry, S. (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics*, 38(2):373–399.
- Rasul, I. and Rogger, D. (2013). Management of bureaucrats and public service delivery: Evidence from the Nigerian civil service. *Mimeo, UCL*.
- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033.

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

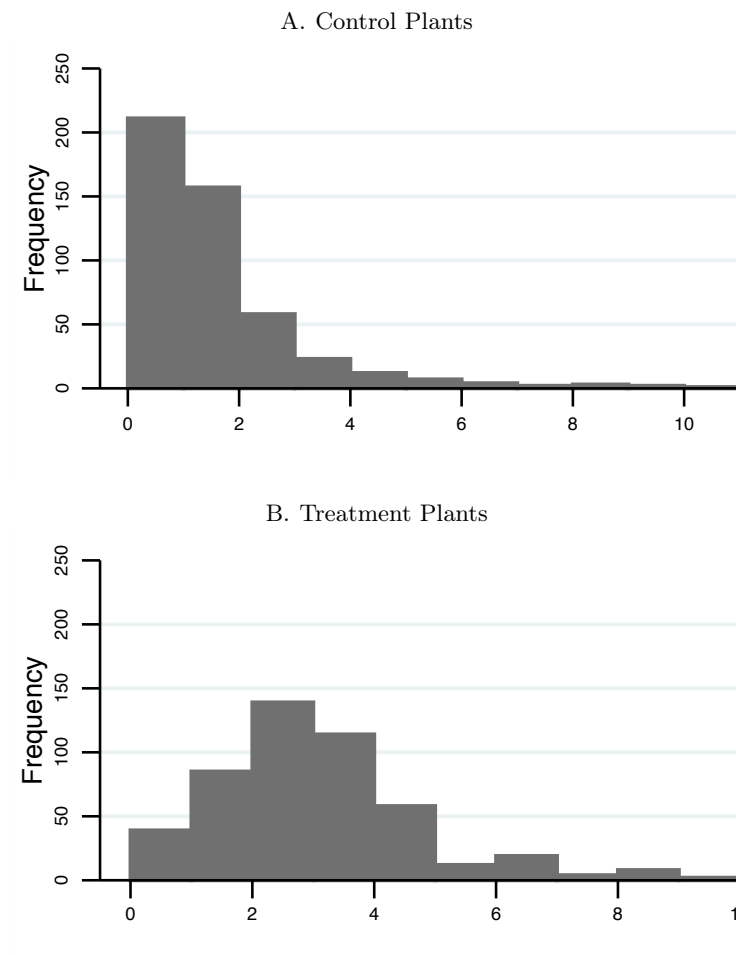
Sitkin, S. and Bies, R. (1994). *The Legalistic Organization*. Sage Publications.

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

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


VII Figures

Figure 1: Distribution of Inspection Rates by Treatment Status



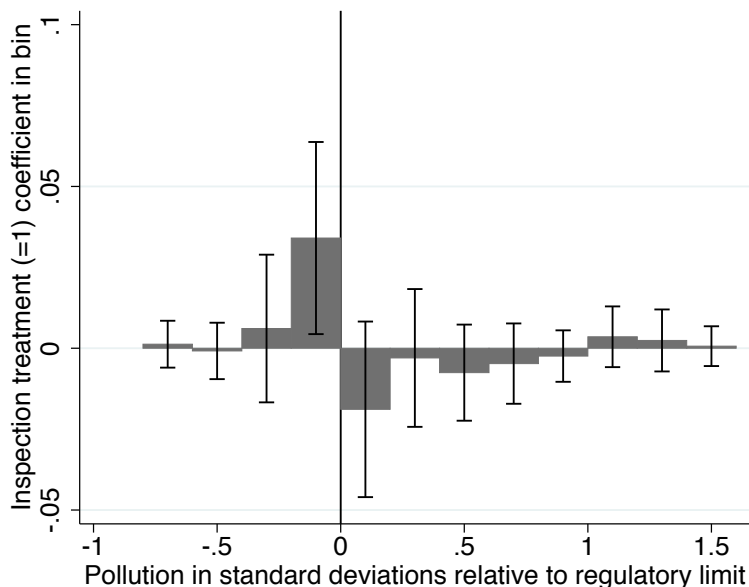
The figure shows the distribution of the annual inspection rate (i.e., inspections per year) over the approximately two years of the experiment by treatment status. The inspection rate is pro rated for the actual duration of the treatment in each of three regions. Data on inspections comes from administrative records of inspection reports filed with the regulator.

Figure 2: Regulatory Targeting of Extreme Polluters



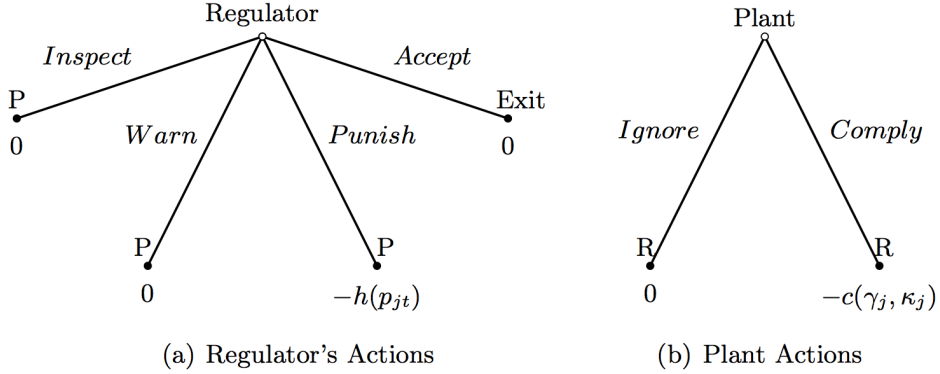
The figure shows the number of plants with pollution readings taken and that fall in various bins, relative to the regulatory standard, during the first year of the intervention for the control and treatment groups, respectively. The first pair of bars shows the number of plants that had at least one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

Figure 3: Inspection Treatment Effect on Pollution Relative to Regulatory Standard



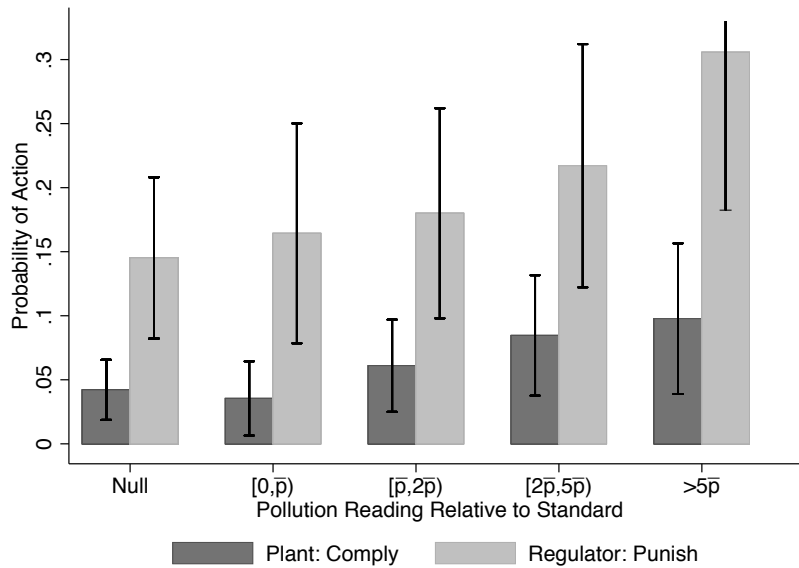
The figure reports coefficients on the inspection treatment assignment from regressions of dummies for a pollution reading being in a given bin relative to the regulatory standard on inspection treatment, audit treatment, inspection \times audit treatment, a dummy for being audit-eligible and region fixed effects. 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. Standard errors are clustered at the plant level and the whiskers show 95% confidence intervals for the inspection treatment coefficient.

Figure 4: Actions of the Players at Each Node



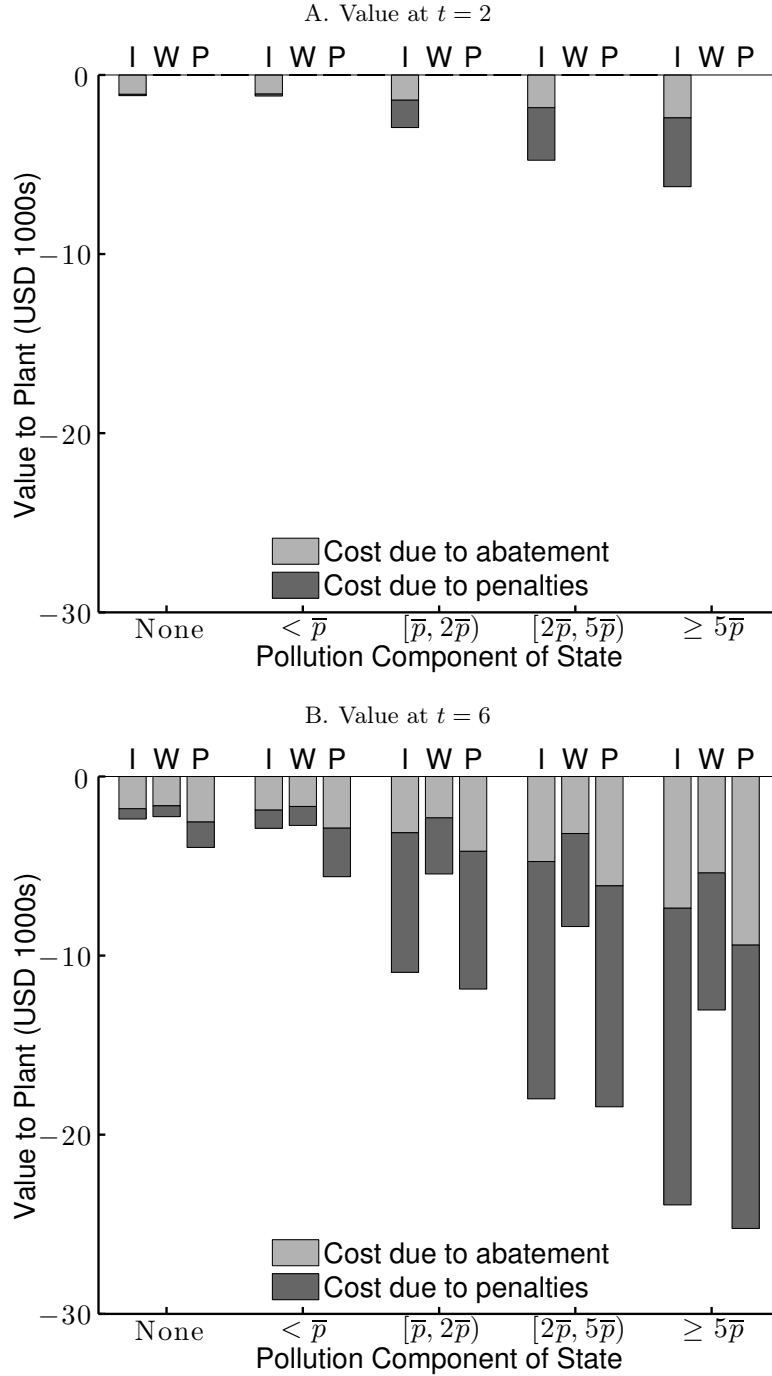
The figure gives the actions of the players at each node and the terminal nodes give the payoffs in the stage game for the plant only. The game begins with an inspection where the regulator (R) observes p_{j1} . The regulator may take four actions. If R *Inspects*, R gets a new signal and moves again. If R *Warns* or *Punishes*, the Plant moves. If R *Accepts*, the game ends. When the plant can move, he *Ignores* or *Complies*. After either action of the plant the regulator moves again.

Figure 5: Predicted Probabilities of Compliance and Punishment



The figure shows the predicted probabilities of the plant taking the action *Comply* and the regulator taking the action *Punish*, respectively within each pair of bars, as calculated in the model of Table 7 for the pollution component of the state taking on different values relative to the regulatory standard. The probabilities are evaluated in round $t = 4$ for the plant and $t = 5$ for the regulator, at a state when the regulator previously *Inspected* and the plant previously *Ignored*.

Figure 6: Total Costs of Environmental Regulation for Regulated Plants



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

VIII Tables

Table 1: Mapping of Player Actions to Raw Documents

Action	Document	Description
<i>Panel A. GPCB Actions</i>		
Inspect	Inspection report	Analysis of air and water samples; report on plant characteristics.
Warn	Letter	Non-threatening letter ordering improvement in pollutant concentrations.
	Citation	Threatening letter demanding explanation for high pollution levels, missing permit to operate, or missing pollution abatement equipment.
	Closure Notice	Notice that the plant will be ordered to close in 15 days if the plant does not take action to improve pollution.
Punish	Closure Direction	Order to close immediately.
	Utility notice	Notice that water or electricity has been disconnected.
Accept	Revocation of Closure Direction	Permission to start operation.
	Implicit	No further GPCB action.
<i>Panel B. Firm Actions</i>		
Ignore	Implicit	Implied by consecutive GPCB actions.
	Protest	Official letter of protestation to GPCB. Challenges parameter readings and other directives.
Comply	Equipment installed	Notice that pollution abatement equipment has been installed.
	Process installed	Notice that a process has been installed.
	Bond posted	Letter from bank to GPCB explaining that the firm has posted a guarantee against future misconduct.

Table 2: Inspection Treatment Covariate Balance

	Treatment (1)	Control (2)	Difference (3)
<i>Panel A. Plant Characteristics</i>			
Capital investment Rs. 50m to Rs. 100m (=1)	0.071 [0.26]	0.087 [0.28]	-0.017 (0.017)
Located in industrial estate (=1)	0.37 [0.48]	0.33 [0.47]	0.032 (0.027)
Textiles (=1)	0.45 [0.50]	0.45 [0.50]	-0.0092 (0.020)
Dyes and Intermediates (=1)	0.16 [0.36]	0.13 [0.34]	0.027 (0.022)
Effluent to common treatment (=1)	0.35 [0.48]	0.37 [0.48]	-0.021 (0.031)
Waste water generated (kl / day)	196.8 [316.4]	192.1 [310.9]	4.30 (16.2)
Lignite used as fuel (=1)	0.35 [0.48]	0.35 [0.48]	-0.0063 (0.018)
Diesel used as fuel (=1)	0.17 [0.38]	0.19 [0.39]	-0.013 (0.024)
Air emissions from boiler (=1)	0.52 [0.50]	0.50 [0.50]	0.019 (0.020)
Bag filter installed (=1)	0.13 [0.34]	0.14 [0.35]	-0.011 (0.021)
<i>Panel B. Regulatory Interactions in Year Prior to Study</i>			
Number of inspections	1.25 [1.32]	1.22 [1.32]	0.026 (0.079)
Any equipment mandated (=1)	0.21 [0.41]	0.19 [0.39]	0.021 (0.021)
Any citation issued (=1)	0.15 [0.36]	0.18 [0.39]	-0.032 (0.024)
Any water citation issued (=1)	0.069 [0.25]	0.077 [0.27]	-0.0086 (0.017)
Any air citation issued (=1)	0.010 [0.10]	0.013 [0.11]	-0.0022 (0.0068)
Any utility disconnection (=1)	0.056 [0.23]	0.063 [0.24]	-0.0064 (0.015)
Any bank guarantee posted (=1)	0.021 [0.14]	0.017 [0.13]	0.0040 (0.0087)
Observations	480	480	

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

Table 3: Regulatory Interactions During Experiment

	Treatment	Control	Difference
<i>Panel A. Inspections by Treatment Status</i>			
Number inspections assigned in treatment, annual	2.12	0	2.12***
	[0.57]	[0]	(0.026)
Total inspections, scaled and annual over treatment	3.11	1.40	1.71***
	[1.77]	[1.59]	(0.11)
Observations	480	480	
<i>Panel B. Perceived Inspections by Treatment Status</i>			
Perceived Inspections by GPCB officials, 2010	3.62	2.92	0.71***
	[1.46]	[1.58]	(0.11)
Total Perceived notices and closures received, 2010	0.30	0.27	0.025
	[0.70]	[0.64]	(0.048)
Observations	403	388	
<i>Panel C. Regulatory Actions by Treatment Status</i>			
Pollution reading ever collected at plant (=1)	0.60	0.38	0.21***
	[0.49]	[0.49]	(0.032)
Any pollution reading above limit at plant (=1)	0.55	0.34	0.22***
	[0.50]	[0.47]	(0.031)
Number of pollution readings above limit at plant	2.84	1.17	1.67***
	[3.67]	[2.58]	(0.20)
Total citations	0.35	0.15	0.20***
	[0.69]	[0.42]	(0.037)
Total water citations	0.12	0.046	0.071***
	[0.37]	[0.22]	(0.020)
Total air citations	0.042	0.021	0.021*
	[0.20]	[0.14]	(0.011)
Total closure warnings	0.17	0.094	0.077***
	[0.48]	[0.34]	(0.027)
Total closure directions	0.20	0.16	0.042
	[0.54]	[0.48]	(0.033)
Total bank guarantees	0.065	0.060	0.0042
	[0.25]	[0.27]	(0.017)
Total equipment mandates	0.040	0.027	0.013
	[0.23]	[0.19]	(0.014)
Total utility disconnections	0.042	0.040	0.0021
	[0.20]	[0.22]	(0.013)
Observations	480	480	
<i>Panel D. Equipment installations since start of experiment</i>			
Firm made at least one equipment installation	0.63	0.61	0.019
	[0.48]	[0.49]	(0.035)
Number of equipment installations made	1.39	1.30	0.090
	[1.65]	[1.61]	(0.12)
Total cost of all equipment installations (USD 1000s)	20.0	21.8	-1.79
	[89.4]	[61.0]	(5.46)
Observations	403	388	

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

Table 4: Endline Pollution and Compliance on Treatments

	Pollution (1)	Compliance (2)
Inspection treatment (=1)	-0.105 (0.0839)	0.0366* (0.0213)
Audit treatment (=1)	-0.187** (0.0849)	0.0288 (0.0258)
Audit \times inspection treatment (=1)	0.286** (0.142)	-0.0365 (0.0353)
Inspection and audit control mean	0.682	0.614
Observations	4168	4168

The table shows regressions of pollution (column 1) and compliance (column 2) on treatment assignments at the plant-by-pollutant level. Pollution consists of air and water pollution readings for each plant, taken during the endline survey, where each pollutant is standardized by dividing by its standard deviation. Compliance is a dummy for each pollutant being below its regulatory standard. The table regresses these outcomes on inspection treatment assignment, audit treatment assignment and inspection \times audit treatment. Specifications also include region fixed effects and a control for the audit sample. Standard errors clustered at the plant level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Structure of Chained Regulator-Plant Interactions

	Regulatory Action				Plant Action		N (7)	% left (8)
	Inspect (1)	Warn (2)	Punish (3)	Accept (4)	Ignore (5)	Comply (6)		
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	31.0	3.2	1.1	29.4	34.6	0.6	25217	

The table reports actions taken by the GPCB and by the firm using administrative data from the regulator. All chains start with an inspection. Action rounds within a chain alternate between GPCB action and firm action. A chain ends when the GPCB accepts, and issues no further inspections, warnings or punishments that follow up directly with the initial inspection. When a chain ends it is dropped from the table. Rounds after the 15th round are not shown and the row 15+ summarizes these rounds; 6 chains go at least 17 rounds and 4 chains go 19 rounds.

Table 6: Model of Plant Cost for Abatement Equipment

	Log Cost of Abatement Equipment, Per Piece				
	(1)	(2)	(3)	(4)	(5)
Medium size industry (=1)	0.888*** (0.230)	1.091*** (0.242)	0.827*** (0.224)	0.858*** (0.225)	0.839*** (0.231)
Textiles (=1)		0.971*** (0.128)		0.119 (0.261)	
Coal or lignite as fuel (=1)			0.772*** (0.179)	0.727*** (0.184)	0.765*** (0.220)
Waste water generated > 100 kl/day			0.473*** (0.174)	0.410 (0.250)	0.546 (0.885)
Constant	1.642*** (0.0701)	1.107*** (0.110)	1.012*** (0.104)	1.000*** (0.110)	0.943 (0.758)
Sigma					1.216 (0.0434)
R^2	0.0439	0.167	0.213	0.214	
Residual SD	1.340	1.251	1.216	1.215	1.275
Mean	1.747	1.747	1.747	1.747	1.747
Observations (plants)	398	398	398	398	912
Observations (equipment)	717	717	717	717	1231

The table presents regression estimates for the fixed cost of abatement equipment on plant characteristics. Columns 1 through 4 present OLS regressions for the set of firms undertaking any abatement investment during the experiment and column 5 reports maximum likelihood estimates of a Heckman selection model among all plants with administrative baseline data on covariates. Plant abatement equipment costs are measured in the endline survey, for all investments over the prior two years of the experiment, and specified in logged thousands of US dollars (US\$ 1 \approx INR 50). Medium size is a regulatory designation for capital stock above INR 50m (the sample includes only small- and medium-sized plants), Textiles is a dummy for belonging to the textile industry, Coal or lignite is a dummy for using one of those fuels, and Waste water generated is a dummy for discharging more than 100 kl/day of effluent. Column 5 reports the estimates of the second stage of a Heckman selection model and Sigma is the estimated variance of the error term in that stage, used to characterize the distribution of idiosyncratic cost shocks. The first, selection stage includes all of the model's independent variables and has coefficients (not reported) of: Medium size industry (=1), 0.504 (0.327); Coal or lignite as fuel (=1), -0.292 (0.213); and Waste water generated > 100 kl/day, 3.136*** (0.363). Standard errors clustered by plant in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Multinomial Logit Model of Action Choice Conditional on State

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

The table reports coefficients from multinomial logit models for the action choice probabilities of each player conditional on the state of the game. See Table 1 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 player taking the specified column action relative to the omitted action. Pollution readings are taken during inspections throughout the treatment period. The omitted pollution reading is null, which occurs when the regulator inspects but does not take a pollution reading. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Estimates of Plant Utility Parameters (US\$ 1000s)

	Assumption on Cost		
	Mean Cost (1)	Cost Model Prediction (2)	Prediction and Cost Shock (3)
<i>Parameters of penalty function</i>			
$h(p) = \tau_0 \mathbf{1}\{\bar{p} < p < 2\bar{p}\} + \tau_1 \mathbf{1}\{2\bar{p} \leq p < 5\bar{p}\} + \tau_2 \mathbf{1}\{5\bar{p} \leq p\}$			
τ_0	40.03 (27.51) [11,81]	10.46 (35.94) [-30,43]	31.25 (21.57) [6,73]
τ_1	54.23 (28.04) [18,110]	34.04 (25.66) [-2,89]	34.15 (18.14) [12,71]
τ_2	41.41 (18.93) [20,81]	23.77 (13.69) [8,50]	29.11 (12.04) [12,51]
<i>Standard deviation of action shock</i>			
σ	5.88 (0.29) [5,6]	5.29 (0.31) [5,6]	2.60 (0.24) [2,3]
Observations	1474	1474	1474

The table presents pseudo-maximum likelihood estimates of the parameters of the plant profit function, from the dynamic discrete game. The three 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 three columns represent different assumptions on the costs of abatement. Column (1) uses the mean abatement capital cost for installing a piece of equipment, measured in the endline survey, column (2) uses the predicted value of that installation based on plant observable characteristics and column (3) uses that prediction and integrates over the distribution of unobserved shocks to installation costs. Observations are those at which the plant moves in rounds $t=4$ and onwards; $t=2$ is omitted because a large number of actions in that round are imputed (see text). Inference is by the bootstrap over 100 samples with replacement, where samples are taken at the level of the plant-chain (i.e., series of interactions) stratified on the maximum pollution reading observed in the chain. Standard errors equal to the standard deviation of bootstrap estimates are in parentheses and 90% confidence intervals giving the 5th and 95th percentiles of the bootstrap estimates are in brackets.

Table 9: Cost of Environmental Regulation to Plants Per Inspection (US\$ 1000s)

	Treatment (1)	Control (2)	Difference (3)
Total cost per inspection	1.34 [0.18]	2.65 [0.44]	-1.31*** (0.02)
Regulatory penalty cost	0.62 [0.12]	0.96 [0.24]	-0.34*** (0.01)
Number of penalties	0.030 [0.00]	0.050 [0.01]	-0.020*** (0.00)
Abatement equipment cost	0.71 [0.14]	1.69 [0.39]	-0.97*** (0.02)
Number of installations	0.066 [0.01]	0.097 [0.02]	-0.032*** (0.00)
Observations	480	480	

The table tests for difference in the expected costs due to penalties and equipment installation requirements as result of a GPCB inspection. Includes firms with at least one recorded inspection. Firm penalty costs are estimated using firm pollution levels in the round where the GPCB issues a punishment. Firm expected equipment costs per inspection are estimated using total observed equipment costs, controlling for plant characteristics. Costs and inspections are aggregated across all plants by treatment and control groups before the cost per inspection is calculated. The cost and number of penalties are recorded in the administrative data from the regulator during the treatment period. The cost and number of equipment installations are recorded in the endline survey. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A Data Appendix

We clean administrative records on interactions between GPCB and sample plants and map the documents to the actions of the dynamic inspection game. Cleaning involves linking documents that are related and putting sets of linked documents into separate chains of related actions. Mapping involves reducing the many possibly signals between the regulator and the plant to the basic purpose of each action.

The documents are mostly letters between GPCB and the plants of various kinds. The Actions in Table 1 represent groupings of underlying documents based on their purpose. Sometimes these correspond directly to a single underlying type of document in the raw data, such as for inspections, and sometimes they represent a category of documents. On the regulatory side, for GPCB, there are four Actions and so four categories of underlying document. The broadest class is for the action *Warn*, which includes a progression from simple letters requesting explanation to closure notices that threaten the plant with closure if it does not reduce pollution. We group these as a single category because, while they represent a range of threats, they impose no immediate costs on the plants. The plant has two possible actions. We infer when the plant has ignored the regulator from written replies that do not have any direct cost or by imputing *Ignore* from other moves in the game (as described below). The plant can comply through installing abatement equipment, installing a new process or posting a bond against future performance (which is rare). These actions are documented in correspondence from plants with verification from third-party equipment vendors.

We link actions into chains, or sequences of actions, based on their underlying documents and references therein. This linking is done through explicit references, dates and some imputations. Some documents explicitly cite a preceding document by date—for example, a closure direction will reference high readings from a pollution analysis report a month before, giving the date and document number of that report. Explicit links are much more likely for documents from the GPCB, which follows formal procedures for citation, than from plants, which often do not cite or cite inaccurate dates. To account for this, we link plant documents with GPCB documents by allowing the date on the plant letter to differ from the reference date by up to one digit.

Even documents that do not explicitly cite others are dated. We use document dating and knowledge of the typical regulatory process to expand chains of documents through a set of

logical rules. The full set of rules for chain construction is given in the list below. The rules are designed to match links in a chain that are both logically appropriate and within the typical range of time lags observed between actions that do have explicit references.

Chain Construction Rules

1. *Explicit matching.*
 - (a) Match documents if document explicitly cites (by document number and/or date) an inspection or other earlier document.
 - (b) If no match, match documents by explicit but fuzzy match that allows dates to differ by one digit.
2. *Date matching.* Order remaining documents, which have been grouped into chain fragments by the explicit matching process, by the date of document within plant. Match chain fragments for each plant if the following conditions on the documents and relative dating are met:
 - (a) GPCB inspection followed by plant reply or protest (within 30 days).
 - (b) GPCB inspection followed by GPCB action (letter, show cause notice, closure notice, closure direction) (within 60 days).
 - (c) GPCB action followed by inspection (within 30 days).
 - (d) GPCB action followed by plant reply or protest (within 30 days).
 - (e) GPCB punishment (closure notice, closure notice, utility confirmation of action) followed by plant compliance (equipment installation, process installation, bank guarantee posted) (within 30 days).
 - (f) Plant compliance followed by GPCB acceptance (revocation of closure direction) (within 60 days).
 - (g) GPCB inspection followed by GPCB acceptance (within 60 days).
3. *Impose logical rules.* Using the raw matching based on explicit references and dates, apply logical rules by imputing actions or truncating chains to enforce the structure of the game.

- (a) *Chains end with an acceptance.* Split chains at any internal acceptance. Impute GPCB *Accept* at chain end if no further follow-up actions.
- (b) *Chains start with an inspection.* Search for available inspections, append to the beginning of chain without an inspection if within 1 month of chain start, else truncate chain before first inspection.
- (c) *Chains alternate moves.* Impute plant move *Ignore* between any consecutive GPCB actions.

The chains of actions constructed in steps (1.) and (2.) do not always adhere to the logical rules of the game we use to characterize the data. In the final stage of data cleaning, we therefore combine or impute actions where appropriate to cast all actions into coherent moves in the game. The main assumptions in imputing moves to complete the game are that the plant had an opportunity to respond to all regulatory actions with compliance, even if it in fact did not comply, and that similarly when regulator had an opportunity to continue pursuing plants at the end of each chain but did not this is equivalent to choosing *Accept*.

Online Appendix for:

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

Esther Duflo*
Michael Greenstone†
Rohini Pande‡
Nicholas Ryan§

September 29, 2014

*MIT, eduflo@mit.edu

†MIT, mgreenst@uchicago.edu

‡Harvard, rohini_pande@harvard.edu

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

I Tables

Table B1: Attrition in the Endline Survey

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

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

Table B2: Endline Attrition by Inspection Treatment Status

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

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

Table B3: Sample Chain

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

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

Table B4: Annual Number of Inspections on Inspection Treatment

	Treatment (1)	Control (2)	Difference (3)
Inspections	2.98 [1.72]	1.40 [1.60]	1.59*** (0.096)
Inspections, as part of treatment	1.62 [0.74]	0 [0]	1.62*** (0.031)
Inspections, not part of treatment	1.37 [1.36]	1.40 [1.60]	-0.029 (0.089)
Inspections, not part of treatment, first day of chain	1.21 [1.18]	1.28 [1.38]	-0.062 (0.077)
Inspections, not part of treatment, after first day of chain	0.12 [0.35]	0.088 [0.34]	0.035 (0.022)
Observations	480	480	

The table shows the differences in the number of different types of inspections between the treatment and control groups during the treatment period of approximately two years. A treatment inspection is an inspection that was assigned as part of the treatment, or a follow-up inspection that follows such an inspection in the same chain. Columns (1) and (2) show means with standard deviations in brackets. Column (3) shows the coefficient from regressions of each variable on treatment, where each regression includes region fixed effects and a control for the audit sample. The table uses administrative data from during the experiment. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B5: Endline compliance on treatment (All Obs)

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

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

Table B6: Endline Pollution and Compliance on Letter Treatment

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

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

Table B7: Model of Plant Revenue

	Log Revenue			
	(1)	(2)	(3)	(4)
Medium size industry (=1)	1.410*** (0.236)	1.582*** (0.266)	1.366*** (0.254)	1.336*** (0.271)
Textiles (=1)		0.732*** (0.189)		-0.109 (0.316)
Coal or lignite as fuel (=1)			0.122 (0.234)	0.163 (0.271)
Waste water generated > 100 kl/day			0.846*** (0.216)	0.905*** (0.249)
Constant	11.11*** (0.0975)	10.70*** (0.174)	10.64*** (0.175)	10.65*** (0.182)
R^2	0.0962	0.160	0.201	0.202
Residual SD	1.358	1.309	1.277	1.277
Observations	694	694	694	694

The table presents regression estimates of plant revenue on plant characteristics. Plant revenue is measured from the endline surveyed and specified in logged thousands of Indian rupees. Medium size is a regulatory designation for capital stock above INR 5m (the sample includes only small- and medium-sized plants), Textiles is a dummy for belonging to the textile industry, Coal or lignite is a dummy for using one of those fuels, and Waste water generated is a dummy for discharging more than 100 kl/day of effluent. Standard errors clustered by plant in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$