

Can Pollution Markets Work in Developing Countries? Experimental Evidence from India*

Michael Greenstone, Rohini Pande, Anant Sudarshan and Nicholas Ryan[†]

September 19, 2022

Abstract

Market-based environmental regulations are seldom used in developing countries, where pollution levels are the highest but state capacity is often weak. We experimentally evaluate a new emissions market for particulate matter covering industrial plants in a large Indian city. We have three main findings. First, the market functioned well: permit trade was active, leaving few permits unused, and plants obtained permits to meet their compliance obligations almost perfectly. Second, treatment plants randomly assigned to the emissions market reduce pollution emissions by 20% to 30%, relative to control plants that remain in the status quo, command-and-control regime. Third, at a constant level of emissions, set to the market cap, variable abatement costs would be 12% higher in the command-and-control regime than under the market. This finding is based on a model of abatement costs fitted in the treatment group on the assumption that plants bid for permits at their expected marginal abatement cost. The findings suggest that markets offer a way to reduce air pollution emissions in India at a relatively low cost.

*We thank the Gujarat Pollution Control Board for collaborating in this research, particularly Rajiv Kumar Gupta, Sanjeev Kumar, Tejas Patel, Manali Bhatt and Aparna Chaubey. We thank the MacArthur Foundation, USAID and the Tata Center for Development at Chicago for financial support. We thank Gargee Goswami, Dipika Gawande, Prajval Jhunjhunwala, Gargi Pal, Shruti Bhimsaria, Vineet Gupta, Winston Hovekamp, Jared Stolove, Fikremariam Gedefaw and Kaixin Wang for excellent research assistance and Sanjana Gorti and Amber Luong for project management.

[†]Greenstone: Energy Policy Institute and Department of Economics, University of Chicago. Pande: Department of Economics, Yale University. Sudarshan: University of Warwick. Ryan (corresponding author): Department of Economics, Yale University (nicholas.ryan@yale.edu).

1 Introduction

Air pollution harms human well-being by shortening people's lives and undermining their formation of human capital (Ebenstein et al., 2017; Isen, Rossin-Slater and Walker, 2017; Ebenstein and Greenstone, 2022). With these harms in mind, India may be considered the center of a global air pollution crisis. Air pollution exceeds the World Health Organization (WHO) standard for very fine particles (PM2.5) in all but one of India's 687 administrative districts. Estimates suggest that reducing air pollution throughout India to the WHO standard would increase average life expectancy by five years for a population of 1.4 billion people (Greenstone et al., 2015; Greenstone, Hasenkopf and Lee, 2022). High pollution levels persist despite that abatement technology is affordable and environmental regulations are stringent, on the books (Piette, 2018). One reason is that regulatory enforcement is imperfect and gives plants only weak incentives to abate pollution (Duflo et al., 2013, 2018). Such gaps in enforcement have been found to be ubiquitous in developing countries and may be a main cause of the pollution crisis (Greenstone and Jack, 2015).¹

Can pollution markets work in developing countries? In the United States, the introduction of markets for pollution has been hailed as a landmark example of turning economic principles into policy (Ellerman et al., 2000). Markets, by allowing plants with low abatement costs to reduce emissions more, can abate pollution at the lowest possible cost (Dales, 1968). This efficiency may be especially valuable in lower-income countries where regulators and citizens are especially sensitive to cost. However, in the same places, the basic assumptions on which emissions markets rely may not apply. In economic models of pollution markets emissions are known, plants always comply and permit trade is unburdened by transactions costs (Coase, 1960). In practice, regulators in developing countries struggle with reliable monitoring and may lack the tools or credibility to force each polluter to hold a permit for every unit of emissions. Emissions markets trade in a commodity created by the state, and these markets cannot function if the state cannot credibly ensure the value of that commodity. Perhaps for these reasons, market-based instruments have

¹A complementary explanation for persistent high pollution in developing countries is that citizens have low willingness to pay for environmental quality (Ito and Zhang, 2020; Baylis et al., 2022; Jack et al., 2022).

rarely been adopted to regulate pollution in developing countries (Stavins, 2003; Blackman, Li and Liu, 2018).

This paper provides the first experimental evaluation of how a pollution market impacts emissions and compliance costs. The experiment we study introduced the world's first market for particulate matter emissions. The evaluation builds on a collaboration over more than a decade with the environmental regulator in Gujarat, India. We authored a white paper on the use of emissions trading for regulation in India in 2010 (Duflo et al., 2010). Thereafter we collaborated with the Gujarat Pollution Control Board (GPCB), the environmental regulator in Gujarat, as it assembled the building blocks for an emissions market. GPCB and its partners set standards for continuous emissions monitoring systems (CEMS) (jointly with the Central Pollution Control Board), mandated that plants install CEMS, laid down the rules for the market and developed a trading platform.

The market for particulate matter emissions from industrial plants was launched in 2019. GPCB mandated CEMS for all eligible large sources of particulates, totaling 317 industrial plants, in the airshed of Surat, a city of 4 million people. The status quo regulation for all of these sample plants is a command-and-control regime in which plants are mandated to install abatement equipment and subject to a standard for the maximum concentration of particulate emissions during infrequent spot checks. After sample plants installed CEMS, a randomly assigned treatment group of 162 plants were shifted into the new emissions market. Treatment plants were allocated and sold permits giving the right to emit one kilogram of particulate matter apiece and could trade permits with other treatment group plants. The control group remained in the status quo regime. Both the treatment and the control group are therefore subject to a new monitoring regime. Conceptually, the treatment group differs from the control group both because treatment plants are subject to standards that limit load, rather than concentration, and because treatment plants were allowed to trade. The evaluation ran from September 2019 to April 2021, with the market suspended from late March of 2020 (when India entered a lockdown against Covid-19 pandemic) until the end of November 2020. The treatment duration is thus roughly a year spread over 18 calendar months.

There are three main findings. The first finding is that the market functioned well, especially

after the first few compliance periods. For example, almost all plants held enough permits to cover emissions in all compliance periods. The regulator assured compliance by establishing a reputation early on. Two firms failed to turn in enough permits at the end of the first compliance period. The regulator responded by levying monetary damages at twice the level of the market price for each unit of excess emissions, as specified in the market rules. Compliance thereafter was essentially perfect. Qualitatively, permit trade appears to have been efficient. Plants traded frequently and trading volume was high, up to 20% of the market cap on some single days. Plant permit holdings at the end of each compliance period differed greatly from initial allocations and plants left little money on the table, in terms of unused and unsold permits at the end of compliance periods. These findings on the large and yet precise changes in permit holdings through trade suggest that the transactions costs of trade in the market were low.

Second, the treatment caused a 20% to 30% decline in particulates emissions, relative to control plants that remained under the status quo regime. The range of point estimates reflects differences in how we impute emissions during periods when plants fail to report data through CEMS. The finding of greater pollution reductions among treatment plants is obviously a function of the emissions cap. The cap had initially been set, based on incomplete data, so that the load standard in the treatment was expected to be as stringent as the concentration standard in the control. On seeing low permit prices, the regulator adjusted the cap downwards over the first several compliance periods. The lower level of emissions in the treatment therefore reflects both improved compliance and an endogenously greater stringency of regulation.

Third, we find that variable abatement costs, holding emissions per plant constant at the level set by the market cap, would be 12% higher in the command-and-control regime than under the emissions market. The exploration of costs begins with a treatment effects analysis of the effect of the market on abatement capital. We find that treatment plants did not increase expenditures on air pollution abatement equipment. This result may not be surprising, because the status quo regime already features a well-enforced set of mandates to *install* equipment. Many of the costs of abatement may instead be variable costs. We found these costs difficult to measure directly,

because they are made up on inputs, like labor and electricity, that are widely used throughout the plant and for which abatement makes up a small share of total input expenditure. We therefore use a model to infer variable abatement costs from plant bids in the permit market. In the model, plants bid for permits at their expected value of marginal abatement costs in a compliance period. We specify a model with rich heterogeneity in abatement costs across plants and time and estimate the elasticity of cost with respect to abatement from the treatment plants' bids and offers for permits. We then apply the estimated abatement cost functions to counterfactually simulate control costs to achieve the treatment level of emissions. The model therefore leverages the permit bidding data, available only in the treatment group, to study the differences in costs between alternate regulatory regimes. The broadest finding of this exercise is that the marginal costs of abatement as revealed in the market appear to be relatively low at emissions levels near that in the treatment.

Our paper contributes to the literature at the intersection of environmental economics and development (Greenstone and Jack, 2015; Jayachandran, 2022). A main theme of this literature has been the Herculean difficulty of effective environmental regulations in low state capacity settings (Greenstone and Hanna, 2014; Blackman, Li and Liu, 2018). Common findings are that poor or corrupted monitoring impedes regulation (Duflo et al., 2013; Oliva, 2015; Duflo et al., 2018; Zou, 2021), and coarse regulations adopted in response to poor monitoring, in turn, are partly undercut through behavioral responses (Davis, 2008; He, Wang and Zhang, 2020). The one prior example of a market targeting particulates with which we are familiar illustrates the conjoint nature of the monitoring and regulation problem (Montero, Sanchez and Katz, 2002).² Our findings suggest that, if the monitoring problem can be addressed, the *private* costs of pollution abatement may not be high, even in settings with very high pollution emissions and ambient pollution levels (see also Greenstone et al., 2022).

²The market in Santiago, Chile targeted particulate matter emissions indirectly and achieved little abatement, because it traded in boiler capacity, a proxy for emissions, rather than directly in emissions load. The problem this created was that particulate abatement is most cheaply achieved through post-combustion actions that leave boiler capacity unchanged. The potential abatement of boiler capacity in this market was therefore extremely inelastic to prices since it is costly to change boiler capacity (a long-term investment). Ultimately, the market disbanded after many covered sources switched fuels in response to a fall in natural gas prices. Pollution monitoring technology has improved in recent years and the particulate market we study relied on modern CEMS.

The focus of the literature on the efficacy of market-based instruments has been on the landmark US environmental markets including the RECLAIM program (introduced 1994 to target SO_2 and NO_X), the Acid Rain program (1995 to target SO_2) and the NO_X Budget Trading Program (introduced in 2003). Although it is widely believed that emissions markets achieve abatement at lower cost than available command-and-control regulations (Ellerman et al., 2000; Burtraw et al., 2005; Fowlie, Holland and Mansur, 2012), most evaluations of emissions markets use engineering estimates of costs because plausible empirical counterfactuals do not exist (Burtraw et al., 2005).³ The best econometric studies of emissions markets use counterfactuals based on the behavior of plants in other areas, smaller plants that are not subject to the same regulation, or time series forecasts based on pre-program emissions (Fowlie, Holland and Mansur, 2012; Martin, De Preux and Wagner, 2014; Borenstein et al., 2019). This study estimates the effect of emissions trading on both pollution and plant costs against a well-defined experimental counterfactual.⁴ The emissions market we study is a policy innovation in a setting where, like in many parts of the world, regulators with low capacity and a constrained set of tools face crisis levels of air pollution.

The remainder of the paper goes as follows. Section 2 introduces environmental regulation in India and the experimental design. We describe the rules of the emissions market started in the treatment. Section 3 discusses the data and the balance of baseline covariates by treatment arm. Section 4 describes trading activity and compliance in the emissions market. Section 5 presents estimates of experimental treatment effects on emissions and abatement capital costs. Section 6 describes the model for pollution abatement and Section 7 its results on the counterfactual costs of abatement under alternate regimes. Section 8 concludes.

³For example, Fowlie, Holland and Mansur (2012) write that “Unresolved disagreements about what constitutes an appropriate measure of counterfactual emissions have resulted in a plurality of opinions regarding RECLAIM’s overall performance. After 15 years of program evaluations, the emissions impacts of RECLAIM vis-à-vis the subsumed CAC rules remain controversial.” (pg. 971)

⁴Martin, Muûls and Wagner (2020) write “An ideal evaluation of the EU ETS would combine a representative firm- or plant-level data set of sufficient detail with a study design that attributes to the EU ETS only those observed behavioral changes it has actually caused. It is difficult to solve this identification problem because there are so many factors that might simultaneously affect firm behavior, thus confounding the impact estimate. The state-of-the-art solution would be to conduct a randomized control trial or field experiment (e.g., Greenstone and Gayer (2009)). As in other real-world settings, however, randomizing participation in the EU ETS is neither desirable nor politically feasible.”

2 Context

2.1 The Command and Control status quo

India's Air Act (1981) defines a command and control framework for regulating industrial pollution. Aspects of this framework, such as emission limits and measurement standards, are set by India's Central Pollution Control Board (CPCB) with enforcement delegated to State Pollution Control Boards (SPCB).⁵ The "command" is a mandate that plants install equipment to reduce pollution, a requirement that is generally checked when a new factory seeks clearance to operate or expand. The "control" is based on manual measurements of pollution, where plants can be sanctioned if their pollution emissions have a concentration above specified limits.⁶

The "command" element of the status-quo is reasonably effective. In our sample, 100% of operating plants had at least one piece of abatement equipment for air pollution installed. Panel B of Table 2 describes the composition of this equipment. In contrast, the "control" element works less well. Figure 2 shows how many plants pollute well above the legal maximum limit. Other research in the same setting has shown that weaknesses in regulatory monitoring result in poor quality official data that typically underreport pollution levels (Duflo et al., 2013). In addition, polluting plants are infrequently penalized (Duflo et al., 2018). As a result, despite the presence of abatement equipment, plants have weak incentives to maintain and operate it.

2.2 The creation of Surat's emissions market

The market in Surat was introduced by the Gujarat Pollution Control Board in the form of a randomized control trial encompassing 317 coal-burning plants located within 30km of the city center. The market represents the culmination of a collaboration with the state regulator over more

⁵The Water Act (1974) established State Pollution Control Boards as state environmental regulators, providing them the power to enforce standards for water pollution. The Air Act (1981) extended these powers to cover air pollution. SPCBs are also empowered to introduce additional regulation in 'highly polluted' regions.

⁶In principle regulators are empowered to impose harsh punishments, including temporary or permanent plant closures and even criminal cases and imprisonment. However, because of the high burden of proof and the difficulty in bringing such cases, prosecutions are conducted very rarely. Meaningful civil fines are *not* part of the portfolio of permitted sanctions in the Air Act.

than a decade.

Prior work involving some of the authors on this paper focused on improving the monitoring and enforcement of status-quo regulation. In Duflo et al. (2013) we describe an experimental evaluation of a reform that increased the independence of third-party auditors and, as a consequence, led audited firms to pollute less. While this reform improved reporting and reduced pollution, it also highlighted the weaknesses in status quo regulatory instruments, particularly monitoring.

In 2010, we authored a white paper addressed to India's Minister for Environment, making the case for using new monitoring technology and testing market-based regulation as an alternative to command-and-control (Duflo et al., 2010). CEMS are widely used in the United States and Europe but had never been mandated in India. A member of the research team (Sudarshan) worked with the Central Pollution Control Board to develop the first technical standards for the use of CEMS in India, focused on using them to measure the mass of pollutants emitted, as would be required to run an emissions market (Central Pollution Control Board, 2013). These standards led to a decision by GPCB to mandate CEMS for a sample of plants in Surat in 2013. A year later, the Central Pollution Control Board (CPCB) followed with a national CEMS mandate for a range of pollutants for large plants.

The rollout of CEMS in Surat was challenging from an implementation point of view and only completed in 2019.⁷ The GPCB convened the research team and an implementing partner to consult on the design of the emissions market as CEMS installations progressed. Our research team collaborated with the regulator and the partner, a private commodity exchange (NCDEX e-Markets Limited (NeML), a branch of the National Commodity and Derivatives Exchange), on the market design. NeML hosted the market on its trading platform after it launched.

3 Sample, Experimental Design and Data

This section describes our sample, data and experimental design.

⁷We discuss the technical, institutional, political, and economic considerations involved in the introduction of CEMS into a low-capacity developing country setting in a forthcoming open-access textbook on Development Engineering (Sudarshan, 2023).

3.1 Sample and experimental design

The sample of industrial plants is selected to include plants with the highest air pollution potential in and around the the industrial city of Surat, Gujarat. Surat is a city of over 4 million people and is known as a hub for the textile industry. Plant air pollution potential was determined on the basis of solid fuel consumption. All 342 plants listed in the regulator’s records that met the following criteria were eligible for the sample: (i) the plant consumed solid fuel (coal or lignite, mainly), (ii) plant boiler capacity of at least one ton per hour (iii) stack diameter of at least 24 centimeters, to allow for CEMS installation and measurement (see Annex Table A5). After the installation of Continuous Emissions Monitoring Systems (CEMS) in all sample plants, half of plants were randomly assigned to the treatment arm. After assignment, but prior to the start of the market, some additional plants closed or were found to be ineligible because they operated only seasonally. The final sample of plants is 317 plants across the two arms with 162 in the treatment.

Figure 1 shows the contribution of particulate matter emissions from sample plants to ambient particulate matter concentrations in Surat, Gujarat in June 2019.⁸ Treatment firms are represented by blue “x” markers and control firms by black “o” markers. The map shows 304 plants with the remaining 13 sample plants lying outside the bounds. The PM concentrations shown in the map (in $\mu g/m^3$) are measured by passing CEMS emissions rates for the sample plants into a simplified Gaussian dispersion model, comparable to the SCREEN3 model used by the US Environmental Protection Agency.⁹

There are two main points to be drawn from the map. First, the contribution of sample plants to urban air pollution is large. The map shows the level of ambient pollution that would be observed in Surat if sample plants were the *only* source of particulates. The implied concentration of total

⁸The river Tapi runs across the northwest corner of the map and the main railway line through the city runs north to south on the western side of the map. The clusters of plants to the east and southeast of the city center are the industrial areas in Kadodara and Palsana. The dense cluster just south of the city center is the Pandesara GIDC (Gujarat Industrial Development Corporation, shorthand for a state-sponsored industrial cluster).

⁹The model used is based on the eddy diffusion theory, considering each plant to be a stationary point emitting source, with CEMS data informing each plant’s emission rate (mass/time) and the location and stack height of each plant determining the point source. The model uses simple assumptions on meteorological conditions such as constant emission rates and wind speeds.

suspended particulates (of all sizes) in and near the city center, ranges from $60 \mu\text{g}/\text{m}^3$ to $180 \mu\text{g}/\text{m}^3$ across space. A typical $PM_{2.5}/TSP$ mass ratio for urban air pollution is around 0.3 (Lall et al., 2004). This ratio implies that the ambient air pollution in parts of Surat, due to industrial point sources alone, would exceed the WHO standard for fine particles by a factor of roughly 2 to 5. Second, the dense clustering of plants and the extent of particulate dispersion in the model imply that pollution from most plants in the market will affect the same areas. The city-level market therefore reduces concerns that emissions trading, even if it reduces pollution overall, may create areas of higher pollution concentration.

Table 1 shows timeline of the experiment and data collection. beginning July 16th, 2019, the market launched with two “mock” trading compliance periods in which no money was at stake, but plants were allocated permits and could buy permits with endowments of fake money. The purpose was to inculcate plants in the market rules and improve the coverage of CEMS monitoring of pollution. After this, there were six real compliance periods, beginning on September 16th, 2019 and running collectively until March 22nd, 2020. At this point, market operations were suspended due to India’s nationwide lockdown, in response to the Covid-19 pandemic. The market was restarted in December 2020 before being affected by a fresh set of restrictions in April 2021 following the second (Delta) wave of Covid. This paper analyzes a data set that includes four additional compliance periods after the first interruption. The duration of treatment is therefore roughly one year of market operation spread over one and a half years of calendar time ending in April 2021.

3.2 Data sources

We use three main sources of data. First, two waves of a survey of plant characteristics, abatement equipment, and economic variables which were conducted prior to market launch and roughly one year after the market started (Table 1). Second, high-frequency pollution data from Continuous Emissions Monitoring Systems (CEMS) and, third, trading data from the market operator.

The in-person baseline survey was conducted from December 21st, 2018 to January 29th, 2019 and had two parts, general and technical. The general part was administered to the plant owner

or manager, and covered plant economic characteristics such as inputs, outputs, sales and energy consumption. The technical part directly observed the abatement equipment installed on every point source of emissions in the plant. Most plants have a single stack, or chimney, though some large plants have more than one. Our team recorded the characteristics of all emissions sources and all abatement equipment attached to those sources and interviewed plant staff about costs of equipment operations. In addition, we simultaneously hired independent environmental labs to take manual stack samples of air pollution emissions at each plant. These samples measure the concentration of particulate matter in stack gas at the time of the survey. We have two waves of manual samples of plant emissions concentration from prior to the start of the emissions trading experiment.

The second source of data is high-frequency data on air pollution from CEMS devices. CEMS are calibrated by comparing CEMS readings to physical pollution samples taken simultaneously in the same stack. CEMS readings measure Suspended Particulate Matter (SPM), which includes particles of all sizes (unlike PM_{10} or $PM_{2.5}$, commonly reported for ambient pollution, which measure the mass only of fine or very fine particles). Pollution readings are reported continuously over the internet to a central server. The main data limitation for CEMS is incomplete reporting during the experiment. The data handling system was designed to store data locally during transient internet outages. However, longer outages and device malfunctions leave gaps in the high frequency data. The mean rate of weekly data reporting rose from initially low levels to over 85% of plants by the end of the sample. Once the treatment assignments were announced, treatment plants had stronger incentives to report than control plants, because their validated emissions in the market would increase if they did not report data continually (see Section 2). Reflecting this, we observe higher rates of data availability for treatment plants than control plants, especially at the start of the market (Annex Figure A1). Control plants caught up in later periods. To account for differential reporting across treatment arms, we will analyze the CEMS data using several different imputation rules for emissions in plants that did not report pollution readings in a given week. The imputation rules are described in Annex Table A1.

The third data source is administrative data on all permit trades in the market. The market operator recorded a ledger of permit purchases and sales for all compliance periods and a complete order book of bids and offers for permits regardless of whether a transaction occurred. These data sets cover all trades and bids, since all permit trades had to occur on the market operator’s platform to be registered. We use this data to describe the market and to estimate marginal costs.

3.3 How did plants fare on pollution and abatement at baseline?

Table 2 summarizes plant covariates by treatment arm. The sample is balanced at baseline across a wide range of measures of inputs, outputs, equipment and pollution.

Panel B shows that sample plants are large factories, and that energy and related inputs comprise a large share of plant expenditures.¹⁰ The average control plant spends USD 350,000 on electricity (panel A). The annual cost of running the boiler, the main source of pollution in the plant, is USD 108,000, excluding direct expenditures on fuel.

Panel C shows that nearly all plants in both treatment and control groups have abatement equipment for air pollution installed at the baseline. For example, 97% of control (98% of treatment) plants have a cyclone installed, 86% of control (81% of treatment) plants have a bag filter installed, 60% of control (64% of treatment) plants have a scrubber installed and 8% of control (11% of treatment) plants have an electrostatic precipitator. The rates of installation move inversely with the expense and efficacy of abatement equipment. All plants must install cyclones, which are inexpensive but relatively low efficacy (reducing SPM emissions by 60-90% but $PM_{2.5}$ by only 0-40%). Larger plants with multiple emissions sources are more likely to be required to install more expensive capital equipment like scrubbers or bag filters (which are rated to remove greater than 90% of SPM load). The “command” portion of regulation, that plants must install equipment, works well, in the sense that these mandates are followed.

Panel C also shows several measures of baseline pollution emissions. PM concentration is the mean particulate matter concentration from manual pollution samples taken during our baseline

¹⁰While many plants are formally classified as “small scale” (71%) this is a government classification, based on the reported capital stock at the time of the plant’s establishment.

survey. The average concentration of SPM in stack gas is 169 mg/Nm^3 in the control group and 179 mg/Nm^3 in the treatment group. Both *average* levels of emissions exceed the SPM *maximum* standard of 150 mg/Nm^3 , and the flow rate of stack gas is balanced across treatment arms. In addition, we also observed pollution by having our enumerators grade the color or opacity of stack gas, from a vantage point outside the factory gate before the survey. This data is gathered without the prior knowledge of plant operators. These grades follow standard “Ringelmann” scores, which range from 0 (no visible air pollution emissions) to 5 (heavy, dark smoke). The Ringelmann readings are also balanced across arms at baseline.

Figure 2 shows baseline pollution distributions in the control and treatment plants using stack samples of pollution concentration (truncated at the 95th percentile, Panel A), and Ringelmann scores (Panel B). The distributions of pollution by either measure are very similar across treatment arms at baseline. About 30% of plants in both arms do not comply with the particulate matter emissions standard (Table 2, panel C).

4 The functioning of the emissions market in the treatment

In this section we describe the emissions market and then establish that the market functioned well in that it hosted a large volume of trade at low prices that yielded a high rate of compliance.

4.1 Market design

Our research team collaborated with the regulator and the market operator to develop the market design for the Gujarat particulate matter market.

Plants.—The market regulates industrial plants, primarily in the textile industry. The characteristics of sample plants are described in Section 3.

Permits and emissions reporting.—Emissions are the total mass of particulate matter during a compliance period. A permit entitles plants to one kilogram of total suspended particulate

(TSP) emissions. As is typical for regulations of particulate emissions, as opposed to ambient pollution levels, the stack-level regulation here does not differentiate by particle size. Emissions are monitored by Continuous Emissions Monitoring Systems (CEMS) on plant stacks.

Compliance.—The compliance period is the period over which permits for emissions are valid. There is no banking or borrowing allowed. At the end of a compliance period, firms must true-up their permit holdings against their cumulative emissions during the compliance period. The compliance period in the market ranged from one month to six weeks, depending on the period considered. Table 1 gives the timeline of ten compliance periods. The India-wide lockdown beginning in May 2020 interrupted the market, which resumed again in December of that year.

Cap.—The regulatory standard for air pollution emissions in sample plants is $150 \text{ mg}/\text{Nm}^3$, which represents the *concentration* of pollution in the gas being emitted from a plant’s stack. The cap in this market instead limits the *load*, or total mass, of pollution emitted. Load is the concentration of pollution multiplied by the volume of gas emitted, which depends on the rate of emissions and plant capacity utilization. Based on ex ante estimates of likely emissions load, GPCB set an initial cap of 280 tons of particulate emissions per month. This cap was set to match the stringency of the existing concentration standard under assumed levels of flow and capacity utilization, which were not accurately measured at the time the cap was first set. The regulator’s goal was to initially set the cap to approximate the stringency of existing regulation and then possibly to tighten over time. As more data arrived, the cap was revised downwards in later compliance periods to 170 tons per month. Appendix Table B1 gives the cap for each compliance period. Figure 6 shows the level of the cap per plant using red horizontal lines. The cap in later periods works out to roughly 1,000 kg of particulate matter emissions per plant-month.

Permit allocation.—Most permits, 80%, were allocated to plants for free (“grandfathered”). The allocation of free permits to plants was done pro rata in proportion to each plant’s share of total emissions capacity. Emissions capacity, in tons per hour at the plant level, is the sum of the capacity of the boiler and thermic fluid heater, the two main fuel-burning pieces of equipment in a

plant. These capacity measures were drawn from administrative data that pre-dated the design or introduction of the market. Plants therefore did not have any opportunity to game or adjust their capacity measures in response to the pro rata allocation rule.

The balance of permits, 20%, were allocated to plants via a uniform price, multi-unit auction run by the Gujarat Pollution Control Board. Each compliance period opened with a uniform price auction in which GPCB offers its entire supply of permits at the market floor price. If the permits offered by the GPCB did not sell out in the first auction, GPCB would offer them again at subsequent weekly auctions, until they were exhausted.

Permit trade.—Plants can trade permits in two ways: via weekly auctions or over-the-counter trades between the auctions. Both of these markets were run by a single operator. At auction, GPCB would offer its permit share, as described above. All participating plants can additionally offer step functions from price to the quantity of permits they wish to buy or sell. The clearing price at auction is the lowest price at which net quantity demanded is weakly negative.

Over-the-counter (OTC) trades could occur between weekly auctions. While the quantity of over-the-counter trades was not restricted, firms could only trade permits at the price revealed by the most recent weekly auction. This restriction on OTC prices was adopted in order to encourage parties to participate in auctions and to limit volatility in permit prices.

Price collar.—Permit prices were restricted to be no less than INR 5 per kilogram and no more than INR 100 per kilogram. The range of the price collar was informed by *ex ante* engineering estimates that abatement of particulate matter, by the equipment commonly in use in the sample, could occur at an *average* (not marginal) cost of between INR 10 and INR 40, depending on the type of equipment installed and the scale of the plant. The ceiling price was therefore seen as sufficiently high that all plants would prefer to abate than to pay the ceiling price per unit of emissions.

The price collar was mechanically enforced in the auction and trading system. Operationally, the floor price was supported by a GPCB commitment to buy back permits at the floor price, in a

quantity up to the value of 20% it initially offered. The ceiling price was supported by a GPCB commitment to sell permits at the ceiling price at the end of each compliance period in unlimited quantity.

Missing data rule.—Plants that do not report emissions for any period of time during the compliance period have their emissions for that period imputed. Non-reporting could occur because of an internet outage, a CEMS device malfunction or other disruptions. The goal of the imputation rule adopted was to incentivize complete and accurate reporting. Missing emissions data was therefore imputed at a high rate that increased in the share of time that the plant did not report over a compliance period. Emissions with imputations are called validated emissions and are used for the determination of compliance.

Non-compliance and penalties.—Plants participating in the market were required to post an environmental bond, called an Environmental Damage Compensation Deposit (EDCD), at the start of the market. At the end of each compliance period there was a one-week true-up period in which an additional auction was held and OTC trade could occur for plants to buy permits if in deficit or sell permits if in surplus. At the end of the true-up period, any plants that had not bought enough permits to cover their emissions during the compliance period were subject to a fine, called Environmental Damage Compensation (EDC), at the rate of twice the ceiling price for every unit of emissions in excess of their permit holdings. The fine was deducted from the EDCD initially posted. This mechanism was established to show the regulator's commitment to penalizing plants that did not comply.

The market, therefore, is of a relatively standard design for emissions trading markets that have been used in other applications. Permits were allocated largely for free, to reduce plant costs, but with a portion auctioned to promote price discovery. Demand for permits is sustained in the market by the threat of fines for emissions in excess of permit holdings.

4.2 Permit prices and quantities

Figure 3 describes the time series of permit prices (panel A) and permit quantities (panel B). Panel A shows the time series of trading prices over time (solid line) and of bid prices (dotted line). Market-clearing prices are generally low, and depending on the compliance period and week range from INR 5 per kg (the price floor) to INR 16 per kg. The market-clearing mechanism (Section 4.1) deliberately reduces price volatility, by constraining over-the-counter trades to occur at prices revealed by auctions held once a week during the compliance period. Therefore, the price time series steps up or down over time. In general, prices were lower in the pre-Covid-interruption compliance periods (1 through 6), when the cap was looser, and higher after the market resumed. In several compliance periods, for example periods 9 and 10, prices are moderately high during the compliance period but then dive at the end, during the true-up period, when emissions are known with certainty. This kind of price behavior is consistent with uncertainty, prior to the end of the compliance period, as to whether the market would be short or long permits in aggregate. When the market closes and this uncertainty is resolved, we expect prices to converge to the ceiling or floor, respectively. Again, the market-clearing mechanism of having a single auction after the close of the compliance period may mute price volatility at the end of each period.

While clearing prices were generally low many individual firms offered bids at higher prices. The dots overlaid on Figure 3, panel A give the average bid price in each weekly permit auction. In the early set of compliance periods average bid prices were commonly in the range from INR 10 to INR 25 per kg, though the market cleared at lower prices. In panel B we show daily permit quantities traded as a fraction of the cap for each compliance period. In each period, there is typically a spike in quantity in the early-going. This represents the double-sided auction held on the first Tuesday of the compliance period. In this auction, the regulator (GPCB) sells a large quantity of permits at the floor price. This sale incentivizes plant participation and establishes a reference price for bilateral trade during the following week, until the next auction is held. Overall the volume of trade is significant, with volumes as high as 20% of the monthly cap, or more, on many single days. Trade volumes are typically larger during the first part of a compliance period

as plants purchase or sell permits to align permit holdings with expected emissions. As plants' uncertainty about their total emissions for the period diminish, toward the end of the period, so do trade volumes.

4.3 Permit allocations and plant emissions

A major question at the start of any market created only for compliance purposes is whether participants will respect the regulatory mandate underlying the market.

Figure 4 plots the distribution across plants of emissions as a fraction of permit holdings, at the end of the true-up period, in 10 separate panels, one for each compliance period. There are two main findings. First, compliance is nearly perfect, where we define compliance as permit holdings at the end of the true-up period exceeding emissions during the compliance period. Plants in bins to the right of 100% are non-compliant because their emissions exceed their permit holdings as of the end of the compliance period. There are only a handful of plants that have emissions exceeding their permit holdings (see periods 1, 3 and 8). In most compliance periods, no plants exceed their permit holdings. Second, the market appears to be efficient in that the overwhelming majority of plants hold permits, at the end of the period, almost exactly equal to their total emissions. A plant that has emissions below permit holdings at the end of the compliance period is “leaving money on the table” in that they bought more permits than needed. These plants have the option of selling excess permits to other plants or to the regulator at the end of the compliance period, if the market clears at the floor price. Looking down the first column of distributions, and then down the second, we see that (i) relatively few plants left money on the table (ii) more plants left money on the table in early compliance periods, when the market was more likely to clear at low prices (iii) in the later compliance periods almost all plants hold only the permits they need to cover their emissions.

We believe compliance was high because the regulator credibly established that it would penalize non-compliance in the early going. In the first compliance period two plants had emissions exceeding their permit holdings. Plant A had emissions of 3928 kg against permit holdings of 3456 kg and Plant B emissions of 4716 kg against permit holdings of a mere 1456 kg. These plants were levied Environmental Damage Compensation (EDC) in accordance with the market rules. Plant A

paid the EDC and then topped up their environmental bond. Plant B had not posted the bond. The regulator ordered plant B to be closed down. Plant B then posted their bond and paid a penalty of INR 652,000. The regulator revoked their closure and allowed the plant to reopen after two weeks.

To what extent is the exact degree of compliance determined by the allocation of permits? Permits totaling 80% of the cap were allocated to plants *pro rata* on the basis of a plant's total heat output, a measure of emissions capacity. If plants had the same capacity utilization, and the same rate of emissions per unit of heat output, then this allocation could in principle equate emissions to permit holdings even absent trade. To test this idea, we plot in Figure 5 the distribution across plants of plant emissions as a percentage of their initial permit allocation in each compliance period. Plants that emit exactly as much as they were allocated would appear as 100% and plants that emit twice what they were allocated would appear as 200%. Because only 80% of the total cap in each period is grandfathered, with the rest being auctioned, we expect the mean emissions as a percentage of the permit allocation to exceed 100%. We find that plant emissions are often much higher or lower than initial permit allocations. In all compliance periods, emissions are fairly widely dispersed, with most of the mass of the distribution between 50% and 200% of the initial allocation. This dispersion implies that plants are unconstrained by their initial permit allocations and trade freely to set permit holdings equal to their emissions (as shown in Figure 4). This dispersion also implies that initial holdings are at best a weak proxy for ultimate emissions. Differences in plant capacity utilization and emissions rates in operation create dispersion relative to the pure capacity-based measure of heat output that is used for permit allocation.

We conclude that the market functioned well. A large volume of trade, at fairly low prices, enabled plants to move from their initial permit allocations to permit holdings that met their emissions. Emissions rarely exceed holdings, but nor do plants leave money on the table by holding extra permits. Finally, the wide dispersion in emissions relative to holdings suggest a broad scope for gains from trade in the emissions market.

5 Empirical results

This section presents empirical results on the average effect of the emissions trading treatment on plant pollution emissions and abatement capital costs.

5.1 Pollution emissions

Graphical analysis.—Figure 6 shows the mean per-plant emissions in kilograms per month, at weekly frequency, over 11 months from April 2019 to April 2021, by treatment arm. Treatment firms are represented by the solid (blue) line, control firms by the dashed (grey line). The vertical grey shaded regions mark the compliance periods during the market, which are separated by vertical lines. The compliance periods are interrupted by the Covid-19 lockdown, denoted interregnum on the chart and shaded in light blue. This interruption divides the sample into *early* (periods 1 to 6) and *late* (periods 7 to 10) compliance periods. The horizontal (red) lines denote the per-plant month market cap for each compliance period, calculated as the aggregate market cap divided by 162 Treatment plants.

In the three months prior to the start of the market, mean emissions loads in the treatment and control plants hover around 2250 to 2500 kilograms per month. At the beginning of the mock trading period, July 16th, emissions begin to decline in *both* groups, but more steeply in the treatment group. By the beginning of compliance period 1, in September (shaded grey), treatment plants are emitting roughly 300 kilograms per month less particulate matter than control plants. The difference between treatment and control average emissions load is then sustained at a similar level during both the early and late sets of compliance periods.

As we discussed in the data section, pollution reporting in the experiment was incomplete and differential across treatment arms: treatment plants reported more reliably because of penalties for non-reporting built in to the market rules. Appendix Figure A1 shows the time series of reporting over time, which was large at the start of the market, but narrowed to only a few percentage points by the end of the experimental period. The main pollution series presented in Figure 6 imputes

emissions for a given plant, when missing, using observations for the same plant at other times within the same week or month. In Appendix A, Figure A2 we show the same time series with alternate imputation rules that admit imputation either within plants across months (panel A) or within treatment arms and months (panel B). As control plants with high imputed emissions are especially likely not to report, we find that these imputations generally increase the estimated difference between treatment and control emissions.

Returning to Figure 6, emissions met the cap in all compliance periods. For each compliance period, we plot a horizontal line showing the mean per plant emissions required to meet the cap exactly. Mean emissions are below this level for all compliance periods, sometimes sharply below (around the Diwali holiday, in November, most plants cease operations for a week and emissions plummet). The apparent over-compliance in early compliance periods is because missing emissions, for the purpose of market compliance, were imputed at a rate much higher than the mean, reflecting a punitive assumption intended to encourage reporting. With those more punitive imputations the cap more closely binds.

Regression analysis.—We now use a regression analysis to estimate the impact of the emissions trading treatment on pollution. We aggregate data to the plant-month level and run the following specification:

$$Pollution_{it} = \beta_1 Treatment_i + \alpha_t + \varepsilon_{it},$$

where $Treatment_i$ is a dummy variable equal to one for plants assigned to the emissions market treatment and α_t are year-month fixed effects.¹¹ Standard errors are clustered at the plant level.

Table 3 reports the results. Treatment markedly reduced pollution emissions, by 20% to 30% depending on the estimate. Columns 1 to 4 use pollution series with no imputation across plant-months. The first pair of columns is unweighted and the second pair is reweighted with probability weights representing the inverse probability of a plant reporting emissions. In column 2, the es-

¹¹We favor this simple specification over difference-in-differences specifications because CEMS data reporting is very sparse in the period before the market started (Figure A1).

timated treatment effect of the market on log emissions load is -0.193 log points (standard error 0.0763) with no imputation. This estimate is very similar with inverse probability weighting (column 4). In columns 5 to 8 we use alternate imputation rules for pollution emissions across plant-month, and generally find higher estimates. This reflects the fact that imputations tend to replace missing control group observations for log particulate emissions load with values higher than the mean among control plants that reported. Rule A, in columns 5 and 6, imputes a stack at its mean emissions from other times in the experiment when emissions are not observed in a given month. Rule B, in columns 7 and 8, imputes a stack at the monthly mean emissions load of its own treatment group for that month when a stack-month observation is missing. Again the odd-numbered columns are unweighted and the even-numbered columns weighted by the inverse probability of reporting. The treatment effect on pollution is -0.282 log points (standard error 0.0744) under Rule A and -0.316 log points (0.0567) under Rule B. Emissions reporting for much of the experiment was lower in the control group.

In summary, treatment caused emissions to decline in the treatment relative to the control group. The estimated decline is observed both in the raw data and under plausible alternative assumptions on missing pollution readings. The estimated magnitude of the pollution decline is similar to that targeted by the regulator.

5.2 Plant costs

In principle, plants can reduce pollution emissions through changes in their level of output, input mix, fuel type or usage, or end-of-pipe abatement expenditures. The relevant abatement action and, therefore, costs depend on the pollutant and level of abatement required. For particulate matter, end-of-pipe abatement equipment can be highly effective. The most common abatement equipment types in our sample are designed to remove 80% (cyclone), 94% (scrubber) and 99% (bag filter) of suspended particulates from stack gas. Whether a given piece of equipment achieves close to this design efficacy depends on how well it is maintained and how reliably it is run.

Input costs at the plant level.—Table 4 shows estimates of treatment effects on total plant costs during our endline survey. Panel A shows measures of total costs and Panel B of costs specific to the boiler house, the part of the plant where the boiler and other fuel-consuming equipment are housed and fuel is consumed to generate heat input. The different columns show the measure of total cost (column 1) and then costs for respective factor inputs: capital (column 2), labor (3), electricity (4), fuel (5) and materials (5). The average plant in the control has annual costs of some USD 1.2m (panel A, column 1).¹² Overall we find no significant effect of the emissions market treatment on plant costs (panel A, column 1), though this estimate is imprecise due to the large variance of plant materials costs. Within the boiler house, we again find no significant effect of the treatment on plant costs (panel B, column 1). The average control plant spends USD 580k per year in the boiler house. The estimated treatment effect is to increase costs by USD 11k (standard error USD 26k). Hence the point estimate of the treatment effect on costs is small and statistically insignificant.

Abatement capital costs.—Table 5 next tests for whether the treatment had any effect on abatement capital investment. Our sample consists of very large plants with millions of dollars in costs and sales. Recognizing that plant costs overall include many factors unrelated to pollution, the survey sought to specifically measure investment in abatement equipment. The dependent variables across panels are: an indicator for whether a plant has any of a given type of abatement equipment (panel A); the number of such devices a plant has (panel B); the estimated capital cost of the devices installed (panel C). The columns refer to different types of air pollution control devices.

There are two findings. First, even in the control group, every plant has installed abatement equipment, and often a lot of it. Cyclones and bag filters are designed to target particulates specifically. The percentage of plants with a cyclone is 95% and with a bag filter 85% (panel A, control

¹²This measure is incomplete because it omits capital costs. Plants are reluctant to report accurate measures of capital investment because plant size classifications and therefore regulatory stringency are based on total investment in property, plant and equipment. We could therefore gather plant-level capital records for too few plants to be useful as an outcome variable. We do separately observe abatement capital.

mean). Plants may have multiple devices when they have multiple pieces of fuel-consuming equipment.¹³ Plants on average have 1.9 cyclones and 1.5 bag filters (panel B, control mean). Second, we cannot reject that there was no change in the cost of abatement capital for treated plants (panel C, column 1). The point estimate is that treatment decreased abatement capital by USD 3,467 (standard error 3,089) against an installed base of USD 44 thousand. There is some evidence of a small change in the composition of abatement capital towards cheaper devices. We estimate marginally statistically significant, but economically very small, increases in abatement capital for cyclones and bag filters. These devices are specific to particulate abatement and lower cost than other equipment. Total costs decline because of a negative point estimate in abatement capital for electrostatic precipitators (ESPs, column 5). An ESP, or dry scrubber, removes particles from stack gas through an electric charge and is much more costly to install and operate than other devices.

Overall our estimates suggest negligible changes to abatement capital and to other inputs to abatement measured at the plant level. Plants have annual costs on the order of millions of US dollars and we observe shifts in abatement capital on the order of thousands of dollars and we cannot reject that there is no change in abatement capital. This result echoes our prior argument that abatement capital is not a constraint on abatement in this setting because the command-and-control regime mandates equipment installation (Duflo et al., 2018). The constraint instead is getting plants to maintain and run the equipment they do have.

Abatement variable costs.—The importance of equipment operation argues that the costs of the market should be measured through variable costs of abatement. Measuring variable costs is difficult because plant inputs used for abatement, such as labor and electricity, are used both for abatement and other purposes and there is no clear delineation of their intention. Electricity, for example, is seldom separately metered for abatement equipment.

Our approach is therefore to use permit prices in the emissions market to estimate the variable

¹³Most commonly, by far, plants will operate both a boiler and a thermic fluid heater or “thermopack.” The boiler burns fuel to heat steam for use in industrial processes like dyeing cloth. The thermic fluid heater heats oil or lubricant for use in keeping process machines running smoothly. Both of these machines send exhaust gas to stacks that are monitored by CEMS. Even if there is a single stack, each machine may have separate abatement equipment to reduce emissions from its fuel use before connecting to the common stack.

costs of abatement. As described in Section 4, the market cleared at prices between the floor of INR 5 per kg and INR 15 per kg, though average bid prices ranged as high as INR 45 per kg. Prices for bids can be benchmarked against data on abatement costs from Indian air pollution control device vendors. Appendix Table A6 presents estimates for abatement costs under *ideal* operating conditions for four kinds of air pollution control devices under four hypothetical plant configurations. Appendix Table A7 presents estimates of abatement costs assuming that plants are already operating a cyclone. Engineering abatement costs vary widely depending mainly on (i) the scale of the plant (ii) the type of equipment that is on the margin (iii) whether any other equipment is assumed to operate simultaneously. At the low end of the scale, average (marginal) cost of abatement for running a cyclone for a large plant that operates no other equipment is estimated to be in the range of 0.60 (0.20) INR per kg. If a plant is already running some equipment, which is more likely, then average (marginal) abatement costs for a bid size plant to operate a bag filter are 10 (3) INR per kg. If a plant is small and already running a cyclone, average (marginal) abatement costs to run a dry scrubber are as high as 71 (20) INR per kg. All of these numbers assume ideal operation of equipment and are therefore likely lower bounds on costs in each scenario. The bid and permit prices suggest that marginal abatement costs in the market were perceived to be within the wide range marked out by engineering estimates.

6 Model of pollution abatement

The analysis above shows that treatment plants reduced pollution in the emissions market relative to control plants. The reduction in pollution was achieved with little or no change in abatement capital. A complete estimate of how the market affected costs should also cover variable abatement costs. The present section therefore takes a model-based approach to infer marginal abatement costs from bidding data in the permit market.

The goal of the model is twofold. First, to incorporate the permit bidding data into the evaluation as a way to measure variable costs. Second, to separate the respective effects of the treatment

on emissions and variable abatement costs. This separation is of interest because the emissions market changes both the level of emissions and how emissions are allocated across plants. Bidding data is only available for the treatment group in which plants bid in the market. We therefore use the data in the treatment group to estimate marginal abatement costs. We then fit the model to out-of-sample regulatory stringency using the distribution of emissions rates in the control group. This comparison allows us to use the model to measure the effects of the experimental treatment on emissions and costs separately.

6.1 Model specification

Abatement technology.—Plant i chooses the level of variable abatement expenditures Z_{it} in each compliance period $t = 1, 2, \dots, 10$. Abatement expenditures could include running abatement equipment more often, changing inputs like filters or chemicals more often, or devoting more labor to the maintenance and operation of a machine. Plants differ in their total heat output H_i . Heat output is the total fuel-burning capacity of a boiler, analogous to the horsepower of a car engine, and is a relevant measure of scale for air pollution emissions. Plants may also differ in other characteristics such as their abatement capital stock. The plant spends a fixed cost Z_{i0} to maintain its abatement capital.

We let $E_{it}(Z_{it})$ be the level of emissions as a function of expenditures. Assume that $E' < 0, E'' > 0$; emissions are decreasing as a function of expenditures but at a rate that decreases in magnitude as expenditures grow.

Regulation under the emissions market.—A plant is allocated permits A_{it} under emissions trading. In the emissions market the allocation rule gave plants permits totaling 80% of the market cap in proportion to their heat output capacity, $A_{it} \propto H_i$. Let P_t be the equilibrium price of permits. Assume that this equilibrium price is known to the plant. The plant seeks to minimize the total cost of compliance:

$$\min_{Z_{it}} Z_{i0} + Z_{it} + P_t(E_{it}(Z_{it}) - A_{it}). \quad (1)$$

The first-order condition for the plant's problem can be written

$$-\frac{\partial Z_{it}(E_{it})}{\partial E_{it}} = MAC(E_{it}) = P_t.$$

This condition is the familiar one that the marginal abatement costs of the plant at the chosen emissions level equal the permit price. This equation will have a unique solution for E_{it}^* under our assumptions on the $E(\cdot)$ function. When all plants equalize their marginal abatement costs the market as a whole reduces emissions at the lowest possible aggregate cost. The level of emissions depends on neither the plant's fixed costs of abatement Z_{i0} nor the allocation of permits A_{it} . The former assumption is justified in our setting as abatement equipment is already installed in the status quo and therefore Z_{i0} is sunk.

Regulation under command and control.—The command and control regime is a concentration limit, i.e. a maximal allowable concentration that the plant is allowed to emit in stack gas. We represent this limit as a plant-specific emissions rate

$$R_{it} = E_{it}/H_i \sim F.$$

The plant spends Z_{it}^c such that $E_{it}(Z_{it}^c) = H_i R_{it}$. This regulatory regime admits two sources of inefficiency. First, emissions rates are limited rather than emissions levels, whereas the damages from pollution depend on emissions levels. Depending on how abatement costs scale with plant size, large plants may have too strong or too weak an incentive to abate pollution. Second, we assume emissions rates are drawn from some distribution F . This assumption is based on the observation, both in our control group and in Duflo et al. (2018), that there is non-compliance and fairly wide dispersion in emissions concentrations under the status quo regime. Our prior work found the regulator had some information on which firms were the most polluting but that this explained a small part of the overall variance in pollution. If, by extension, the enforced emission rates are not much related to marginal abatement costs, then different plants will be expected to

have different marginal abatement costs at their assigned emissions rate.

6.2 Estimation

The main object of estimation is plant abatement technology and particularly the heterogeneity across plants in marginal abatement costs. We adopt a simple form for the abatement cost function to allow for simple estimation of marginal abatement costs using data from plant bids. We assume the abatement cost function

$$Z_{it}(E_{it}) = e^{\beta_0 + \tilde{\xi}_{it}} H^{\beta_2} \left(\frac{1}{\beta_1 + 1} \right) \left(\bar{E}_i^{\beta_1 + 1} - E_{it}^{\beta_1 + 1} \right), \quad \beta_1 \in (-1, 0). \quad (2)$$

The constant \bar{E}_i represents emissions for a plant of size $H = 1$ when no variable abatement expenditures are made. Even with zero expenditures, pollution emissions will never be infinite, but will approach some high, uncontrolled level. The parameter restriction on β_1 ensures that costs are strictly decreasing in E_{it} and marginal costs of abatement are decreasing in emissions (increasing in abatement).

The abatement cost function implies that the log of marginal abatement cost is

$$\log MAC(E_{it}) = \beta_0 + \beta_1 \log E_{it} + \beta_2 \log H_i + \tilde{\xi}_{it}. \quad (3)$$

The parameter β_1 is the elasticity of marginal abatement costs with respect to emissions. With $\beta_1 < 0$ abatement costs are decreasing in emissions.

There are two difficulties in the estimation of (3). First, the marginal abatement cost of a plant is typically not observed. Second, emissions are endogenous to abatement costs. We expect ordinary least squares estimates of equation (3) would yield positively-biased estimates of β_1 .¹⁴

Our approach to estimation is to use within-period variation in plant bids to estimate the

¹⁴The level of emissions E_{it} is endogenous to the abatement cost shock $\tilde{\xi}_{it}$. Plants with a high abatement cost shock will choose high levels of emissions E_{it} such that $\mathbb{E}[\tilde{\xi}_{it}|E_{it}] \neq 0$ even if $\mathbb{E}[\tilde{\xi}_{it}] = 0$.

marginal abatement cost function. We estimate a version of (3) as

$$\log b_{itk} = \beta_1 \log E_{itk} + \xi_{it} + \varepsilon_{itk}, \quad \mathbb{E}[\varepsilon_{itk} | E_{itk}, \xi_{it}] = 0. \quad (4)$$

The plant-period effects $\xi_{it} = \beta_0 + \beta_2 \log H_i + \tilde{\xi}_{it}$ subsume the effect of plant heat output on emissions. Here E_{itk} is the emissions level (permit holdings) the plant would have if bid k in period t were executed. The plant bids its marginal cost plus an error term that is additively separable in logs. This assumption can be econometrically justified, in our setting, if plants form rational, unbiased expectations of their emissions at the time of bidding but have uncertainty about the exact emissions level.¹⁵ We find this assumption plausible in the emissions market we study. Plants had CEMS devices and could read the output but validated emissions for the market were released at a lag. Moreover, plants may not know at the start exactly how much they will operate in a compliance period. Plants therefore likely had some uncertainty about their contemporaneous emissions level up until the true-up period when all emissions had been recorded.

Equation (4) is our main estimating equation. Estimates by OLS will be unbiased under the given assumption that the *bid quantity* is uncorrelated with the unobserved determinants of bid prices conditional on plant-period effects. This specification allows that marginal abatement costs are unobservably higher for some plants and in some periods, but rules out that innovations in marginal costs of abatement within a period are related to emissions levels within a period. Variation in $\log E_{itK}$ comes from different bids that plants submit within the same compliance period as information about their own emissions and hence permit demand arrives.

We additionally estimate the stringency of regulation in the command and control regime. While *de jure* the status quo is a constant concentration (rate of emissions per gas volume) across plants i and time t , that is not the best representation of the status quo *de facto*. The command and control regime involves regulatory discretion in both inspections and penalties and varying concentration levels across the population of plants (Duflo et al., 2018). When the abatement function

¹⁵For example, assume firms anticipate emissions of $\tilde{E}_{itk} = E_{it} v_{itk}$ with $v_{itk} \perp E_{itk}, \xi_{it}$ and $\mathbb{E}[\log v_{itk}] = 0$. Then firms bidding their expected marginal costs yields the above specification (3) with a residual $\varepsilon_{itk} = \beta_1 \log v_{itk}$ based on the forecast error.

is convex, these details may matter, because a distribution of emissions rates across plants may yield a higher total abatement cost than a uniform rate, at the mean of that distribution, imposed on all plants. We therefore use several alternative representations of the status quo to capture the fact that emissions rates vary in the command-and-control regime.

The simplest command and control regime we estimate is a log-normal distribution of emissions rates $\log R_{it} \sim \mathcal{N}(\mu_t, \sigma_t)$ with a mean and standard deviation fit separately in each compliance period. This regime is probably too simple to represent the status quo, because larger plants are cleaner: the status quo emissions rate is declining in heat output capacity H_i . This fact is consistent with a regulatory regime that inspects large plants more often and so imposes greater expected penalties on large plants for high emissions rates. We therefore model a more flexible regime where the emissions rate depends on plant heat capacity

$$\log R_{it} = \beta_{0t} + \beta_{1t} \log H_i + \varepsilon_{it}.$$

We fit this regression in the control group separately in each compliance period. We then set counterfactual emissions rates for the treatment group plants, had they been regulated like the control plants, using the predicted values from this regression as $\widehat{E}_{it} = H_i \exp(\widehat{\log R_{it}})$. In some regimes we factor in variation in emissions rates in the command and control regime. We do this by multiplying the predicted log emissions rate by plant-specific heterogeneity using draws from a log normal distribution fit to the variance of $\widehat{\varepsilon}_{it}$ in each period.

7 Model results

This section estimates the abatement cost model and uses it to characterize the cost savings from emissions trading.

7.1 Marginal abatement cost function estimates

Table 6 presents estimates of variants of the specification (4). The data set is at the plant-period-bid level for all bids offered in the treatment group in all compliance periods. The first column controls for heat output and day-of-period. The second through third columns add period fixed effects, plant and period fixed effects, and plant-period fixed effects, respectively.

Our preferred estimate in Table 6, column 4 is that the elasticity of bid prices (marginal abatement costs) with respect to emissions is -0.491 (standard error 0.0897). As emissions increase, the marginal cost of reducing emissions decreases. This estimate satisfies the parameter restriction in (2), which was not imposed in estimation. It also accords with intuition that, for example, the marginal cost of further improving the operation of abatement equipment rises as emissions are cut. The pattern of results across columns 1 through 4 shows the importance of using within-period data to estimate this elasticity. Our preferred estimate uses plant-period fixed effects. Estimates with basic controls or only period fixed effects are severely positively biased and statistically indistinguishable from zero (columns 1 and 2). Even two-way plant-and-period fixed effects yield an elasticity less than half as great in magnitude as our preferred estimate (column 3). The positive bias of these estimates accords with intuition that plant bid quantities (emissions) across periods are correlated with plant-period shocks to marginal abatement cost.

The Table 6 estimates impose one key sample restriction: we use only bids in the first half of each compliance period to estimate the elasticity. The rationale for this restriction is that the problem (1) assumes plants can make a choice of whether to comply by reducing emissions or buying permits. As the end of a compliance period approaches, plant emissions within a period are largely sunk, so this choice is no longer possible. (At the extreme, during the true-up period when the compliance period has ended, emissions are fixed.) Consistent with the idea that plants can no longer choose emissions flexibly as a compliance period ends, Figure A6 shows that the elasticity of bids with respect to emissions quantity that we estimate goes from being negative and significant early in the compliance period to close to zero as the time left within a compliance period wanes.

7.2 Counterfactual permit market equilibrium

The parameters of the abatement cost function (4) allow us to solve for an equilibrium in the permit market for any set of plants and level of the emissions cap. Each plant sets emissions such that its expected marginal abatement cost equals the permit price, which yields emissions

$$E_{it}(P_t) = P_t^{1/\hat{\beta}_1} e^{-\hat{\xi}_{it}/\hat{\beta}_1}. \quad (5)$$

For $\hat{\beta}_1 < 0$, as estimated, emissions are decreasing as a function of the permit price and increasing as a function of the plant-specific marginal abatement cost shock. The permit market equilibrium price P_t^* for a cap \bar{Q}_t is given by

$$E_t(P_t^*) = \sum_i E_{it}(P_t^*) = \bar{Q}_t. \quad (6)$$

The equilibrium price is unique given that emissions monotonically decrease in price. The emissions for any plant are given by (5) evaluated at P_t^* and their abatement cost by (2). The same estimated cost functions are used to evaluate emissions and costs under alternate command-and-control regimes described above.

7.3 Model fit

Figure 7 shows the fit of the model to market-clearing prices by compliance period. The dashed black line shows data on mean bids; the dotted black line shows data on mean clearing prices; the solid blue line shows the model simulation of market-clearing prices.

The model's predicted market prices have a very good fit to fluctuations in mean bid prices, which are used to estimate abatement costs. Bids and simulated prices are relatively high in the first period, fall to about INR 8-10 for periods two through six, and then rise in the final four compliance periods to about INR 10-12.

The model fits less well to market-clearing prices: the model prediction for market prices in the first period through sixth periods are consistently above actual clearing prices. This gap in fit

in explicable as a feature of the simplicity of the static model that we use to represent the dynamic permit market. The data series shown in the plot is the average market-clearing price over each compliance period. The model, by contrast, is fit only to bids offered in the first half of each compliance period, when plants face a meaningful choice between abatement and purchasing permits. The model can therefore be thought of as representing plants' expected abatement costs at the beginning of each compliance period. In period one, when all plants were new to the market, expected abatement costs were relatively high at the start (see Figure 3), so bids were high before the permit market converged to a lower clearing-price at the end of the period. In compliance periods two through six a similar convergence happened relative to a smaller initial gap. In compliance periods seven through ten permit bids early in the market were more reflective of realized market-clearing prices. The model of expected abatement costs therefore fits firm bids well in all periods but actual market-clearing prices better during later compliance periods.

7.4 Counterfactual analysis

The main purpose of the model is to quantify the change in variable cost savings due to moving from a command and control regime to an emissions market.

Figure 8 compares total variable abatement costs under the two regimes. The command and control regime uses a capacity-based emissions rate with error to set emissions targets for each plant. In this regime counterfactual emissions for each treatment plant, had they been in the control group, are assigned using the fitted equation (5) augmented by an idiosyncratic emissions shock for each plant. The various levels of emissions on the horizontal axis are achieved by scaling up or down the distribution of emissions targets proportionally. The emissions market regime sets an emissions cap at each level of emissions on the horizontal axis. The cost curves for both regimes are drawn for the estimated levels of marginal abatement cost shocks in compliance period four. The dashed horizontal and vertical lines intersect the command and control abatement cost curve at approximately the level of aggregate emissions observed in the control group.

The Figure 8 total abatement cost curves show that abatement costs rise only slowly in response to reductions in pollution, under either regime. On the command-and-control cost curve, the arc

elasticity of total variable abatement costs with respect to total emissions, evaluated around the control level of emissions, is -0.20, indicating that costs rise fairly slowly with cuts in emissions. This elasticity depends on, but is naturally lower than, the elasticity of marginal abatement costs estimated in Table 6. Total abatement costs include inframarginal abatement costs and are therefore less elastic to emissions than marginal costs.

At any given level of pollution, the emissions market achieves lower variable abatement costs than the command and control regime. The market has costs about 10-15% lower than the command and control regime during this compliance period at emissions levels in the range of the data. Treatment emissions were typically capped around 170 tons. An estimated treatment effect of 30% on emissions (Table 3) then implies control emissions of approximately 240 tons. The model enables a comparative static comparison of what emissions levels could be achieved holding constant control abatement costs. The filled black circle represents emissions under the status quo command and control regime. The corresponding cost under the emissions market is 13% lower (moving down the vertical dashed line). Because total abatement costs are not very elastic with respect to emissions, the cost curves imply that the emissions market, at the same variable abatement costs as under the command and control regime, would cut total emissions nearly in half (moving left along the horizontal dashed line). Alternately, a range of outcomes with both lower emissions and lower costs are available along the arc of the emissions trading cost curve between the horizontal and vertical dashed lines.

Table 7 uses the same comparative static analysis to compare costs under different regulatory regimes across all compliance periods. Each row in the table represents a different regulatory regime. The command and control regimes in rows two through five differ in how they assign emissions targets to plants. The *Constant emissions rate* regime sets one emissions rate for all plants, which is similar to the *de jure* regime of a single concentration standard. The *Constant emissions rate with error* regime allows log-normal shocks to plant emissions around this emissions rate. The *Capacity-based rate* regime sets the emissions rate as a function of capacity according to (5). The *Capacity-based rate with error* regime allows log-normal shocks to plant emissions around

the capacity-based rate. We use this regime as our default representation of the status quo. Finally, *Capacity-based emissions rate with correlated error* allows a log-normal shock to plant emissions around the capacity-based rate and draws the shocks such that they are slightly negatively correlated ($\rho = -0.1$) with estimated plant-period marginal abatement costs. This implies that high-cost plants will have somewhat lower emissions targets. We introduce this correlation to capture, in a simple way, the observation that the regulator does have some information about plant emissions and targets more polluting plants more aggressively (Duflo et al., 2018).

There are three main findings from Table 7. First, total variable abatement costs are 12% higher under the status quo (column 3, row 5) than under emissions trading (column 3, row 1), at the treatment emissions level (170 tons per month, which was the cap in the majority of compliance periods). Since the gap between market and command and control abatement costs increases in emissions (Figure 8), this cost difference rises to 15% at the control level of emissions.

Second, the cost differences among alternative representations of the command and control regime are smaller than the difference in cost with respect to the market regime. The differences in costs in the command and control regime can be understood as the result of two forces: (i) heterogeneity in emissions rates interacting with convex abatement costs (ii) scale effects. On heterogeneity, command and control regimes that allow idiosyncratic shocks across plants, as seen in the data, have higher costs than regimes that do not because abatement costs are convex. This convexity pushes up costs for plants that are assigned lower rates of emissions more than it reduces costs for plants with higher rates. The status quo costs would therefore be INR 10.37m per month under a flat capacity-based emissions rate but INR 10.70m per month under a heterogeneous capacity-based rate. On scale, we estimate that the emissions rate in the control is decreasing fairly steeply in plant heat capacity (not reported). Marginal abatement costs per unit of emissions *also* decrease with scale; however, they do not decrease as quickly as do emissions rates. As a result larger plants tend to have higher marginal abatement costs in regimes that condition emissions rates on scale, raising overall costs. For example, in column 5, a constant emissions rate with error yields costs of INR 10.58m per month whereas a capacity-based rate with error has costs of INR

10.70m per month.

The third finding is that the cost difference between the regimes is large enough that the estimated abatement costs in the treatment group—at a substantially lower level of emissions—are lower than the abatement costs in the control group. The treatment group has abatement costs of INR 10.08m per month for emissions of 170 tons (column 2, row 1) whereas the control group has abatement costs of INR 10.70m per month for emissions of 240 tons (column 5, row 5). The treatment group therefore has total variable abatement costs 6 percent below the control despite that treatment emissions are 30 percent lower than in the control group.

8 Conclusion

High regulatory costs are one reason why air pollution emissions may be high in developing countries. We evaluate the effects of a new emissions market on particulate matter air pollution and abatement costs. Our estimates show that the introduction of emissions trading reduced pollution by between 20% and 30%, relative to a control group of plants that remained in the command-and-control regime. We estimate that this reduction in pollution was achieved for no additional expenditure on abatement capital and at low permit prices within the range of engineering estimates of abatement costs.

We use permit trading data from the emissions market to expand our measurement of costs to include variable abatement costs, which are otherwise difficult to measure. The level and elasticity of marginal abatement costs are estimated in the treatment group under the assumption that plants bid their expected marginal abatement cost in the market. Control plants did not trade permits as they did not participate in the market. It is therefore not possible to simply compare permit prices across the two treatment arms. Instead, we apply the model estimated in the treatment group to calculate abatement costs at the control level of emissions. Our model estimates imply that the emissions market reduced variable abatement costs at the same time that it sharply cut emissions. At constant cost, the emissions market would have reduced emissions load by nearly half.

The results of the present experiment, in combination with prior work, suggest that a large part of the costs of environmental regulation are due not to abatement costs per se, which are the object of most theory on regulation, but to the fixed costs of monitoring and enforcing regulation. Dufflo et al. (2018) describe and model in detail how enforcement is conducted in the command and control status quo. The regulatory regime of visiting plants, taking observations and levying penalties has extremely high costs, but even in that paper we estimated that abatement equipment had high efficacy when it was run. Here, similarly, the emissions market in Gujarat took years of preparation to set up and required plants to install a continuous emissions monitoring systems (CEMS). These investments came on top of high levels of pre-existing investment in abatement equipment itself. Taking these estimates of fixed regulatory costs as sunk, our model estimates imply that variable abatement costs impose an additional expense on the order of USD 1,000 per plant-month, which is very small next to the total costs of the plants concerned.

Many plants in India and other developing countries emit pollution at high rates that contribute to low levels of air quality. Historically, the United States has been able to dramatically reduce pollution emissions from manufacturing by relying on abatement equipment without cutting output (Shapiro and Walker, 2018). China has recently taken a turn to more stringent environmental regulation (Greenstone et al., 2021). Our estimates suggest that India, with an appropriate level of investment in monitoring and a flexible regulatory regime, could also cut emissions at a low cost.

References

- Baylis, Patrick, Michael Greestone, Kenneth Lee, and Harshil Sahai.** 2022. “Pollution Masks and the Demand for Clean Air: Experimental Evidence from Delhi.”
- Blackman, Allen, Zhengyan Li, and Antung A Liu.** 2018. “Efficacy of command-and-control and market-based environmental regulation in developing countries.” *Annual Review of Resource Economics*, 10: 381–404.
- Borenstein, Severin, James Bushnell, Frank A Wolak, and Matthew Zaragoza-Watkins.** 2019. “Expecting the unexpected: Emissions uncertainty and environmental market design.” *American Economic Review*, 109(11): 3953–77.
- Burtraw, Dallas, David A Evans, Alan Krupnick, Karen Palmer, and Russell Toth.** 2005. “Economics of Pollution Trading for SO₂ and NO_x.” *Annu. Rev. Environ. Resour.*, 30: 253–289.
- Central Pollution Control Board.** 2013. “Specifications and Guidelines for Continuous Emissions Monitoring Systems (CEMS) for PM Measurement with Special Reference to Emission Trading Programs.” *CPCB/e-PUBLICATION/2013-14*.
- Coase, R. H.** 1960. “The Problem of Social Cost.” *The Journal of Law & Economics*, 3: 1–44. Publisher: [University of Chicago Press, Booth School of Business, University of Chicago, University of Chicago Law School].
- Dales, John Harkness.** 1968. *Pollution, property & prices: an essay in policy-making and economics*. Toronto: University Press.
- Davis, Lucas W.** 2008. “The effect of driving restrictions on air quality in Mexico City.” *Journal of Political Economy*, 116(1): 38–81.
- Dufo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan.** 2010. “Towards an Emissions Trading Scheme for Air Pollutants in India.” *MIT Center for Energy and Environmental Policy Research*, 10(11).
- Dufo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan.** 2013. “Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India.” *The Quarterly Journal of Economics*, 128(4): 1499–1545.
- Dufo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan.** 2018. “The value of regulatory discretion: Estimates from environmental inspections in India.” *Econometrica*, 86(6): 2123–2160.
- Ebenstein, Avraham, and Michael Greenstone.** 2022. “Childhood Exposure to Particulate Air Pollution, Human Capital Accumulation, and Income: Evidence from China.” *Working Paper*.
- Ebenstein, Avraham, Maoyong Fan, Michael Greenstone, Guojun He, and Maigeng Zhou.** 2017. “New evidence on the impact of sustained exposure to air pollution on life expectancy from China’s Huai River Policy.” *Proceedings of the National Academy of Sciences*, 114(39): 10384–10389. Publisher: Proceedings of the National Academy of Sciences.
- Ellerman, A Denny, Paul L Joskow, Richard Schmalensee, Juan-Pablo Montero, and Elizabeth M Bailey.** 2000. *Markets for clean air: The US acid rain program*. Cambridge University Press.
- Fowle, Meredith, Stephen P Holland, and Erin T Mansur.** 2012. “What do emissions markets deliver and to whom? Evidence from Southern California’s NO_x trading program.” *American Economic Review*, 102(2): 965–93.

- Greenstone, Michael, and B. Kelsey Jack.** 2015. “Envirodevonomics: A Research Agenda for an Emerging Field.” *Journal of Economic Literature*, 53(1): 5–42.
- Greenstone, Michael, and Rema Hanna.** 2014. “Environmental regulations, air and water pollution, and infant mortality in India.” *American Economic Review*, 104(10): 3038–72.
- Greenstone, Michael, and Ted Gayer.** 2009. “Quasi-experimental and experimental approaches to environmental economics.” *Journal of Environmental Economics and Management*, 57(1): 21–44.
- Greenstone, Michael, Christa Hasenkopf, and Kenneth Lee.** 2022. “Air Quality Life Index Annual Update.”
- Greenstone, Michael, Guojun He, Ruixue Jia, and Tong Liu.** 2022. “Can Technology Solve the Principal-Agent Problem? Evidence from China’s War on Air Pollution.” *American Economic Review: Insights*, 4(1): 54–70.
- Greenstone, Michael, Guojun He, Shanjun Li, and Eric Zou.** 2021. “China’s War on Pollution: Evidence from the First Five Years.” National Bureau of Economic Research Working Paper 28467.
- Greenstone, Michael, Janhavi Nilekani, Rohini Pande, Nicholas Ryan, Anant Sudarshan, and Anish Sugathan.** 2015. “Lower pollution, longer lives: life expectancy gains if India reduced particulate matter pollution.” *Economic and Political Weekly*, 40–46.
- He, Guojun, Shaoda Wang, and Bing Zhang.** 2020. “Watering down environmental regulation in China.” *The Quarterly Journal of Economics*, 135(4): 2135–2185.
- Isen, Adam, Maya Rossin-Slater, and W. Reed Walker.** 2017. “Every Breath You Take—Every Dollar You’ll Make: The Long-Term Consequences of the Clean Air Act of 1970.” *Journal of Political Economy*, 125(3): 848–902. Publisher: The University of Chicago Press.
- Ito, Koichiro, and Shuang Zhang.** 2020. “Willingness to Pay for Clean Air: Evidence from Air Purifier Markets in China.” *Journal of Political Economy*, 128(5): 1627–1672. Publisher: The University of Chicago Press.
- Jack, Kelsey, Seema Jayachandran, Namrata Kala, and Rohini Pande.** 2022. “Money (Not) to Burn: Payments for Ecosystem Services to Reduce Crop Residue Burning.”
- Jayachandran, Seema.** 2022. “How economic development influences the environment.” *Annual Review of Economics*, 14: 229–252.
- Lall, Ramona, Michaela Kendall, Kazuhiko Ito, and George D Thurston.** 2004. “Estimation of historical annual PM_{2.5} exposures for health effects assessment.” *Atmospheric Environment*, 38(31): 5217–5226.
- Martin, Ralf, Laure B De Preux, and Ulrich J Wagner.** 2014. “The impact of a carbon tax on manufacturing: Evidence from microdata.” *Journal of Public Economics*, 117: 1–14.
- Martin, Ralf, Mirabelle Muûls, and Ulrich J Wagner.** 2020. “The impact of the European Union Emissions Trading Scheme on regulated firms: what is the evidence after ten years?” *Review of environmental economics and policy*.
- Montero, Juan-Pablo, Jose Miguel Sanchez, and Ricardo Katz.** 2002. “A market-based environmental policy experiment in Chile.” *The Journal of Law and Economics*, 45(1): 267–287.
- Oliva, Paulina.** 2015. “Environmental regulations and corruption: Automobile emissions in Mexico City.” *Journal of Political Economy*, 123(3): 686–724.
- Piette, Lauren.** 2018. “Improving environmental enforcement in India.” *University of Chicago Law School*.
- Shapiro, Joseph S, and Reed Walker.** 2018. “Why is pollution from US manufacturing declining?”

The roles of environmental regulation, productivity, and trade.” *American Economic Review*, 108(12): 3814–54.

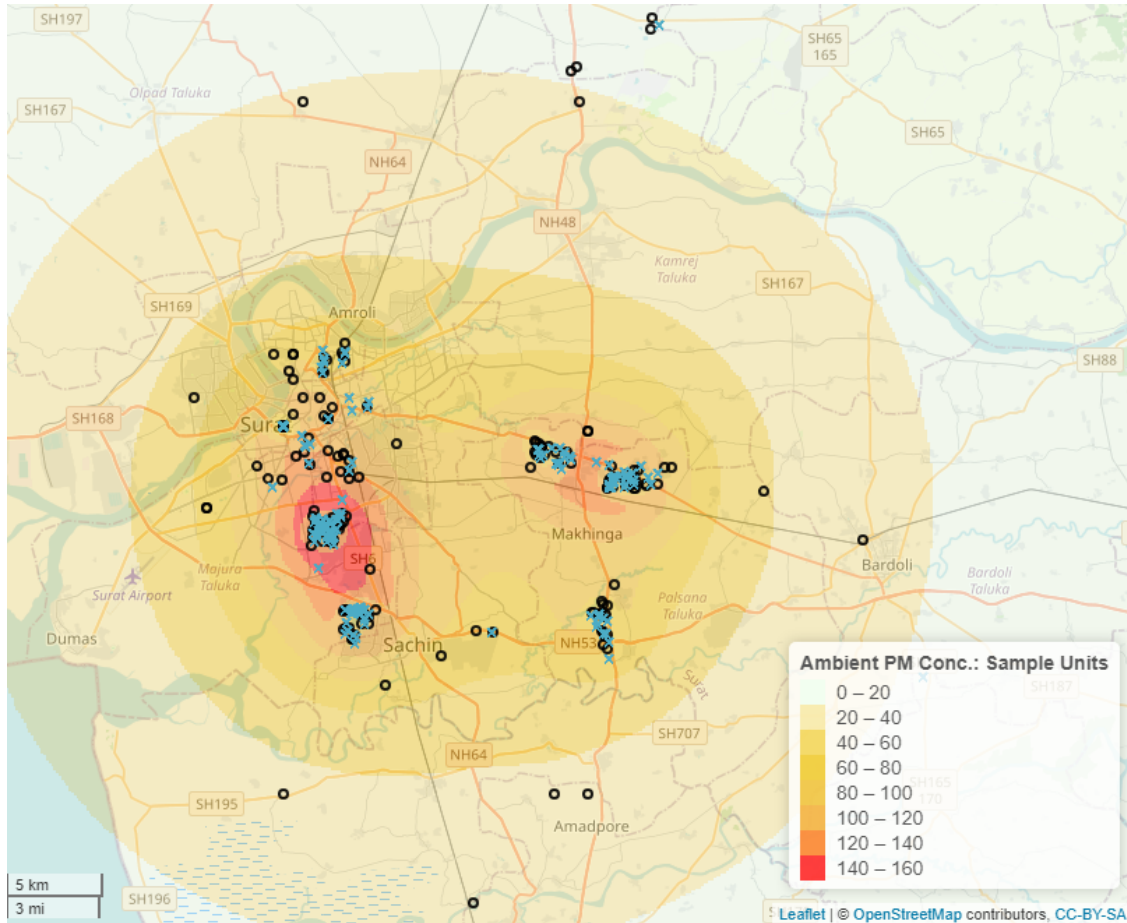
Stavins, Robert. 2003. “Experience with market-based environmental policy instruments.” Elsevier Handbook of Environmental Economics.

Sudarshan, Anant. 2023. “Monitoring Industrial Pollution in India.”

Zou, Eric Yongchen. 2021. “Unwatched pollution: The effect of intermittent monitoring on air quality.” *American Economic Review*, 111(7): 2101–26.

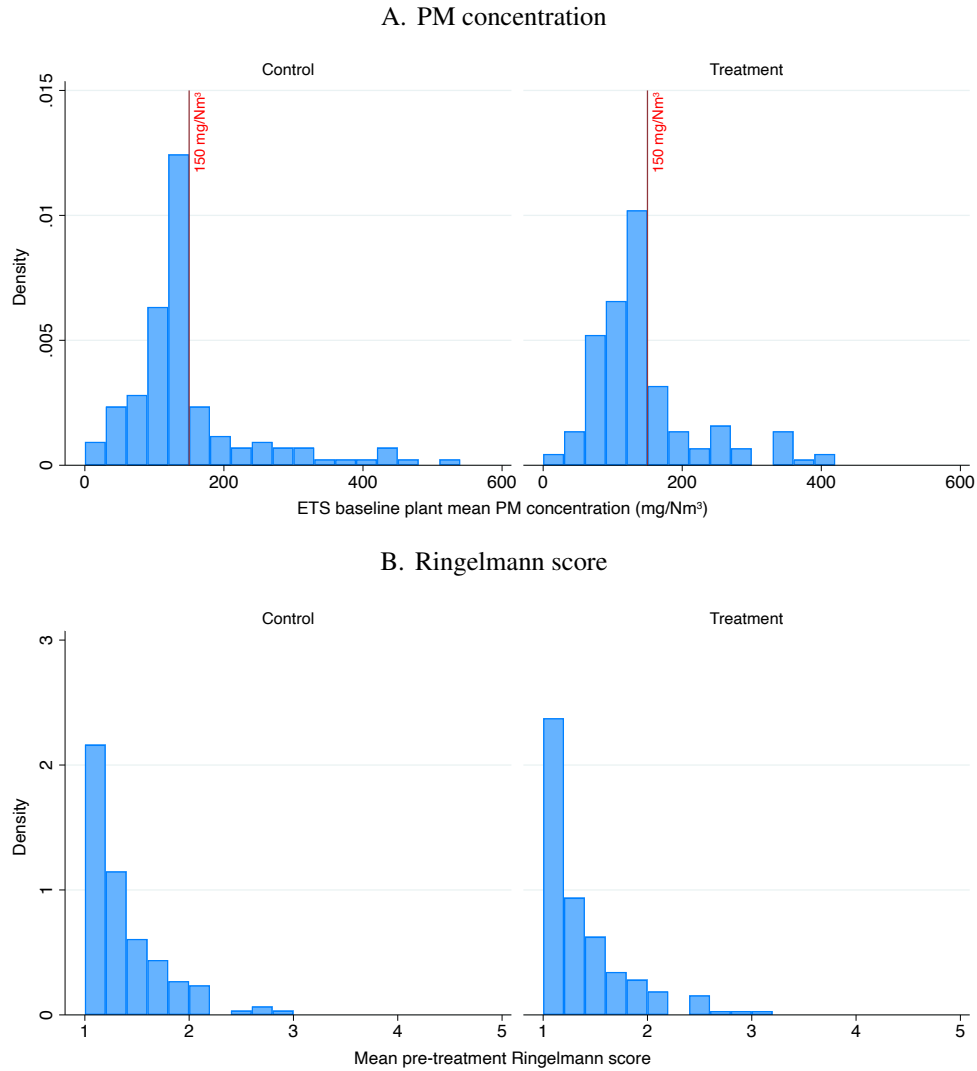
9 Figures

Figure 1: Modeled contribution of sample plants to ambient pollution levels at baseline



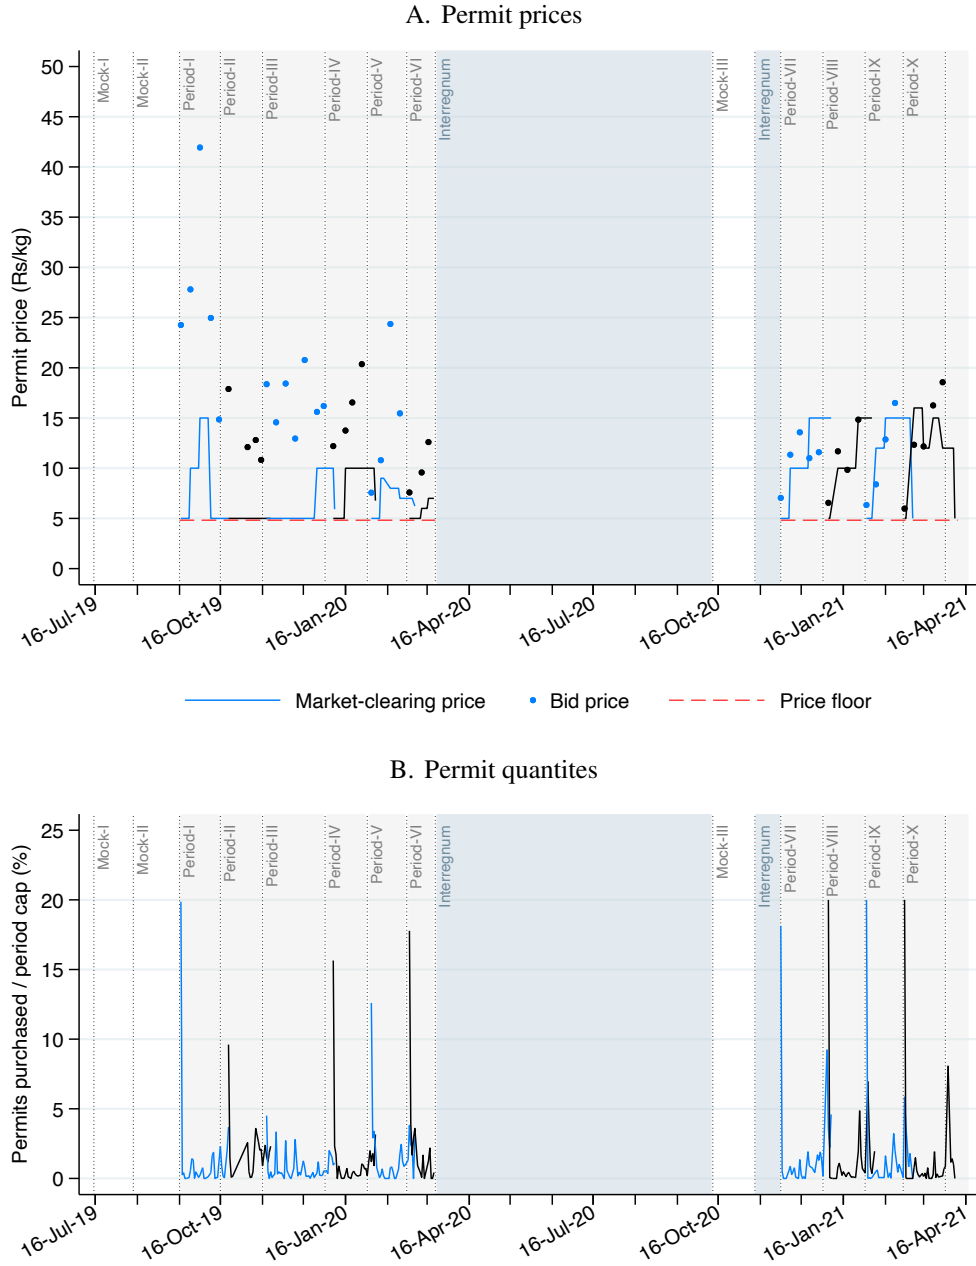
Note. The figure shows the contribution of PM emissions from sample plants to ambient PM concentrations in Surat, Gujarat in June 2019. Treatment plants are represented by blue \times markers and control plants by black \circ markers. 304 plants are shown, with 13 sample plants lying outside the bounds of the map. The ambient PM concentrations (in $\mu\text{g}/\text{m}^3$) are measured by passing CEMS emission rates for the sample plants into a simplified Gaussian dispersion model, comparable to the SCREEN3 model used by the US Environmental Protection Agency. The model used is based on the eddy diffusion theory, considering each plant to be a stationary point emitting source. The CEMS data informs each plant's emission rate (mass/time). The location and stack height of each plant determine the point source. The model uses simple assumptions on meteorological conditions such as constant emission rates and wind speeds.

Figure 2: Distribution of pollution before the experiment



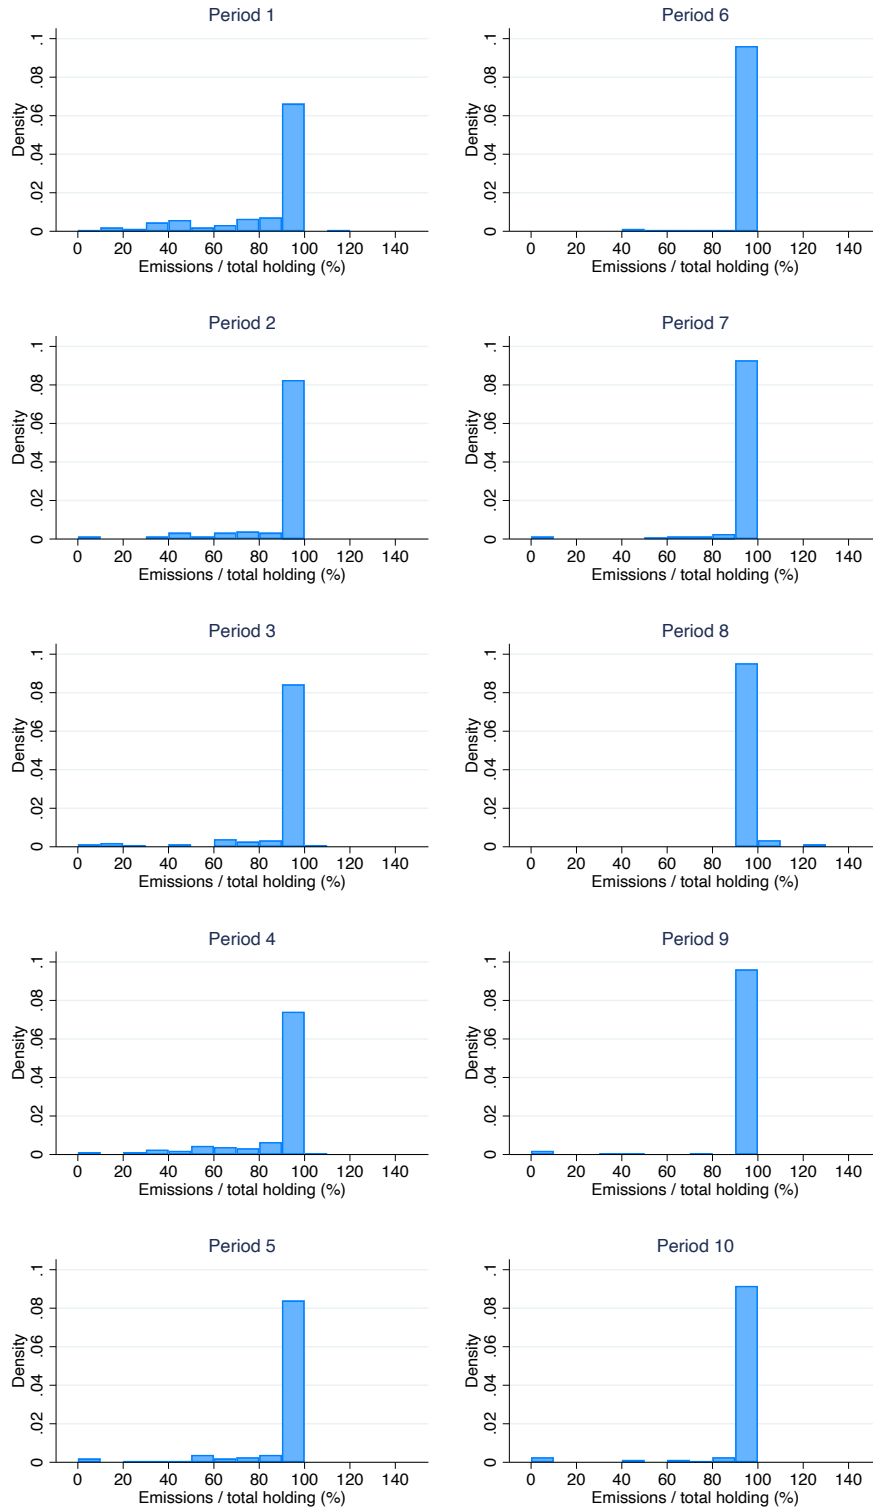
Note. The top panel presents the distributions of the mean PM concentration by treatment status as measured by manual iso-kinetic stack sampling at ETS baseline (December 2018 to January 2019). One PM sample was collected from each industrial stack by a third-party laboratory. We truncated the PM concentration at the 95th percentile, or 520 mg/Nm³. As a result, we dropped 14 observations. The red, vertical lines indicate the regulatory concentration standard of 150 mg/Nm³. An iso-kinetic measurement of PM concentration above this threshold indicates non-compliance under the status quo regime. At ETS baseline, 28% of sampled plants in the control group and 34% of sampled plants in the treatment group were reported to be out of compliance with the 150 mg/Nm³ regulatory limit. The bottom panel presents the distributions of the mean pre-treatment Ringelmann score based on four rounds of Ringelmann surveys conducted from April 2019 to June 2019, by treatment status. The Ringelmann score is a scale for measuring the apparent density of smoke. The scale has five levels of density. Score 1 to 5 correspond to an opacity of 20%, 40%, 60%, 80% and 100%.

Figure 3: Permit prices and quantities purchased



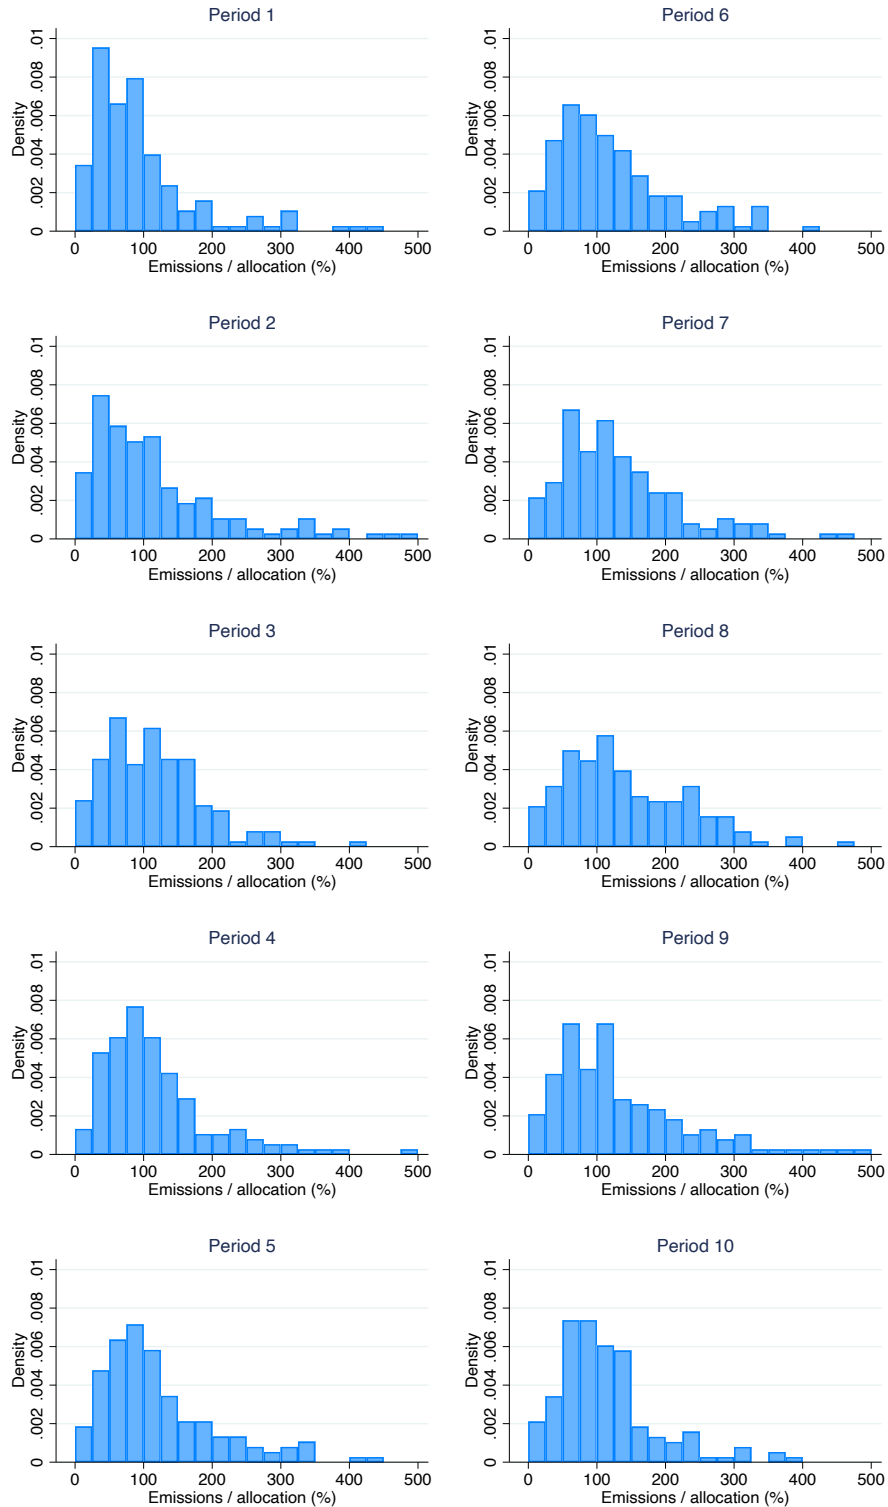
Note. This figure shows the daily permit prices (panel A) and quantities of permits purchased expressed as a fraction of period cap (panel B), from September 2019 to April 2021. Since permits of two consecutive compliance periods were traded simultaneously on some days, we use blue and black colors to differentiate them. On panel A, the points represent the average bid prices of weekly auctions, during which the market-clearing prices were determined. The red, dashed line indicates the price floor of 5 Rs/kg.

Figure 4: Distribution of emissions / total permit holding by compliance period




Note. This figure plots the distributions of [emissions (monthly pro-rated) / total permit holding (monthly pro-rated) * 100%] across treated plants ($N = 156$) by compliance period, truncated at 135% (about 99.5th percentile). Total permit holding is defined as the total number of permits a plant had by the end of a period. Emissions data is directly from NeML inventory data set. Permit holdings is constructed by summing the inventory and consumption from NeML inventory. The bin width is 10. 13 observations with zero permit holdings are dropped.

Figure 5: Distribution of emissions / permit allocation by compliance period




Note. This figure plots the distributions of [emissions (monthly pro-rated) / permit allocation (monthly pro-rated) * 100%] across treated plants ($N = 156$) by compliance period, truncated at 500% (about 97.5th percentile). Emissions and allocation data are directly from the NeML inventory data set. The bin width is 25. One observation with a negative emission value is dropped. Additionally, five observations with zero allocation values are dropped.

Figure 6: PM emissions by treatment status



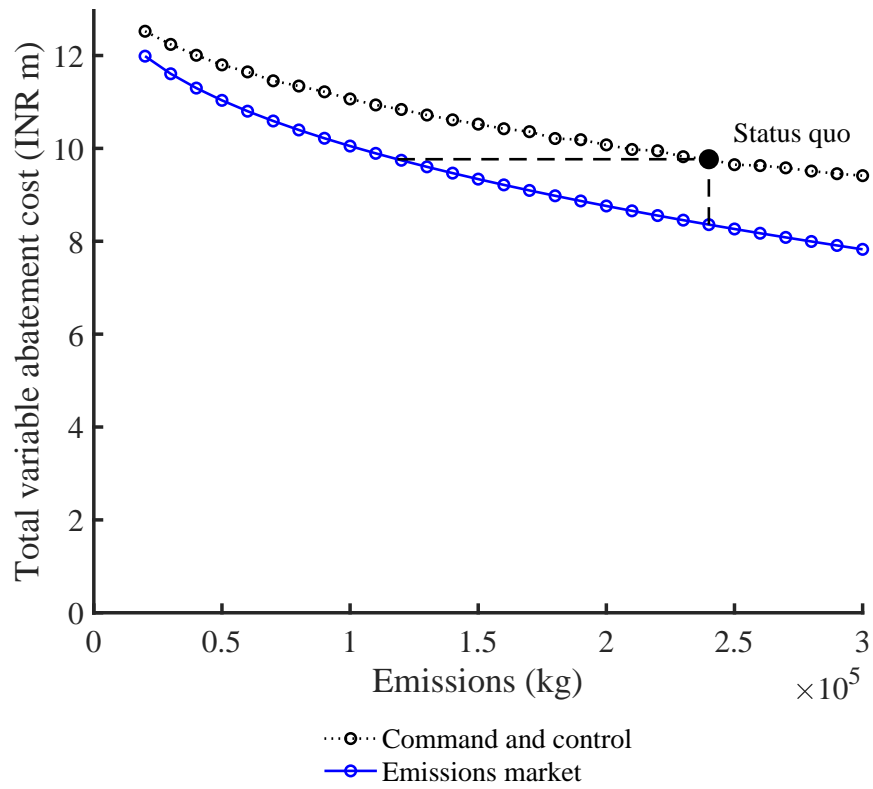
Note. The figure shows the weekly mean per-plant PM emissions in kilograms calculated at a monthly rate equivalent, from April 2019 to March 2021. The missing pollution readings are imputed within a stack-week, and then within a stack-month. Appendix provides a detailed note on the construction of the PM emission variable. This sample consists of 292 plants that had at least one day of PM data from CEMS devices during the ETS experiment. The treatment group is represented by the solid (blue) line, control group by the dashed (grey) line. The grey regions mark the ten compliance periods in the emissions market. The light blue regions mark the two interregnum periods when the emissions market was closed. The aggregate market caps for each compliance period were: 280 tons per 30 days (for Mock-I, Mock-II, and Period-I), 200 tons per 30 days (for Period-II), 180 tons per 30 days (for Period-III), and 170 tons per 30 days thereafter. The horizontal (red) lines denote the per-plant month market cap for each period.

Figure 7: Model fit to market-clearing prices



The figure shows the fit of the model to the time series of market and bid prices by compliance period. The solid (blue) line is the time series of market-clearing prices in the fitted model. The model is fit based on bids in the first half of each compliance period. The dashed (black) line is the time series of mean prices bid in each compliance period. The dotted (black) line is the time series of mean prices for any permit transaction in each compliance period.

Figure 8: Variable abatement costs by regime



The figure shows the total (not marginal) variable abatement costs by regime by regulatory regime. The command and control regime uses a capacity-based emissions rate with error to set emissions targets for each plant, as described in Section 6. The various levels of emissions on the horizontal axis are achieved by scaling up or down the distribution of emissions targets proportionally. The emissions market regime sets an emissions cap at each level of emissions on the horizontal axis. The cost curves for both regimes are drawn for the estimated levels of marginal abatement cost shocks in compliance period four. The dashed horizontal and vertical lines intersect the command and control abatement cost curve at approximately the level of aggregate emissions observed in the control group.

10 Tables

Table 1: Intervention timeline

Compliance Period		Data Collection	
		Survey	CEMS
Dec-2018		Baseline survey	
Apr-2019			CEMS data begins
Jul-2019	Mock-I		
Aug-2019	Mock-II		
Sep-2019	Period-I		
Oct-2019	Period-II		
Nov-2019	Period-III		
Jan-2020	Period-IV		
Feb-2020	Period-V		
Mar-2020	Period-VI		
Apr-2020			
Interregnum (COVID-19)			
Oct-2020	Mock-III		
Nov-2020	Interregnum (Diwali)	Endline survey	
Dec-2020	Period-VII		
Jan-2021	Period-VIII		
Feb-2021	Period-IX		
Mar-2021	Period-X		

Note. Compliance periods were of heterogeneous length, though most lasted approximately one month; of particular note, Period-III began in the middle of November and lasted 45 days until early January. Baseline and endline surveys collected data on plant and boiler house costs, revenue, and emissions abatement mechanisms. While CEMS device readings were collected from April 2019 onward, data availability was low until the emissions trading scheme commenced in July 2019. During mock periods, plants simulated live period transactions with monetary vouchers. We had two interregnum periods where the market was closed: the first wave of the COVID-19 pandemic and shutdowns, and Diwali in 2020. Plant production remained sufficiently high during Diwali in 2019 to continue market operations.

Table 2: Balance of plant characteristics by treatment status

	Treatment	Control	Difference
<i>Panel A: Plant Measures</i>			
Total electricity cost (1,000 USD)	467.6 [869.0]	345.8 [327.0]	121.9 (78.5)
Log(plant total heat output)	15.6 [0.62]	15.6 [0.50]	0.012 (0.065)
Size as recorded on environment consent (1 to 3)	1.37 [0.64]	1.37 [0.62]	0.0052 (0.075)
Small-scale (size=1)	0.72 [0.45]	0.71 [0.46]	0.0063 (0.054)
Large-scale (size=3)	0.086 [0.28]	0.075 [0.26]	0.011 (0.032)
Number of stacks	1.08 [0.41]	1.04 [0.21]	0.035 (0.038)
Textiles sector (=1)	0.85 [0.36]	0.87 [0.33]	-0.025 (0.041)
<i>Panel B: Plant Abatement and Investment Cost</i>			
Boiler house employment	36.9 [32.9]	32.3 [29.4]	4.62 (3.69)
Boiler house capital expenditure (1,000 USD)	199.9 [405.0]	171.4 [196.6]	28.5 (38.3)
Boiler house operating cost (1,000 USD)	140.4 [206.3]	112.4 [84.2]	28.0 (18.3)
APCD: Cyclone present	0.98 [0.14]	0.97 [0.17]	0.0100 (0.019)
APCD: Bag filter present	0.80 [0.40]	0.88 [0.33]	-0.079* (0.043)
APCD: Scrubber present	0.64 [0.48]	0.61 [0.49]	0.030 (0.058)
APCD: ESP present	0.12 [0.33]	0.075 [0.26]	0.045 (0.035)
<i>Panel C: Plant Pollution Measures</i>			
Plant total PM mass rate (kg/hr)	3.62 [4.94]	3.60 [3.82]	0.027 (0.52)
Plant mean PM concentration (mg/Nm ³)	179.0 [156.1]	168.8 [150.2]	10.2 (18.2)
Plant mean Ringelmann score (1 to 5)	1.37 [0.43]	1.35 [0.37]	0.017 (0.047)
Above regulatory standard at ETS baseline (=1)	0.34 [0.47]	0.28 [0.45]	0.054 (0.055)

Note. This table shows differences in plant measures (panel A), plant abatement and investment cost (panel B), and plant pollution (panel C) between the treatment and control groups of plants in the baseline survey conducted from December 2018 to January 2019. This sample consists of 292 plants that had at least one day of PM data from CEMS devices during the ETS experiment (See Table A4 for the balanced table of the full sample). In panel B, cyclone, bag filter, scrubber, and electrostatic precipitator (ESP) are different devices used to reduce emissions. Some plants did not respond to some questions in the survey. For the control group, the numbers of observations are 125 for boiler house capital expenditure, 129 for gross sales revenue, 136 for plant total heat output and number of stacks, and 134 for the rest. For the treatment group, the numbers of observations are 142 for boiler house capital expenditure, 144 for gross sales revenue, 155 for Ringelmann score, 156 for plant total heat output and number of stacks, and 151 for the rest. The first and second columns show means with standard deviations given in brackets. The third column shows the coefficients from regressions of each variable on treatment, with robust standard errors in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 3: Treatment effects on PM emissions (log(PM mass/month))

	No Imputation				With Imputation			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ETS Treatment=1	-0.178** (0.0764)	-0.193** (0.0763)	-0.177** (0.0752)	-0.194** (0.0751)	-0.282*** (0.0744)	-0.282*** (0.0745)	-0.316*** (0.0567)	-0.316*** (0.0568)
Month FE		Yes		Yes		Yes		Yes
Imputation rule					Rule A	Rule A	Rule B	Rule B
Reweighted			Yes	Yes				
Mean dep. var (control)	6.67	6.67	6.66	6.66	6.80	6.80	6.88	6.88
R ²	0.13	0.17	0.14	0.17	0.18	0.22	0.16	0.25
Plants	292	292	292	292	292	292	292	292
Observations	3235	3235	3235	3235	3796	3796	3796	3796

Note. This table reports the estimated treatment effects on PM emissions. The outcome variable is the log of plant-level PM mass (kg) per month. A detailed note on the construction of the outcome variable is in Appendix . To construct this variable, we first truncate the stack-level daily mean PM mass rate (kg/hr) data by setting values greater than the 99th percentile to missing. Missing stack-level daily PM mass rates are imputed with the stack’s own weekly mean PM mass rate. All remaining missing values of a stack’s daily PM mass rate are imputed using the stack’s monthly mean PM mass rate. For the panel imputed with Imputation Rule A: *Stack-Experiment*, all remaining missing values of a stack’s daily PM mass rate are imputed using the stack’s mean PM mass rate across the experiment (July 2019 to March 2021, excluding interregnum). For the panel imputed with Imputation Rule B: *Treatment-Month*, all remaining missing values of a stack’s daily PM mass rate are imputed using the monthly mean PM mass rate of the stack’s treatment group. The stack-level imputed PM mass (kg) is then calculated by taking the product of stack-level daily PM mass rate (kg/hr) and daily non-report hours, and the stack-level daily PM mass (kg) is the sum of the actual PM mass (kg) and the imputed PM mass (kg). Finally, the plant-level monthly PM mass is obtained by summing up the stack-level daily PM mass by month and stacks in the plant. Here, a month is defined as the 16th of this month to the 15th of next month. All columns control for plant characteristics including capital expenditure, operating cost, log(total heat output), mean boiler installation year, and their corresponding indicators for missing values. In addition to plant controls, columns 2, 4, 6, and 8 add month-year fixed effects to control for time variant differences common in each plant. We also apply the inverse probability weighting method in columns 3 and 4. The probability of reporting in a month is predicted using a probit model where the only explanatory variable is an indicator variable that takes value 1 if the CEMS phase of a plant is either 3 or 4, and takes 0 if the CEMS phase is either 1 or 2. Robust standard errors in parentheses are clustered at the plant level with statistical significance indicated by * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 4: Treatment effects on plant costs using survey data

	Total Costs (1)	Components				
		Capital (2)	Labor (3)	Electricity (4)	Fuel (5)	Materials (6)
<i>Panel A: Plant, Outside Boiler House</i>						
ETS Treatment=1	258.0 (233.8)		19.08 (34.17)	25.21* (13.53)		302.4 (250.1)
R ²	0.03		0.02	0.65		0.01
Control mean Plants	1193.24 224		306.37 249	162.13 247		764.17 283
<i>Panel B: Boiler House</i>						
ETS Treatment=1	11.26 (26.31)	-7.178 (19.05)	1.561 (3.332)		26.87* (15.35)	-0.142 (0.596)
R ²	0.93	0.63	0.05		0.98	0.19
Control mean Plants	578.48 185	190.88 218	47.86 262		299.50 225	4.33 283

Note. This table reports the effects of treatment assignment on plant costs outside boiler house (panel A) and boiler house costs (panel B). Specifications use our best estimates for plant costs from the endline survey (FY 2019-20) and control for a relevant baseline survey estimate (FY 2017-18, unless otherwise noted). Both surveys, and especially the baseline, were geared toward costs relevant to emissions; these costs likely account for about half the full firm costs. Survey data is subject to some noise as firms may classify costs differently. Variable notes: At endline, boiler house capital costs include annualized installation and modification costs as well as annual operating and maintenance costs; at baseline, they only include annual operating and maintenance costs. At endline, labor costs outside the boiler house consist of office and production house worker salaries, and labor costs within the boiler house include all boiler house managerial and daily wage workers. Both sets of labor costs control for boiler house daily wage workers at baseline. Electricity costs are only reported at the plant-level at both endline and baseline; consequently, this specification also includes boiler house electricity costs. Fuel costs include costs for all boiler house fuels, most of which are for coal and lignite. At baseline, fuel is of FY 2018-19. Materials consist of production house materials (e.g., cloth), water, and chemical at both baseline and endline; materials are handled almost exclusively outside the boiler and vary substantially due to different products and supplier arrangements (e.g., some firms obtain free water). Robust standard errors are given in parentheses with statistical significance indicated by * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 5: Treatment effects on abatement capital using survey data

	All APCDs (1)	Components			
		Cyclone (2)	Bag (3)	Scrubber (4)	ESP (5)
<i>Panel A: Any</i>					
ETS Treatment=1	0 (.)	0.0233* (0.0134)	0.0650*** (0.0231)	-0.0151 (0.0310)	-0.0311 (0.0207)
R ²	.	0.66	0.68	0.71	0.75
Control mean	1.00	0.95	0.85	0.67	0.12
Plants	276	276	276	276	276
<i>Panel B: Count</i>					
ETS Treatment=1	0.123 (0.0892)	0.135** (0.0574)	0.0733* (0.0440)	-0.0482 (0.0548)	-0.0542 (0.0346)
R ²	0.81	0.77	0.80	0.77	0.79
Control mean	4.86	1.93	1.53	1.23	0.17
Plants	276	276	276	276	276
<i>Panel B: Capital Cost (\$1000s)</i>					
ETS Treatment=1	-3.467 (3.089)	0.602** (0.266)	0.530* (0.318)	-0.222 (0.407)	-4.281 (3.344)
R ²	0.90	0.85	0.83	0.84	0.89
Control mean	44.04	7.80	9.85	9.69	16.70
Plants	276	276	276	276	276

Note. This table reports the effects of treatment assignment on the presence of APCDs (panel A), the count of APCDs (panel B), and the capital cost of APCDs (panel C). All specifications control for the corresponding baseline count or value. In panel C, costs are the product of the number of abatement devices with its average industry-standard cost by boiler house capacity. We use the same cost value at the baseline and the endline. Robust standard errors are given in parentheses with statistical significance indicated by * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 6: Elasticity of marginal cost with respect to emissions

	log(Bid price)			
	(1)	(2)	(3)	(4)
log(Emissions as bid)	-0.0445 (0.0540)	-0.0933 (0.0565)	-0.206*** (0.0730)	-0.491*** (0.0897)
log(Plant total heat output)	0.0395 (0.0367)	0.0953** (0.0432)		
Bid day (normalized)	1.441*** (0.130)	1.644*** (0.116)	1.554*** (0.116)	1.670*** (0.141)
Period FE		Yes	Yes	
Plant FE			Yes	
Plant \times Period FE				Yes
R ²	0.10	0.18	0.36	0.55
Plants	146	146	138	127
Observations	3120	3120	3112	2775

Note. This table reports the results of regressing log(bid price) on log(emissions as bid). Emissions as bid is defined as the permit holdings if the bid is executed. We run regressions using bids placed in the first half of a compliance period. We include compliance period fixed effects in columns 2 and 3, plant fixed effects in column 3, and plant \times period fixed effects in column 4. Bid day is the day of the compliance period on which the bid is placed. It is normalized by dividing the total length of the compliance period and the true-up period. Robust standard errors in parentheses are clustered at the plant level with statistical significance indicated by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Variable abatement costs under alternative regulatory regimes

	Emissions = 170 tons			Emissions = 240 tons		
	Price	Cost	Δ Cost	Price	Cost	Δ Cost
	(INR/kg)	(INR m)	(%)	(INR/kg)	(INR m)	(%)
	(1)	(2)	(3)	(4)	(5)	(6)
Emissions market	12.23	10.08	0	9.91	9.31	0
Constant emissions rate		10.89	8.04		10.24	9.99
Constant emissions rate, with error		11.23	11.41		10.62	14.07
Capacity-based rate		10.91	8.23		10.27	10.31
Capacity-based rate, with error		11.27	11.81		10.67	14.61
Capacity-based rate, correlated error		11.39	13		10.8	16

The table shows the results of counterfactual simulations under different regulatory regimes. Each row represents a different regime. The first row is the emissions market. The second through final rows are different command and control regimes that vary in how the emissions target is set for each plant. Constant emissions rate sets a single fixed ratio of emissions to heat output capacity for all plants. Constant emissions rate with error allows for idiosyncratic variation in the constant rate across plants. Capacity-based rate sets an emissions rate as a function of plant capacity, such that larger plants can have higher or lower rates of emission per unit capacity. Capacity-based rate with error allows for the capacity-based rate to idiosyncratically vary across plants. Finally, capacity-based rate with correlated error is the same as capacity-based rate with error except that the idiosyncratic error is drawn with a negative -0.1 correlation with estimated plant marginal abatement cost shocks. Columns 1 to 3 show results for emissions of 170 tons per month (the treatment level) and columns 4 to 6 for emissions of 240 tons per month (the control level). Within each set of three columns the variables show the market price (if applicable), the total variable abatement costs per month, and the change in abatement costs relative to the emissions market.

A Data Appendix

A.1 Pollution Data From Continuous Emissions Monitoring Systems (CEMS)

Construction of the emissions variable.—We describe how we construct the plant-level monthly average PM mass (kg). CEMS provides stack-level daily reporting hours and uncalibrated daily average PM mass rate (kg/hr) or PM concentration (mg/Nm³). A plant might have multiple stacks. A month in our analysis is defined as the 16th of this month to the 15th of next month. We follow four steps: calibration, truncation, imputation, and aggregation.

Calibration

The raw data set consists of 242,303 daily observations of 337 stacks (318 plants) from April 16th, 2019 to April 3rd, 2021. Stacks are assigned to install either Type-1 or Type-2 CEMS devices. The Type-1 devices measure the daily average PM mass rates (kg/hr), and the Type-2 devices measure the daily average PM concentration (mg/Nm³). The PM mass rate and concentration are calibrated according to the device type. For a stack i (j) that uses Type-1 (2) devices, we calibrate its average PM mass rate (concentration) on the day d using the formula

$$\text{PM_Rate}_{i,d} = m_i \text{PM_Rate}_{i,d}^{\text{raw}} + c_i,$$

$$\text{PM_Conc}_{j,d} = m_j \text{PM_Conc}_{j,d}^{\text{raw}} + c_j,$$

where m and c are stack's calibration factors. Any negative calibrated value is set to missing. We convert the mass rate to concentration, or vice versa, using

$$\text{PM_Conc}_{i,d}^{\text{cal}} = \frac{1000^2 \text{PM_Rate}_{i,d}^{\text{cal}}}{(3600 \text{max_velocity}_i) \times \text{stack_area}_i},$$

where max_velocity is the maximum flue velocity (m/s) of calibration samples, and stack_area is the stack cross-sectional area (m²).

Truncation

A stack-day observation is an outlier if its concentration is greater than the 99th percentile of the calibrated stack-level daily average PM concentration in the month of that day. We set outliers' calibrated PM mass rates and concentrations to missing. Truncation is based on concentration because the concentration is comparable across stacks while the mass rate is not. We drop all observations of a plant if it has no non-missing value for PM mass rate during the ETS experiment. The result is a panel of daily observations of 310 stacks (292 plants) from April 16th, 2019 to April 3rd, 2021 ($N = 222,890$).

Imputation

We impute the stack-level daily average PM mass rate (kg/hr). Let $\text{PM_Rate}_{i,d}^*$ denote the imputed PM mass rate of plant i on day d , and let $\text{Hour}_{i,d}$ denote the reporting hour. If $\text{PM_Rate}_{i,d}^*$ is available for (i, d) , then the *validated* stack-level daily PM mass (kg) is given by

$$\text{PM_Mass}_{i,d}^{\text{val}} = \begin{cases} \text{PM_Rate}_{i,d} \cdot \text{Hour}_{i,d} + \text{PM_Rate}_{i,d}^* \cdot (24 - \text{Hour}_{i,d}) & \text{if } \text{PM_Rate}_{i,d} \text{ is not missing,} \\ \text{PM_Rate}_{i,d}^* \cdot 24 & \text{if } \text{PM_Rate}_{i,d} \text{ is missing.} \end{cases}$$

Otherwise, we will leave $\text{PM_Mass}_{i,d}^{\text{val}}$ as missing.

The first step is imputing daily average PM mass rate with the stack's weekly average PM mass rate. If the weekly average is not available, we use different averages to impute as summarized in Table A1.


Table A1: Summary of imputation rules

Step	Consideration	No Imputation	Imputation Rule A: Stack-Experiment	Imputation Rule B: Treatment-Month
	Imputation Level	Stack daily mean PM mass rate (kg/hr)	Stack daily mean PM mass rate (kg/hr)	Stack daily mean PM mass rate (kg/hr)
1	Truncation	99th percentile	99th percentile	99th percentile
2	Impute for missing values	Stack weekly mean PM mass rate	Stack weekly mean PM mass rate	Stack weekly mean PM mass rate
3	Impute for the rest of missing values	Stack monthly mean PM mass rate	Stack mean PM mass rate across ETS experiment	Treatment group monthly mean PM mass rate

Aggregation.


We first aggregate the validated stack-level daily PM mass to the stack-level monthly PM mass. We set the stack-level monthly PM mass as missing if there is one (or more) missing observations in that month. We then aggregate the stack-level monthly PM mass to the plant-level monthly PM mass. For a plant with multiple stacks, we set the plant-level monthly PM mass missing if one (or more) stack has a missing monthly value. The final product is a panel of monthly observations of 292 plants from April 2019 to March 2021 ($N = 7,008$).

Figure A1: Data availability from CEMS by treatment status



Note. The figure shows the percentage of plants reporting, at weekly frequency, from April 2019 to March 2021. The missing pollution readings are imputed within a stack-week, but not across stacks or weeks. This sample consists of 292 plants that had at least one day of PM data from CEMS devices during the ETS experiment. The treatment group is represented by the solid (blue) line, and control group by the dashed (grey) line. The grey regions mark the ten compliance periods in the emissions market. The light blue regions mark the two interregnum periods when the emissions market was closed.

Figure A2: PM emissions by treatment status



Note. The figure shows the weekly mean per-plant PM emissions in kilograms calculated at a monthly rate equivalent, from April 2019 to March 2021. In the top panel, the missing pollution readings are imputed within stack-week, and then within stack-experiment. In the bottom panel, they are imputed within stack-week, and then within treatment-month. Appendix provides a detailed note on the construction of the PM emission variable. This sample consists of 292 plants that had at least one day of PM data from CEMS devices during the ETS experiment. The treatment group is represented by the solid (blue) line, control group by the dashed (grey) line. The grey regions mark the ten compliance periods in the emissions market. The light blue regions mark the two interregnum periods when the emissions market was closed. The aggregate market caps for each compliance period were: 280 tons per 30 days (for Mock-I, Mock-II, and Period-I), 200 tons per 30 days (for Period-II), 180 tons per 30 days (for Period-III), and 170 tons per 30 days thereafter. The horizontal (red) lines denote the per-plant month market cap for each period.

Table A2: Distribution of number of stacks by plant

Number of Stacks	All	Treatment	Control
1	289	149	140
2	12	5	7
3	1	1	0
4	2	2	0
Total	304	157	147


Note. This table shows the distribution of number of stacks by plant for 304 in-sample plants surveyed at ETS baseline.

Table A3: Mean of the log(PM emissions) by imputation rules

	Control	Treatment	All
No Imputation	6.67 [1336]	6.52 [1899]	6.58 [3235]
Rule A: Stack-Experiment	6.80 [1768]	6.54 [2028]	6.66 [3796]
Rule B: Treatment-Month	6.88 [1768]	6.59 [2028]	6.72 [3796]


Note. The table shows the mean \ln [PM emissions (kg/month)] with number of observations given in the brackets by different imputation rules in the control group, treatment group, and the whole sample.

Figure A3: Kernel density of PM emissions by treatment status



Note. This figure plots the kernel density of PM emissions (kg/month) by treatment status in different stages of imputation described in Table A1. Stack-Week corresponds to the emissions variable after step 2. Stack-Month, Stack-Experiment, and Treatment-Month correspond to the variables constructed based on the No Imputation Rule, the Imputation Rule A, and the Imputation Rule B, respectively. Note that imputing the treatment group mean causes values to converge to the group mean. Since the distribution of emissions is highly positive-skewed, the emissions of most plants are less than the group mean. Rule B, therefore, inflates the emissions of those plants. As a result, the peak of the kernel density curve under Treatment-Month for the control group shifts to the right.

Figure A4: Kernel density of log(PM emissions) by treatment status



Note. This figure plots the kernel density of log[PM emissions (kg/month)] by treatment status in different stages of imputation described in Table A1. Stack-Week corresponds to the emissions variable after step 2. Stack-Month, Stack-Experiment, and Treatment-Month correspond to the variables constructed based on the No Imputation Rule, the Imputation Rule A, and the Imputation Rule B, respectively. Imputing the treatment group mean causes values to converge to the group mean, so the distribution of PM emissions and that of log(PM emissions) should have less dispersion under Rule B. As the distribution of PM emissions is more clustered near the mean under Rule B, the mean of log(PM emissions) should be closer to the log of mean PM emissions for Rule B. By the concavity of log function, the log of mean is no less than the mean of log values. Hence, the mean of log(PM emissions) should be higher for Rule B than others.

A.2 Survey Data

The ETS baseline survey was conducted from December 2018 to February 2019. The unit of analysis is a plant, which has at least one stack. The survey consists of three main sections: a general section, a technical section, and an isokinetic stack sampling section. In the general section, researchers at J-PAL South Asia asked the plant managers questions about plant operations. Researchers then spoke to boiler engineers to collect information about the machinery specifications for the technical section. For the last part, environmental labs collected samples from the stack attached to the boiler and/or thermopack to measure the PM concentration and PM mass rate. Participation in the survey is voluntary. Plants were notified by J-PAL South Asia that their name and data would not be published in any report, and their data would never be shown to the Gujarat Pollution Control Board (GPCB). J-PAL covered the cost of stack sampling and surveys. In addition to stack sampling, J-PAL South Asia had conducted ten rounds of Ringelmann surveys from February 2018 to June 2019. The Ringelmann score is a scale for measuring the apparent density of smoke. The scale has five levels of density. Score 1 to 5 correspond to an opacity of 20%, 40%, 60%, 80% and 100%. Prior to Ringelmann surveys, GPCB informed plants that the information collected would not be used for determining compliance with the GPCB norms or any other legal/regulatory purpose.

In Table 2 and Table A4, variables in panel A are from the general section of the ETS baseline survey, and those in Panel B are from the technical section. In panel B, cyclones, bag filters, scrubbers, and electrostatic precipitators (ESPs) are air pollution control devices (APCDs) used to abate PM emissions. In panel C, the plant's total PM mass rate is the sum of the plant's stacks' PM mass rates measured from stack sampling, and the plant's mean PM concentration is the mean of the plant's stacks' PM concentrations from sampling. The plant's mean Ringelmann score is the average of scores from the four pre-treatment rounds of Ringelmann surveys conducted from April 2019 to June 2019.

Table A4: Balance of plant characteristics by treatment status, full sample

	Treatment	Control	Difference
<i>Panel A: Plant Measures</i>			
Total electricity cost (1,000 USD)	456.2 [853.1]	389.1 [660.7]	67.1 (89.6)
Log(plant total heat output)	15.6 [0.61]	15.5 [0.59]	0.085 (0.067)
Size as recorded on environment consent (1 to 3)	1.36 [0.63]	1.40 [0.65]	-0.038 (0.073)
Small-scale (size=1)	0.72 [0.45]	0.69 [0.47]	0.033 (0.053)
Large-scale (size=3)	0.083 [0.28]	0.088 [0.28]	-0.0056 (0.032)
Number of stacks	1.08 [0.41]	1.05 [0.21]	0.035 (0.037)
Textiles sector (=1)	0.85 [0.36]	0.85 [0.36]	-0.0032 (0.041)
<i>Panel B: Plant Abatement and Investment Cost</i>			
Boiler house employment	36.8 [32.5]	31.7 [30.0]	5.13 (3.59)
Boiler house capital expenditure (1,000 USD)	198.3 [398.6]	164.2 [190.9]	34.0 (36.7)
Boiler house operating cost (1,000 USD)	138.1 [202.6]	111.0 [84.9]	27.1 (17.6)
APCD: Cyclone present	0.98 [0.14]	0.97 [0.16]	0.0081 (0.017)
APCD: Bag filter present	0.80 [0.40]	0.86 [0.35]	-0.055 (0.043)
APCD: Scrubber present	0.64 [0.48]	0.61 [0.49]	0.032 (0.056)
APCD: ESP present	0.11 [0.32]	0.082 [0.27]	0.033 (0.034)
<i>Panel C: Plant Pollution Measures</i>			
Plant total PM mass rate (kg/hr)	3.62 [4.86]	3.51 [3.76]	0.11 (0.50)
Plant mean PM concentration (mg/Nm ³)	177.9 [153.6]	168.5 [151.5]	9.37 (17.5)
Plant mean Ringelmann score (1 to 5)	1.36 [0.42]	1.35 [0.37]	0.0090 (0.045)
Above regulatory standard at ETS baseline (=1)	0.33 [0.47]	0.28 [0.45]	0.052 (0.053)

Number of plants	162	156
------------------	-----	-----

Note. This table shows differences in plant measures (panel A), plant abatement and investment cost (panel B), and plant pollution (panel C) between the treatment and control groups of plants in the baseline survey conducted from December 2018 to January 2019. This sample consists of 318 plants in the ETS experiment. In panel B, cyclone, bag filter, scrubber, and electrostatic precipitator (ESP) are different devices used to reduce emissions. Some plants did not respond to some questions in the survey. For the control group, the numbers of observations are 137 for boiler house capital expenditure, 141 for gross sales revenue, 148 for Ringelmann score, 156 for plant total heat output, and 147 for the rest. For the treatment group, the numbers of observations are 147 for boiler house capital expenditure, 150 for gross sales revenue, 160 for Ringelmann score, 162 for plant total heat output and number of stacks, and 157 for the rest. The first and second columns show means with standard deviations given in brackets. The third column shows the coefficient from regressions of each variable on treatment, with robust standard errors in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table A5: Sample determination and attrition by treatment status

	Control	Treatment	Total
Plants that received treatment assignment	168	174	342
Closed/extinct plants with treatment assignment	10	10	20
Operational-at-baseline plants with treatment assignment	158	164	322
Plants removed from ETS sample by GPCB	2	2	4
In-sample plants	156	162	318
Plants incompletely treated due to closure	7	6	13
Plants completely treated	149	156	305
In-sample plants surveyed at ETS Baseline	147	157	304
In-sample plants manually stack sampled at ETS Baseline	147	157	304
In-sample plants with GPCB administrative data	156	162	318
In-sample plants reporting CEMS data	136	156	292
In-sample plants surveyed at ETS Endline	142	153	295
Treated plants with market trading data	-	155	155

Note. This table reports the sample determination and attrition during the ETS experiment. Of the original ETS-CEMS sample of 373 plants, 342 operational plants received treatment assignment in May 2019 (row 1). Of these 342 plants included in the ETS treatment randomization, 20 plants were extinct or permanently closed (row 2). The permanent shutdown status of these 20 plants has been verified with Ringelmann survey panel data covering the sample from March 2018 to June 2019, as well as regulatory inspection and audit documentation on the GPCB administrative portal. The 342 plants that received treatment assignment, less the 20 plants who received assignment while extinct or shutdown, yield 322 operational plants with treatment assignment at baseline (row 3). Four of these 322 operational-at-baseline plants were officially removed from the ETS sample by GPCB after the treatment assignment (row 4). Three of the removed plants (2 in control, 1 in treatment) are seasonal sugar cooperatives, operational for only four months of the year; the fourth treatment plant is a particle-board producing plant which uses bagasse, rather than coal, as fuel. Of the 318 in-sample plants, 13 are known to have been incompletely treated by the intervention, due to temporary financial closure before or after the treatment assignment was done (row 6). The 304 plants surveyed at baseline are distinct from the 304 plants manually sampled, and are therefore reported separately (rows 8, 9). This paper reports experimental results from the sample of 292 plants reported at least one day of CEMS data from April 16, 2019 to April 3rd, 2021 (row 11). Of the 162 in-sample plants in the treatment group, 153 plants have market trading data (row 13).

Table A6: Engineering estimates of abatement costs under ideal operating efficiency

	Cyclone (1)	Bag Filter (2)	Scrubber (3)	ESP (4)
<i>Total Boiler Capacity = 3 TPH</i>				
Capital costs (Rs/month, amort.)	6953.33	6518.75	10430.00	78225.00
Variable costs (Rs/month)	3000.00	2812.50	4500.00	33750.00
Emission reduction (%)	80.00	99.00	94.00	99.70
Assumed pollution (kg/month)	7879.48	7879.48	7879.48	7879.48
Emission abatement (kg/month)	6303.59	7800.69	7406.71	7855.85
Average abatement cost (Rs/kg)	1.58	1.20	2.02	14.25
Variable abatement cost (Rs/kg)	0.48	0.36	0.61	4.30
<i>Total Boiler Capacity = 6 TPH</i>				
Capital costs (Rs/month, amort.)	9560.83	15645.00	16514.17	104300.00
Variable costs (Rs/month)	4125.00	6750.00	7125.00	45000.00
Emission reduction (%)	80.00	99.00	94.00	99.70
Assumed pollution (kg/month)	11616.85	11616.85	11616.85	11616.85
Emission abatement (kg/month)	9293.48	11500.68	10919.84	11582.00
Average abatement cost (Rs/kg)	1.47	1.95	2.16	12.89
Variable abatement cost (Rs/kg)	0.44	0.59	0.65	3.89
<i>Total Boiler Capacity = 8 TPH</i>				
Capital costs (Rs/month, amort.)	11299.17	19990.83	26075.00	173833.33
Variable costs (Rs/month)	4875.00	8625.00	11250.00	75000.00
Emission reduction (%)	80.00	99.00	94.00	99.70
Assumed pollution (kg/month)	18061.92	18061.92	18061.92	18061.92
Emission abatement (kg/month)	14449.54	17881.30	16978.21	18007.73
Average abatement cost (Rs/kg)	1.12	1.60	2.20	13.82
Variable abatement cost (Rs/kg)	0.34	0.48	0.66	4.16
<i>Total Boiler Capacity = 15 TPH</i>				
Capital costs (Rs/month, amort.)	13906.67	20860.00	26075.00	234675.00
Variable costs (Rs/month)	6000.00	9000.00	11250.00	101250.01
Emission reduction (%)	80.00	99.00	94.00	99.70
Assumed pollution (kg/month)	43907.43	43907.43	43907.43	43907.43
Emission abatement (kg/month)	35125.95	43468.36	41272.98	43775.71
Average abatement cost (Rs/kg)	0.57	0.69	0.90	7.67
Variable abatement cost (Rs/kg)	0.17	0.21	0.27	2.31

Note. Table displays engineering estimates of abatement cost for different APCDs and boiler capacities. We assume no prior operational APCDs and each APCD is purchased in isolation. Costs can be compared with those in other tables at a rate of INR 70 to USD 1. Capital costs are amortized to a monthly flow value. All plants are assumed to have a raw inlet concentration of 2,000 mg/Nm³; in practice it can vary between 1,000 mg/Nm³ and 10,000 mg/Nm³. This is converted to a monthly mass rate via a volumetric flow rate collected at baseline, assuming continuous operation for 16 hours/day and 25 days/month. Of plants with boilers in our analysis sample, the boiler capacity (BC) distribution is: 11% have 2-3 TPH BC, 47% have 4-7 TPH BC, 36% have 8-14 TPH BC, 6% have 15+ TPH BC.

Table A7: Engineering estimates of abatement costs under ideal operating efficiency, if a cyclone is already operating

	Cyclone (1)	Bag Filter (2)	Scrubber (3)	ESP (4)
<i>Total Boiler Capacity = 3 TPH</i>				
Capital costs (Rs/month, amort.)	6953.33	6518.75	10430.00	78225.00
Variable costs (Rs/month)	3000.00	2812.50	4500.00	33750.00
Emission reduction (%)	80.00	99.00	94.00	99.70
Assumed pollution (kg/month)	1575.90	1575.90	1575.90	1575.90
Emission abatement (kg/month)	1260.72	1560.14	1481.34	1571.17
Average abatement cost (Rs/kg)	7.89	5.98	10.08	71.27
Variable abatement cost (Rs/kg)	2.38	1.80	3.04	21.48
<i>Total Boiler Capacity = 6 TPH</i>				
Capital costs (Rs/month, amort.)	9560.83	15645.00	16514.17	104300.00
Variable costs (Rs/month)	4125.00	6750.00	7125.00	45000.00
Emission reduction (%)	80.00	99.00	94.00	99.70
Assumed pollution (kg/month)	2323.37	2323.37	2323.37	2323.37
Emission abatement (kg/month)	1858.70	2300.14	2183.97	2316.40
Average abatement cost (Rs/kg)	7.36	9.74	10.82	64.45
Variable abatement cost (Rs/kg)	2.22	2.93	3.26	19.43
<i>Total Boiler Capacity = 8 TPH</i>				
Capital costs (Rs/month, amort.)	11299.17	19990.83	26075.00	173833.33
Variable costs (Rs/month)	4875.00	8625.00	11250.00	75000.00
Emission reduction (%)	80.00	99.00	94.00	99.70
Assumed pollution (kg/month)	3612.38	3612.38	3612.38	3612.38
Emission abatement (kg/month)	2889.91	3576.26	3395.64	3601.55
Average abatement cost (Rs/kg)	5.60	8.00	10.99	69.09
Variable abatement cost (Rs/kg)	1.69	2.41	3.31	20.82
<i>Total Boiler Capacity = 15 TPH</i>				
Capital costs (Rs/month, amort.)	13906.67	20860.00	26075.00	234675.00
Variable costs (Rs/month)	6000.00	9000.00	11250.00	101250.01
Emission reduction (%)	80.00	99.00	94.00	99.70
Assumed pollution (kg/month)	8781.49	8781.49	8781.49	8781.49
Emission abatement (kg/month)	7025.19	8693.67	8254.60	8755.14
Average abatement cost (Rs/kg)	2.83	3.43	4.52	38.37
Variable abatement cost (Rs/kg)	0.85	1.04	1.36	11.56

Note. The table is the same as Table A6 except one cyclone is already assumed to be operating when calculating the quantity of abatement.

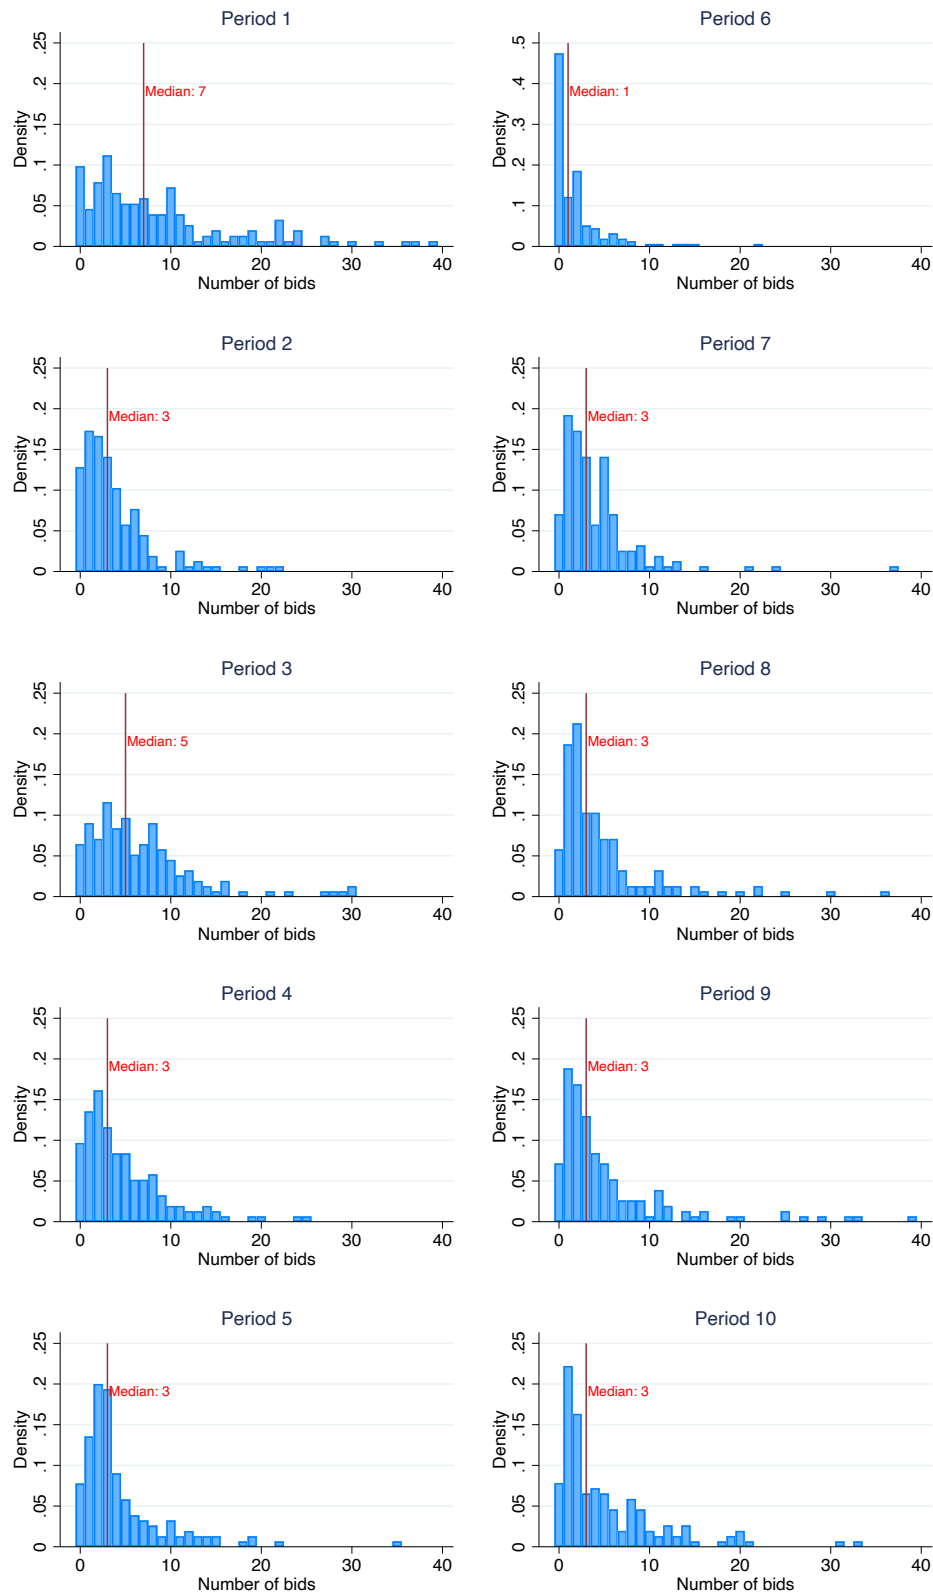
A.3 Trading Data from the Emissions Market

Table A8: Trading data summary statistics

	All	Purchase	Sale
<i>Panel A: Order</i>			
Order quantity (kg)	411.61 (707.98)	429.50 (565.09)	398.78 (794.52)
Order price (Rs/kg)	11.25 (11.56)	9.47 (10.50)	12.52 (12.10)
Order price (Rs/kg), weighted by quantity	9.23 (8.49)	8.42 (8.71)	9.86 (8.27)
Observations	8433	3520	4913
<i>Panel B: Trade</i>			
Trade quantity (kg)	360.23 (563.25)	389.58 (543.76)	326.64 (583.10)
Trade price (Rs/kg)	9.32 (7.38)	9.21 (9.30)	9.45 (4.21)
Trade price (Rs/kg), weighted by quantity	8.44 (6.17)	8.19 (7.26)	8.78 (4.23)
Observations	3799	2027	1772


Note. This table shows the mean of order quantity and price (panel A) and trade quantity and price (panel B), with the standard deviation given in the brackets.

Figure A5: Distribution of number of bids placed per plant by compliance period



Note. This figure presents the distributions of number of bids placed per plant by compliance period, truncated at 40 (about 99th percentile). The bin width is 1. The red line indicates the median number of bids placed.

Figure A6: Elasticity estimate by weeks remaining in the order period



Note. The top panel presents the coefficients of log(emissions as bid) from regressing $\ln(\text{bid price})$ on $\log(\text{emissions as bid})$ and $\text{plant} \times \text{period}$ fixed effects, estimated with different sample truncations defined by the number of weeks remaining in the order period. The bottom panel shows the number of bids placed in different sample truncations.

B Emissions Market Design Appendix

Figure B1: Location of Surat market within India



Note. This figure is a map of India with state outlines. The state of Gujarat is shaded and the city of Surat is marked by a × symbol.

Table B1: Compliance periods and market caps

Period	Start Date	End Date	Days	Cap (kg/30 days)	Per-plant Cap (kg/30 days)	Total Cap (kg)
Mock-I	2019/07/15	2019/08/12	29	280,000	1,728	270,667
Mock-II	2019/08/13	2019/09/15	34	280,000	1,728	317,333
Compliance-I	2019/09/16	2019/10/15	30	280,000	1,728	280,000
Compliance-II	2019/10/16	2019/11/15	31	200,000	1,235	206,667
Compliance-III	2019/11/16	2019/12/31	46	180,000	1,111	276,000
Compliance-IV	2020/01/01	2020/01/31	31	170,000	1,049	175,667
Compliance-V	2020/02/01	2020/02/29	29	170,000	1,049	164,333
Compliance-VI	2020/03/01	2020/03/21	21	170,000	1,049	119,000
Interregnum-I	2020/03/22	2020/10/11	204	-	-	-
Mock-III	2020/10/12	2020/11/11	31	170,000	1,049	175,667
Interregnum-II	2020/11/12	2020/11/30	19	-	-	-
Compliance-VII	2020/12/01	2020/12/31	31	170,000	1,049	175,667
Compliance-VIII	2021/01/01	2021/01/31	31	170,000	1,049	175,667
Compliance-IX	2021/02/01	2021/02/28	28	170,000	1,049	158,667
Compliance-X	2021/03/01	2021/03/31	31	170,000	1,049	175,667

Note. This table reports the start and end date of compliance periods and the market cap of each period. The market cap is the total amount of PM emissions – summed up across all market participants - that is allowed *per month (30 days)* under the Emissions Trading scheme. The total market cap vary across compliance periods, due to duration of the compliance period. Specifically, the total market cap in a compliance period is the market cap $\times 30 /$ (number of days in the compliance period). The per-plant cap is calculated by dividing the market cap by 162, the number of in-sample plants in treatment arm. The market was closed during Interregnum-I due to the COVID-19 pandemic and during Interregnum-II following the Divali festival.