

The Benefits and Costs of Emissions Trading: Experimental Evidence from a New Market for Industrial Particulate Emissions*

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

March 23, 2022

Abstract

Market-based environmental regulations have the potential to abate pollution at a low cost, but are seldom used in developing countries, where pollution levels are the highest. We report the results of an experiment to test the efficacy of a new emissions market, for particulate matter, in the abatement of air pollution from industrial plants in an Indian city. We find that plants randomly assigned to participate in the market reduce pollution emissions by 20% to 30% relative to control plants that remain in the command-and-control status quo regime. This reduction in pollution was achieved without any added costs of plant abatement capital and at permit prices near low *ex ante* estimates of marginal abatement costs. In ongoing work, we will use permit trading activity in the new market to estimate variable abatement costs.

*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 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: Energy Policy Institute, University of Chicago. Ryan (corresponding author): Department of Economics, Yale University (nicholas.ryan@yale.edu).

1 Introduction

Many developing countries today have extremely high levels of air pollution. The World Health Organization sets a standard for fine particles ($PM_{2.5}$) of 10 micrograms per cubic meter. In the United States, 73 out of 3,142 counties, home to 36 million people, exceed that standard. In India, air pollution exceeds the same standard in *all* 687 administrative districts, home to more than 1.3 billion people. If the level of air pollution in India were reduced to the WHO standard, Indian citizens would see life expectancy increase by five years, on average (Energy Policy Institute at Chicago, 2020).

One reason that air pollution may be high is that environmental regulation is costly or ineffective at bringing emissions down. To say that regulation is ineffective does not mean that there is a technical problem in pollution abatement: the experience of the United States has shown that it is possible to increase manufacturing output while dramatically reducing pollution emissions (Shapiro and Walker, 2018). Rather, developing countries may choose not to reduce pollution when environmental regulations, by their structure or uneven enforcement, put a high cost on firms and regulators relative to how much they bring down pollution.

Market-based instruments such as pollution taxes and emissions trading are one way to blunt the trade-off between environmental quality and abatement cost. Theoretically, emissions markets achieve abatement at the lowest possible cost, by allowing plants with low abatement costs to achieve greater reductions in pollution. Despite this promise, market-based instruments are seldom used in developing countries (Blackman, Li and Liu, 2018). A plausible reason is that the functioning of an emissions market depends on reliable emissions monitoring and the transparent enforcement of penalties. Existing environmental regulations may lack both of these prerequisites (Duflo et al., 2013, 2018). Emissions markets trade in a commodity created by the state; if the state cannot credibly ensure the value of that commodity, by enforcing standards, then markets cannot function.

There are therefore two distinct empirical questions regarding the use of market-based instru-

ments for environmental regulation in developing countries. First, can an emissions market improve compliance and reduce pollution, even where command-and-control regulation is enforced imperfectly? Second, how does the adoption of a market affect abatement costs? The first question is often overlooked, since theory typically assumes that compliance will be complete, whatever regulatory regime is adopted.

This paper provides an empirical test of whether a new emissions market can reduce air pollution emissions in a developing-country context. We collaborated with the environmental regulator in Gujarat, one of India's most industrialized states, to design and implement a market for particulate matter air pollution emissions from industrial plants. We believe this market is the first true particulate matter market in the world.¹ To allow for a rigorous evaluation, the market was introduced in a randomized control trial. All eligible sources, totaling 317 plants, in a large industrial city and the surrounding airshed, were first connected to Continuous Emissions Monitoring Systems (CEMS) to measure pollution emissions. Then a treatment group of 162 plants were shifted into the new emissions market. The control group remained in the status quo command-and-control regulatory regime.

Our data come from several sources meant to characterize both the benefits and costs of the emissions market. First, we have a baseline survey of plant characteristics that covers abatement capital and economic variables like employment and sales. At the time of the baseline survey, we also took independent measurements of air pollution in the stack (chimney) of each sample plant. Second, we have administrative data on plant pollution reporting via CEMS. The CEMS readings measure the pollution load from all sample plants at high frequency. Third, we have administrative data on plant participation in the market, including all permit bids and offers and records of cleared trades as well as any regulatory penalties. Fourth, an endline survey to measure abatement costs for all sample plants.

We have three main findings from the analysis, pertaining to compliance, emissions and abate-

¹Chile introduced what was nominally a market for particulates from point sources; however, due to the costs of monitoring, the market was based on boiler capacity rather than measured pollution (Montero, Sanchez and Katz, 2002). We discuss this case in our literature review below.

ment costs.

The first main finding is that compliance with the market mechanism was nearly perfect. Nearly all treated plants held permits exceeding their emissions in all compliance periods. This result was possible for two reasons. First, the regulator imposed penalties for the few observed violations that exceeded the cost of compliance. These penalties established a reputation for enforcement that could sustain the value of permits. Second, the permit market functioned well. The market saw a large volume of permit trade, up to 20% of the monthly market cap on some single *days*. Plants traded often and plant permit holdings at the end of each compliance period differed greatly from initial allocations.

The second main finding is that the emissions market caused a large drop in air pollution. After the introduction of the market, pollution in treatment plants declined by an estimated 20% to 30%, relative to that in control plants, which remained under the status quo command and control regime. The difference in these point estimates depends on the degree of imputation that is used for calculating the emissions for non-reporting sample plants. Emissions reporting in the experiment was incomplete and rates of reporting were higher in treatment than control plants. Higher rates of non-reporting by dirtier control plants make the point estimate for the market reduction in emissions larger when emissions are imputed for these plants. We find point estimates for the effect of the emissions trading treatment on pollution that are large and statistically significant for a variety of plausible imputation rules.

Our experiment spanned a period from September 2019 to April of 2021 that includes India's lockdown of economic activity against the Covid-19 pandemic, beginning in late March of 2020. During the lockdown the plants in our sample ceased operation with the rest of the economy. The market resumed in December 2021 and we observe similar treatment effects on pollution in both the pre- and post-lockdown periods.

Our third finding is that the above reduction in pollution was achieved with no measureable increase in plant abatement capital costs. The plants in our sample are large and have average annual input expenditures of roughly USD 1.8 million. Our point estimate for the effect of the

treatment on air pollution control device capital costs is USD -3.5 thousand (standard error USD 3.1 thousand). This finding agrees with our prior work, which argued that abatement capital is not a constraint on abatement in the command-and-control regime because plants are effectively mandated to install equipment, even if they do not run it (Duflo et al., 2018). Moreover, in the present experiment we observe that permit prices are fairly low, suggesting that marginal costs of abatement are also low.

Our tentative interpretation of these results is that the market achieved significant reductions in pollution at low abatement costs. We also provide evidence that permit trade happened at low prices near *ex ante* estimates of marginal abatement costs. In ongoing analysis we are developing a model to calculate the social surplus from the adoption of the market, including both the benefits of reducing pollution and the costs of pollution abatement. The model is necessary to supplement the experiment because the experiment changed both (i) regulatory stringency (average pollution emissions) (ii) the allocation of abatement. A finding that the costs of abatement are small is consistent with increases in allocative efficiency offsetting the higher costs due to added stringency. The purpose of the structural analysis is to decompose these effects.

This paper contributes to two distinct literatures in environmental economics and development. The first literature pertains to the efficacy of market-based instruments. The literature has focused in particular 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). There is a broad consensus among economists that emissions markets have achieved abatement at lower cost than would have been possible through command-and-control regulations (Ellerman et al., 2000; Burtraw et al., 2005; Fowlie, Holland and Mansur, 2012). The empirical basis for this consensus is broad, but not especially well-founded, due to a fundamental problem: the markets under study often regulate practically all plants of a given type in a given area. For example, the Acid Rain program regulated large power plants in the Eastern United States. This makes it difficult to develop a counterfactual for plant emissions

or costs in the absence of trade that is based on empirical observation of other plants.² Most evaluations of emissions markets use engineering estimates of costs (Burtraw et al., 2005). 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). The counterfactual problem has loomed over the evaluation of markets for greenhouse gas (GHG) emissions in particular.³ GHG emissions are dependent on aggregate shocks and GHG abatement is inelastic in the short term, making it hard to know if a market achieved any abatement relative to emissions in the counterfactual economy without a GHG market (Ellerman and Buchner, 2008; Borenstein et al., 2019; Martin, Muûls and Wagner, 2020).

Our contribution, against this backdrop, is to provide estimates of the benefits and costs of emissions trading against a sharply defined experimental counterfactual. We find that emissions trading indeed does lower pollution relative to a clearly-defined command-and-control counterfactual that represents the status quo regulatory regime for the same plants. The pollution reductions achieved by any market will depend on the stringency of the cap it imposes. Our estimates of the reduction in emissions in the Surat particulates market are similar to the best prior estimates for emissions reductions due to RECLAIM (Fowlie, Holland and Mansur, 2012).

This paper also contributes to the literature at the intersection of environmental economics and development. A main theme of this literature has been the Herculean difficulty of enforcing environmental regulations that bind on pollution in developing countries (Greenstone and Hanna,

²For 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)

³Martin, 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.”

2014; Blackman, Li and Liu, 2018). Common findings are that: (a) poor or corrupted monitoring impedes regulation (Duflo et al., 2013; Oliva, 2015; Duflo et al., 2018; Zou, 2021) (b) the coarse regulations adopted in response to poor monitoring, in turn, are partly undercut through behavioral responses (Davis, 2008; He, Wang and Zhang, 2020). These forces together may create large differences between the private cost of an abatement action, such as driving a newer car or running an air pollution control device, and the social cost of enforcing a regulation so that this abatement action is chosen. 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) describe a market in Santiago, Chile that targeted particulate matter emissions indirectly. The market functioned poorly and achieved little abatement because it traded in boiler capacity, a proxy for emissions, rather than in directly in emissions load.⁴ The results of our paper 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. This accords with other recent work on the social value of improvements in pollution monitoring (Greenstone et al., 2022).

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 the main empirical results on emissions and abatement costs. Section 6 concludes.

2 Context and experimental design

This section introduces the context of the experiment and describes the experimental design.

⁴Boilers burning solid fuel were the main source of particulates in Chile. Because particulates could not be accurately and economically measured at the time, the Chilean market traded in boiler capacity instead of particulate emissions. The problem this creates is 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). The market disbanded after many covered sources switched fuels in response to a fall in natural gas prices.

2.1 Institutional context

India's system of environmental regulation is a traditional command-and-control system. The Water Act (1974) established state environmental regulators, called State Pollution Control Boards (SPCBs), and gave them the power to enforce standards for water pollution. The Air Act (1981) extended these powers to cover air pollution. These laws provide a criminal framework in which polluters can be prosecuted for exceeding emissions standards. However, because of the high burden of proof and the difficulty in bringing such cases, this framework is seldom used. The main teeth of environmental regulation come from State Pollution Control Boards (SPCBs). SPCBs follow a traditional command-and-control model. The "command" is a mandate that industrial plants must install equipment to reduce pollution emissions. The "control" is that plants can then be closed if they are found not to run this equipment to control pollution.

Some of our own research, joint with Esther Duflo, has tested the efficacy of reforms within this command-and-control framework in the state of Gujarat. Duflo et al. (2013) experimentally tested a reform that changed the incentive structure in the market for third-party environmental auditors, to make auditors independent of the firms on which they report. We found that this reform increased the accuracy of pollution reports and that plants that were assigned to receive more independent reports reduced pollution. The Gujarat Pollution Control Board adopted this reform permanently. Duflo et al. (2018) experimentally studied an increase in inspection frequency from the regulator's own staff. This experiment found a marginal increase in compliance with pollution standards. The effect of inspections on pollution would have been greater if inspections had been targeted at plants that faced a higher risk of sanction.

The interventions studied in this work reduced pollution emissions at the margin, but also showed the limits of the command-and-control regime. In particular, the above studies highlight that while sanctions for violators can be large, they are very infrequently applied. Therefore, despite a broad mandate for the installation of abatement equipment, this equipment is often not run. In trying to target the plants most likely to run their equipment, the regulator can do better than randomly assigning inspections, but still has relatively weak information with which to predict

pollution emissions. All of these points argue that there may be advantages to an emissions market in which pollution is directly measured and firms are incentivized to reduce pollution by a steady, predictable price.

2.2 Market design

The intervention studied is an emissions trading market, also known as cap-and-trade, for particulate matter air pollution. Emissions markets have been used to control pollutants such as sulfur dioxide in the United States and carbon dioxide in the European Union. Our research team collaborated with the regulator and the market operator to develop the market design for the Gujarat particulates market. This part reviews the basic market rules that comprise the experimental treatment.

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 cap in the market was set by GPCB to target a decrease of 29% in pollution relative to the baseline level of pollution. We forecast at the time that this reduction would achieve

the existing regulatory standard for pollution emissions concentrations. 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 market 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. 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.

2.3 Sample and experimental design

The experiment introduced the emissions market in phases for industrial plants in the city of Surat, Gujarat and surrounding industrial areas.

In the first phase, Continuous Emissions Monitoring Systems (CEMS) were installed in all sample plants. The mandate for CEMS reporting installation itself was constructed as an experiment. In the second phase, a randomly selected group of plants were assigned to the emissions market treatment. The present section describes the common sample that was used for both experiments and the experimental design for the second, emissions trading experiment. The emissions trading experiment was cross-randomized with the phases of the CEMS experiment, and begun only after all plants in the CEMS experiment had completed CEMS installation.

The sample of industrial plants was selected to include the plants with the highest air pollution potential in and around the city of Surat, Gujarat. Surat is a city of over 4 million people and

is known as a prosperous industrial 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). This group of 342 plants was randomly assigned to the treatment arm with a probability of one half and the control arm otherwise. 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 in the market is therefore 317 plants, of which 304 were covered in our baseline survey.

Figure 1 shows the contribution of particulate matter emissions from sample plants to ambient particulate matter concentrations in Surat, Gujarat in June 2019. The 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). 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 $\mu\text{g}/\text{m}^3$) 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.⁵

There are two main points to be drawn from the map. First, the contribution of sample plants to urban air pollution is very 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 suspended particulates (of all sizes) in and near the city center, ranges from $60 \mu\text{g}/\text{m}^3$

⁵The model used is based on the eddy diffusion theory, considering each plant to be a stationary point emitting source, with CEMS data informing each plants emission rate (mass/time) and the location and stack height of each plant determining the point source. The model uses simple assumptions on meterological conditions such as constant emission rates and wind speeds.

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 the timeline of the experiment and data collection. The market began with two compliance periods of “mock” trading, beginning July 16th, 2019, in which no money was at stake, but plants were allocated permits and could buy permits with endowments of fake money. The purpose of this period was to inculcate plants in the market rules and to improve the coverage of CEMS monitoring of pollution. After the mock trading period, 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 a nationwide lockdown in India, in response to the Covid-19 pandemic, that stopped most industrial activity in the country. The market was restarted in December 2020. This paper analyzes a data set that includes four additional compliance periods after the restart. The duration of treatment is therefore roughly one year of market operation spread over one and a half years of calendar time.

3 Data and summary statistics

This section describes our data sources and tests for the balance of plant characteristics at baseline by treatment arm.

3.1 Data sources

We use three main sources of data. The first is a survey of plant characteristics, abatement equipment, and economic variables such as sales and capital. The survey was conducted prior to market launch and repeated roughly one year after the market started (Table 1). The second is

high-frequency pollution data from Continuous Emissions Monitoring Systems (CEMS). The third is trading data from the market operator.

The baseline survey was conducted from December 21st, 2018 to January 29th, 2019 in person at sample plants. The survey has both general economic parts and technical parts. The general part of the survey was administered to the plant owner or manager as a respondent. This part covered plant characteristics such as inputs, outputs, sales and energy consumption. The technical part of the survey 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 survey team recorded the characteristics of all emissions sources and all abatement equipment attached to those sources and interviewed plant staff about the costs of equipment operations. At the time of the baseline survey, in addition, we hired independent environmental labs to take manual samples of air pollution emissions from the stack of each factory. 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. CEMS, generically, refers to any *in situ* device for reporting on pollution at high frequency. As part of the development of this project, a member of our research team (Sudarshan) participated in a technical review process with the Central Pollution Control Board (CPCB) to establish standards for particulate matter CEMS for industrial use in India (Central Pollution Control Board, 2013). The Gujarat Pollution Control Board (GPCB) thereafter mandated that all sample plants install CEMS devices that met this standard. CEMS devices are calibrated by comparing CEMS readings to physical pollution samples taken in the same stack at the same time. 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 then reported continuously over the internet to a central server.

The main limitation in the CEMS data is that reporting during the experiment is incomplete.

The mean rate of weekly data reporting began at roughly one-third of plants, before the market started, but rose to 85% of plants by the end of the sample. The data handling system was designed to store data locally during transient internet outages. However, longer outages, device malfunctions and the like leave gaps in the high frequency data. In addition, once the treatment assignments were announced, treatment plants had a much stronger incentive to remedy non-reporting than control plants, because their validated emissions in the market would increase if they did not report data continually (see Section 2). Consistent with this incentive, 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 to a good extent 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 was operated by NCDEX e-Markets Limited (NeML), a branch of the National Commodity and Derivatives Exchange (NCDEX), a private Indian company founded in 2003. The market operator recorded a ledger of permit purchases and sales for all compliance periods as well as a complete order book of bids and offers for permits regardless of whether a transaction took place. 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 in the description of the market and the estimation of marginal costs.

3.2 Balance of baseline covariates by treatment arm

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

Sample plants are large factories. While 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. Energy and related inputs comprise a large share of plant expenditures. The average control plant spends USD 350 thousand on electricity (panel A). The boiler, the main

source of pollution in the plant, costs USD 108 thousand to run each year, not including direct expenditures on fuel (panel B).

Panel C shows that nearly all plants in both the treatment and control groups have abatement equipment for air pollution installed at the baseline. For example, with respect to air pollution abatement equipment (panel C), 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.

Table 2, panel C shows several measures of baseline pollution emissions. PM concentration is the mean particulate matter concentration from manual pollution samples taken during our baseline survey. The average concentration of SPM in stack gas is $169 \text{ mg}/\text{Nm}^3$ in the control group and $179 \text{ mg}/\text{Nm}^3$ in the treatment group. Both of these *average* levels of emissions exceed the SPM *maximum* standard of $150 \text{ mg}/\text{Nm}^3$. The flow rate of stack gas is balanced across the two treatment arms. In addition to physical measurements of pollution, 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. 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 the distributions of pollution in the control and treatment plants at baseline. Panel A compares the distribution of pollution concentrations as measured by stack samples of pollution concentration (truncated at the 95th percentile). Panel B shows the distribution of Ringel-

mann scores, a proxy for pollution based on visual grading of stack emissions. The distributions of pollution by either measure are very similar across treatment arms at baseline. About 30% of the plants in both treatment arms exceed the particulate matter emissions standard (Table 2, panel C).

4 The functioning of the emissions market in the treatment

This section describes trade in the market that the treatment founded. We establish that the market functioned well in that (i) it hosted a large volume of trade (ii) at low prices (iii) with a high rate of compliance.

4.1 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, in the range from INR 5 per kg (the price floor) to INR 16 per kg, depending on the compliance period and week. The market-clearing mechanism (Section 2.2) 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.

Figure 3, panel B shows the permit quantities traded each day 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. The volumes of trade are typically larger during the first part of a compliance period as plants purchase or sell permits to align permit holdings with expected emissions. The volume of trade diminishes towards the end of the period when plants have less uncertainty about what their total emissions for the period will be.

4.2 Permit allocations and plant emissions

The volume of permit trade helped plants to achieve nearly perfect compliance. Compliance is defined as permit holdings at the end of the true-up period exceeding emissions during the compliance period. 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 from the figure. First, compliance is nearly perfect. Plants in bins to the right of 100% are non-compliant in that 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 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 this exact compliance determined by the allocation of permits? Permits were allocated to plants *pro rata* on the basis of a plant's total heat output, a measure of emissions capacity. If plants all 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 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%.

The main finding from Figure 5 is 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 not at all constrained by their initial permit allocations but 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 main result of the paper for pollution. The figure 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.

Mean emissions loads in the treatment and control plants, prior to the experiment, hover around 2250 to 2500 kilograms per month in the three months prior to the start of the market. 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. This difference is sustained throughout the rest of the sample. The difference between treatment and control average emissions load is similar during both the early and late sets of compliance periods.

We investigate the effect of non-reporting on emissions on the difference between treatment and control pollution loads. 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). We find that these imputations generally increase the estimated difference between treatment and control emissions. The reason is that control plants with high imputed emissions are especially likely not to report.

Returning to Figure 6, we find that emissions met the level of the cap in all compliance periods. During each compliance period, we put a horizontal line on the graph showing the mean per plant emissions required to meet the cap exactly. In the figure, 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 emissions, for the purpose of market compliance, were imputed at a rate higher than the mean. With those more punitive imputations the cap more closely binds.

Regression analysis.—To estimate the impact of the emissions trading treatment on pollution, we now turn to a regression analysis of the pollution data. 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. We 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). Standard errors are clustered at the plant level.

Table 3 reports the results. The columns from 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 columns 5 to 8 we use alternate imputation rules for pollution emissions. 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 main finding of the table is that the treatment markedly reduced pollution emissions, by 20% to 30% depending on the preferred estimate. In column 2, the estimated 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). Specifications that allow imputations across plant-months generally lead to higher estimates of treatment effects (columns 5 to 8). 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. Imputations tend to replace missing control group observations for log particulate emissions load with values higher than the mean among control plants that reported. Therefore accounting for under-reporting by control plants tends to increase the size of the estimated treatment effect on pollution emissions.

We conclude that the 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

This part estimates the costs at which the emissions market reduced pollution using data from plant surveys.

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 will 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. Our baseline and endline surveys covered plant costs both for general inputs, like labor, and for abatement-specific inputs, like air pollution control devices.

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).⁶ 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

⁶This 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.

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. Because plant costs overall include many factors unrelated to pollution, they will necessarily be noisy measures of abatement effort. Our survey was therefore designed specifically to measure investment in abatement equipment. The different panels of Table 5 show treatment effects for different dependent variables: 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 three findings from Table 5. 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 mean). Plants may have multiple devices when they have multiple pieces of fuel-consuming equipment.⁷ Plants on average have 1.9 cyclones and 1.5 bag filters (panel B, control mean). Second, we estimate small increases in abatement capital for air pollution control devices designed to control particulate matter (panel C). In particular, we estimate that cyclone capital increased by USD 602 (standard error USD 266, column 2) in the treatment on a base of USD 7,801 and bag filter capital by USD 530 (standard error USD 318, column 3) on a base of USD 9,846. These treatment effects correspond to roughly one in ten plants in the treatment installing a new piece of equipment for particulate abatement (panel B). Third, overall we cannot reject that there was no change in the cost of abatement capital for treated plants (panel C, column 1). While the added capital costs for cyclones and bag filters are statistically significant, they are small, and

⁷Most 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.

offset in total costs by 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 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 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 were sometimes as high as INR 45 per kg.

We use data on abatement costs from Indian air pollution control device vendors to provide engineering estimates of abatement costs as a benchmark for permit prices. Appendix Table A6 presents estimates for abatement costs under *ideal* operating conditions from the operation of the four kinds of air pollution control devices for 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 run-

ning 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 range suggested by engineering estimated, if not at the absolute low end of abatement costs under ideal operation.

6 Conclusion

High regulatory costs are one reason why air pollution emissions may be high in developing countries. We design, implement and evaluate the effects of emissions trading on particulate matter air pollution from industrial point sources. 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.

In ongoing work we will analyze how the introduction of the market affected plant abatement costs and abatement efficiency. A structural analysis of the market operation will complement the results from the randomized experiment, for two reasons. First, it is difficult directly to measure marginal costs of abatement from survey data. Data on permit trades, in the context of a model of abatement, may be used to discipline estimates of marginal abatement cost functions and thereby to calculate the contribution of variable costs to total abatement costs. Second, our experiment changed two aspects of regulation simultaneously: stringency and the means of allocating pollution load across plants. A model would be useful to quantify the separate effects of these two changes

on plant costs.

References

- 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*.
- 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.** 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.
- Ellerman, A Denny, and Barbara K Buchner.** 2008. “Over-allocation or abatement? A preliminary analysis of the EU ETS based on the 2005–06 emissions data.” *Environmental and Resource Economics*, 41(2): 267–287.
- 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.
- Energy Policy Institute at Chicago.** 2020. “The Air Quality Life Index.” Energy Policy Institute at the University of Chicago, <https://aqli.epic.uchicago.edu/the-index/>.
- Fowlie, 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 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, 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.
- He, Guojun, Shaoda Wang, and Bing Zhang.** 2020. “Watering down environmental regulation in China.” *The Quarterly Journal of Economics*, 135(4): 2135–2185.
- 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.

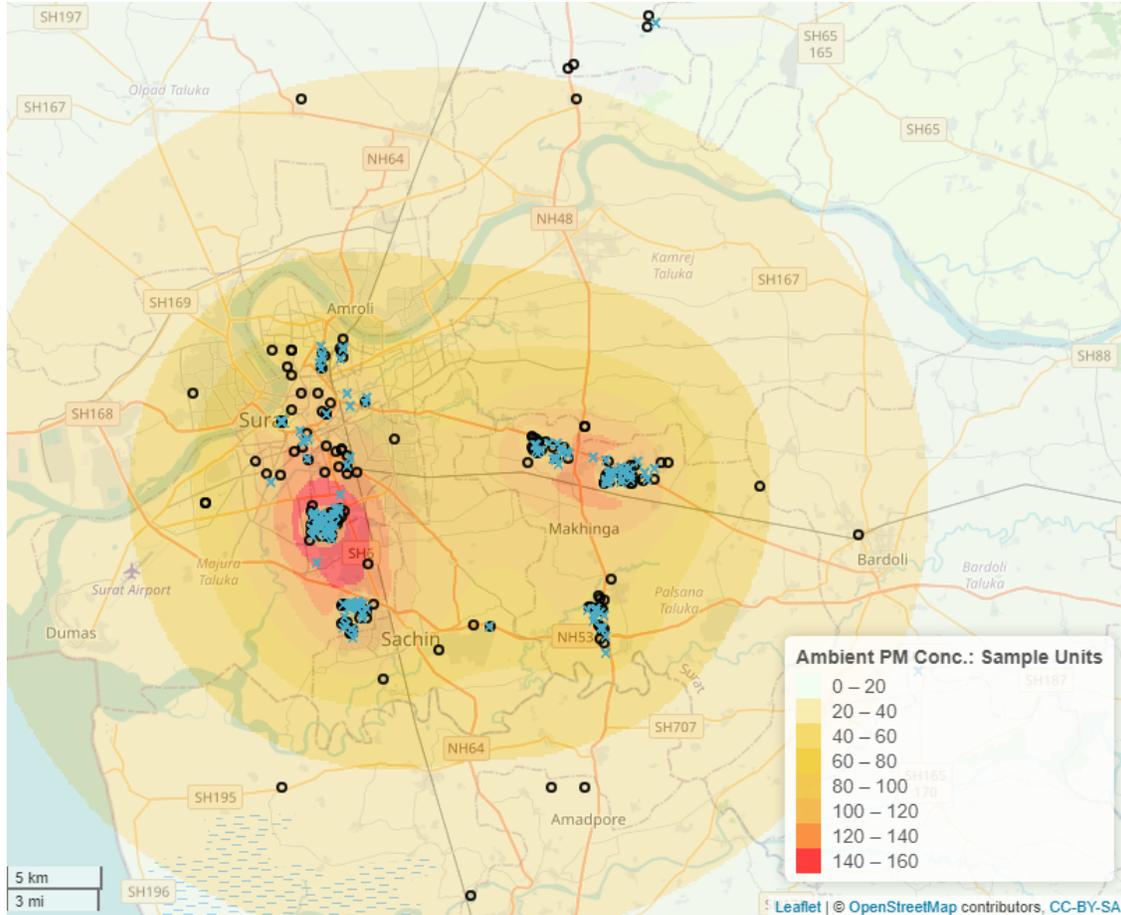
Oliva, Paulina. 2015. “Environmental regulations and corruption: Automobile emissions in Mexico City.” *Journal of Political Economy*, 123(3): 686–724.

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.

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

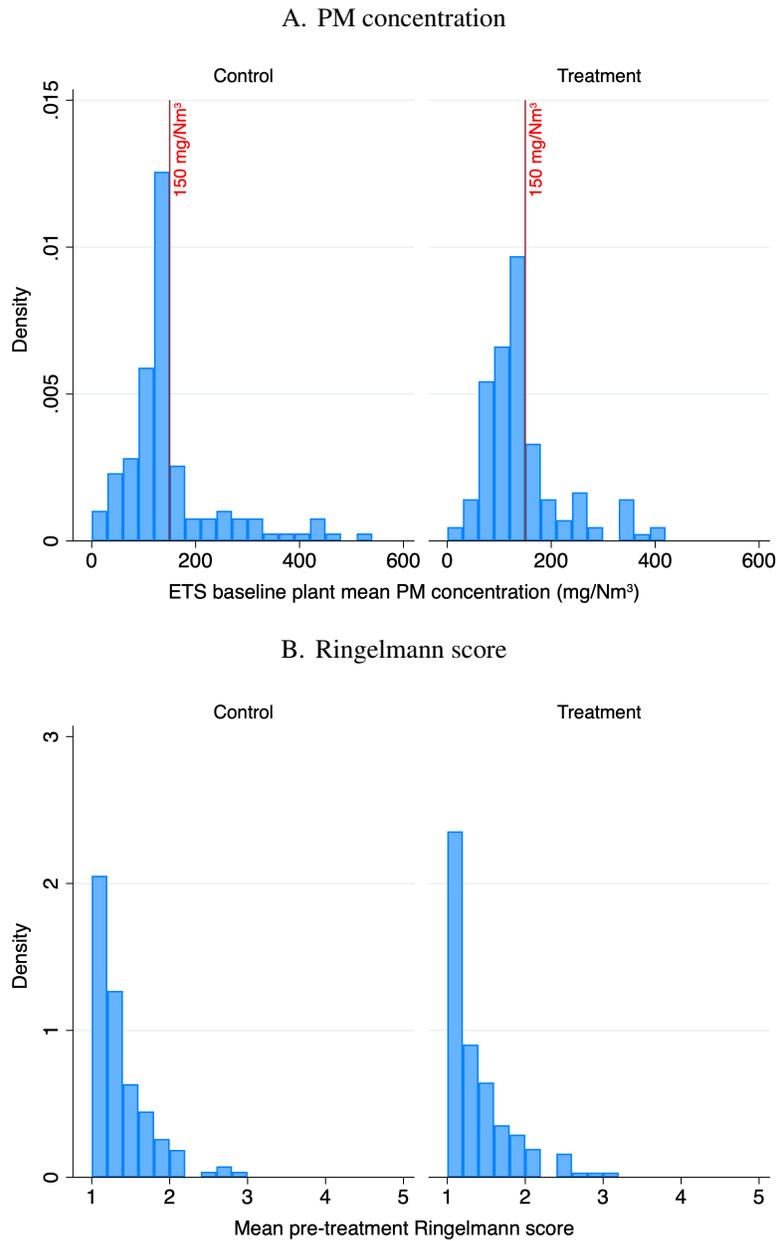
7 Figures

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 × markers and control plants by black o 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 plants 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.

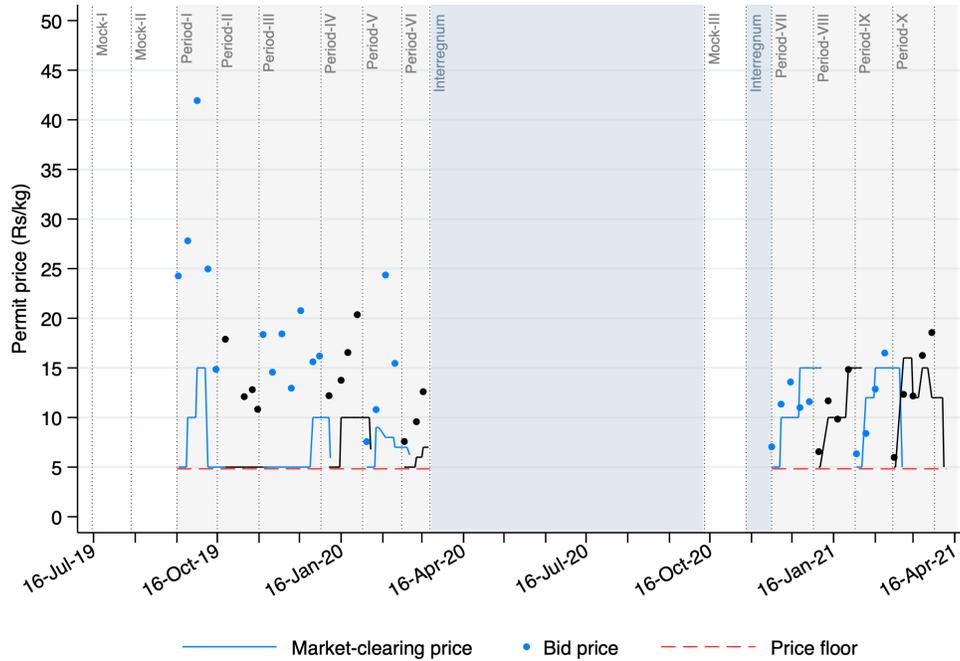
Figure 2: Distribution of pollution before the experiment



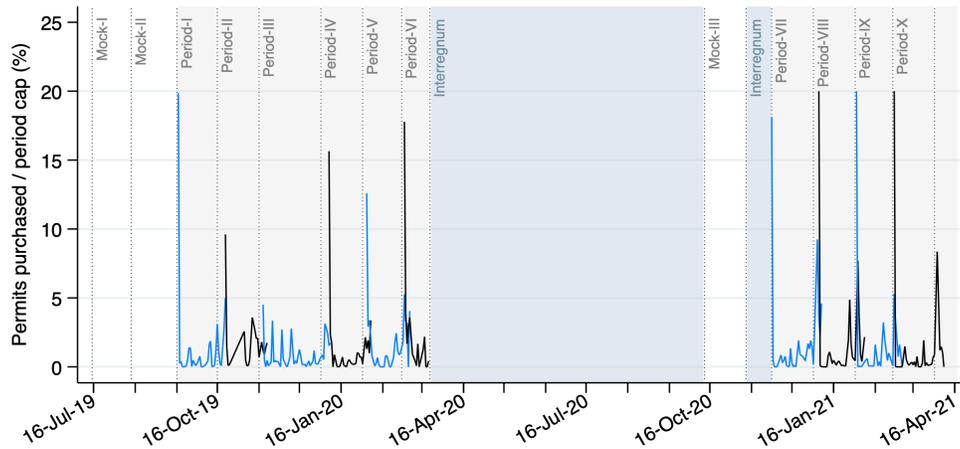
Note. The top panel presents the distributions of the mean PM concentration by treatment status as measured by manual iso-kinetic stack sampling at ETS baseline (December 2018 to January 2019). One PM sample was collected from each industrial stack by a third-party laboratory. We truncated the PM concentration 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

A. Permit prices

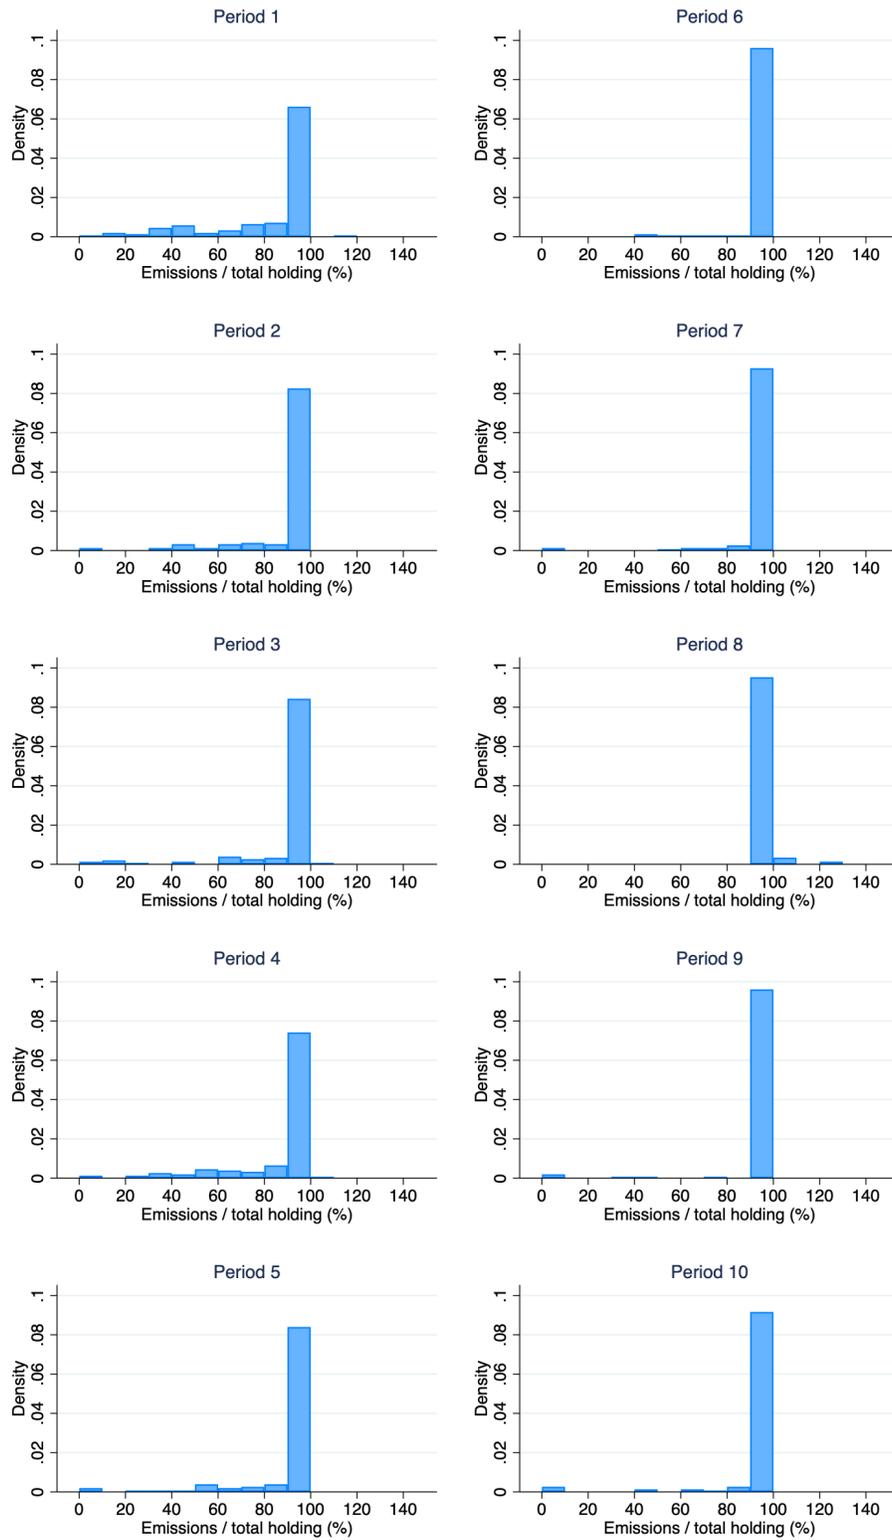


B. Permit quantities



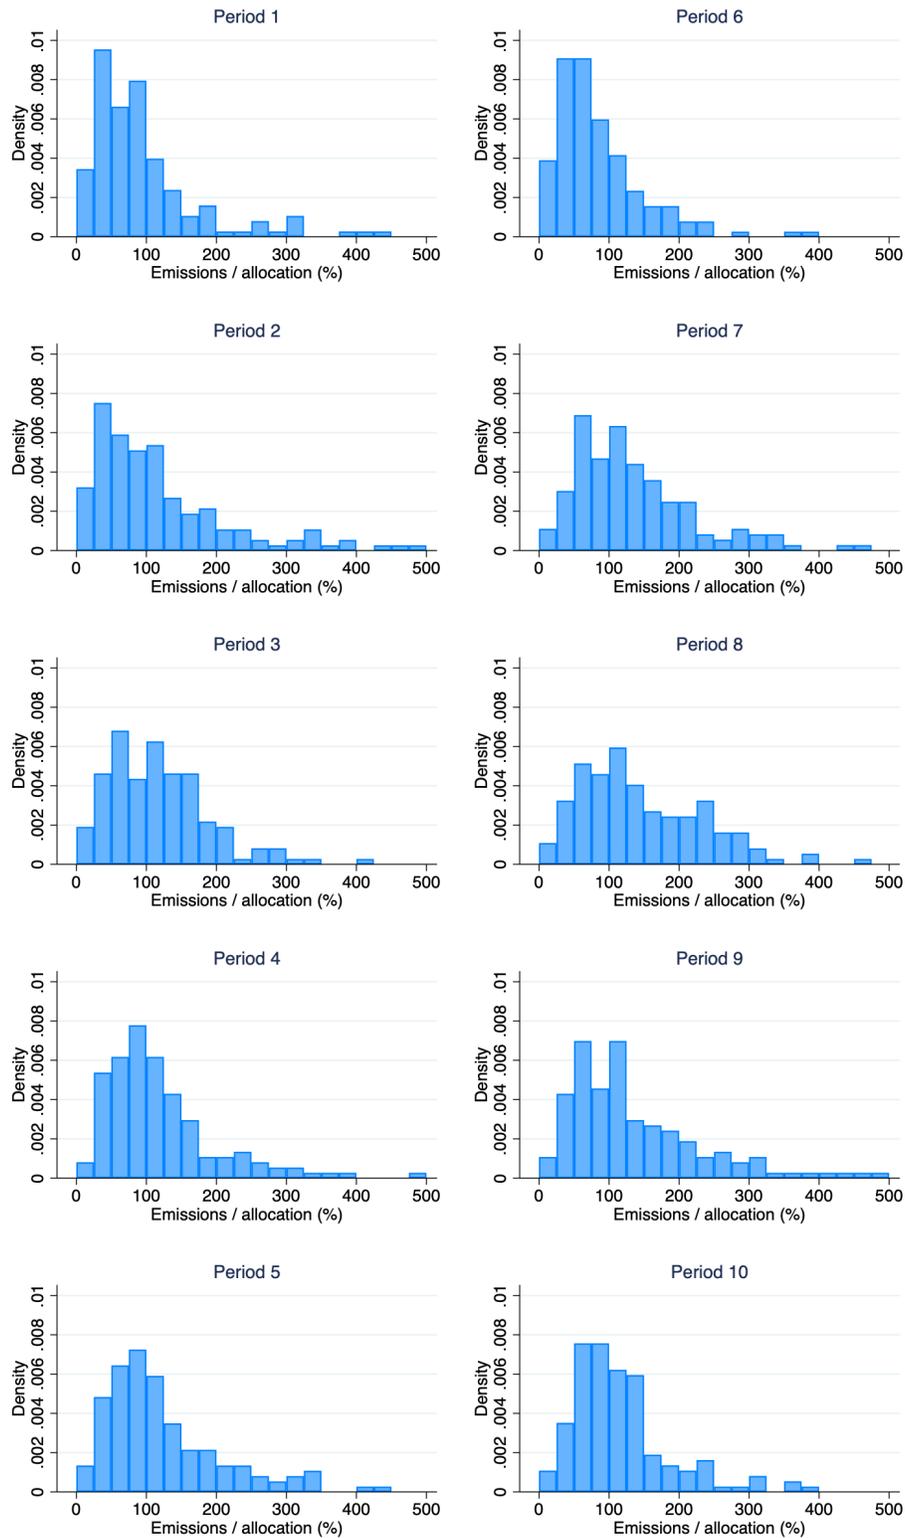
Note. This figure shows the daily permit prices (panel A) and quantities of permits purchased expressed as a fraction of period cap (panel B), from September 2019 to April 2021. Since permits of two consecutive compliance periods were traded simultaneously on some days, we use blue and black colors to 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



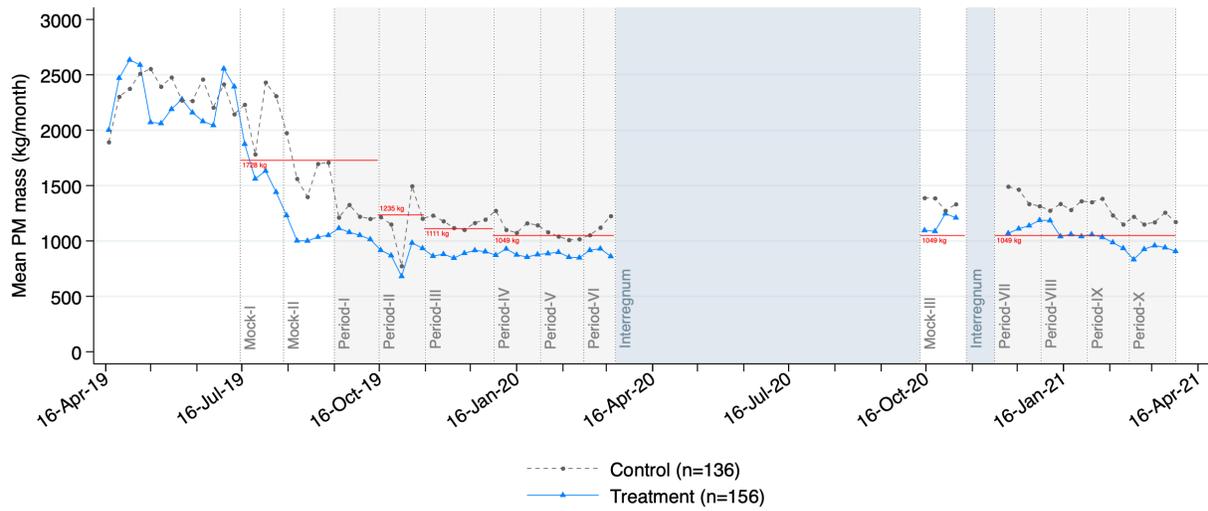
Note. This figure plots the distributions of [emissions (monthly pro-rated) / total permit holding (monthly pro-rated) * 100%] across treated plants ($N = 156$) by compliance period, truncated at 135% (about 99.5th percentile). Total permit holding is defined as the total number of permits a plant had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

8 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 (<i>COVID-19</i>)		
Oct-2020	Mock-III		
Nov-2020	Interregnum (<i>Diwali</i>)	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

	Control	Treatment	Difference
<i>Panel A: Plant Measures</i>			
Total electricity cost (1,000 USD)	345.75 [327.02]	467.64 [869.04]	121.89 (78.50)
Ln(plant total heat output)	15.59 [0.50]	15.60 [0.62]	0.01 (0.07)
Size as recorded on environment consent (1 to 3)	1.37 [0.62]	1.37 [0.64]	0.01 (0.07)
Small-scale (size=1)	0.71 [0.46]	0.72 [0.45]	0.01 (0.05)
Large-scale (size=3)	0.07 [0.26]	0.09 [0.28]	0.01 (0.03)
Number of stacks	1.04 [0.21]	1.08 [0.41]	0.03 (0.04)
Textiles sector (=1)	0.87 [0.33]	0.85 [0.36]	-0.03 (0.04)
<i>Panel B: Plant Abatement and Investment Cost</i>			
Boiler house employment	32.30 [29.37]	36.92 [32.89]	4.62 (3.69)
Boiler house capital expenditure (1,000 USD)	171.41 [196.56]	199.91 [405.01]	28.50 (38.27)
Boiler house operating cost (1,000 USD)	112.40 [84.16]	140.38 [206.26]	27.99 (18.29)
APCD: Cyclone present	0.97 [0.17]	0.98 [0.14]	0.01 (0.02)
APCD: Bag filter present	0.88 [0.33]	0.80 [0.40]	-0.08*** (0.04)
APCD: Scrubber present	0.61 [0.49]	0.64 [0.48]	0.03 (0.06)
APCD: ESP present	0.07 [0.26]	0.12 [0.33]	0.04 (0.03)
<i>Panel C: Plant Pollution Measures</i>			
Plant total PM mass rate (kg/hr)	3.60 [3.82]	3.62 [4.94]	0.03 (0.52)
Plant mean PM concentration (mg/Nm ³)	168.81 [150.21]	179.01 [156.07]	10.21 (18.16)
Plant mean Ringelmann score (1 to 5)	1.35 [0.37]	1.37 [0.43]	0.02 (0.05)
Above regulatory standard at ETS baseline (=1)	0.28 [0.45]	0.34 [0.47]	0.05 (0.05)
Number of plants	136	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 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, scrubber, bag filter, 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	0.133	0.175	0.135	0.173	0.181	0.216	0.163	0.255
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	0.032		0.020	0.647		0.007
Control mean	1193.238		306.366	162.132		764.165
Plants	224		249	247		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	0.925	0.627	0.045		0.979	0.190
Control mean	578.475	190.880	47.860		299.496	4.332
Plants	185	218	262		225	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	.	0.664	0.685	0.711	0.755
Control mean	1.000	0.947	0.848	0.674	0.121
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	0.811	0.772	0.800	0.774	0.794
Control mean	4.864	1.932	1.530	1.227	0.174
Plants	276	276	276	276	276
<i>Panel C: 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	0.903	0.849	0.827	0.844	0.887
Control mean	44.040	7.801	9.846	9.689	16.705
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 * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

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.

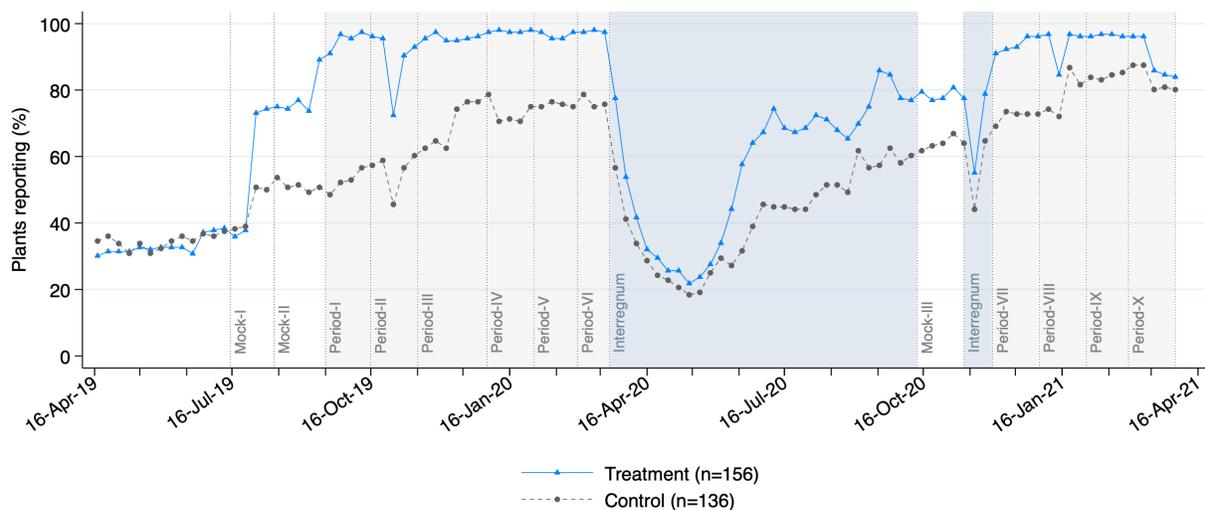
Table A1: Summary of Imputation Rules

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.

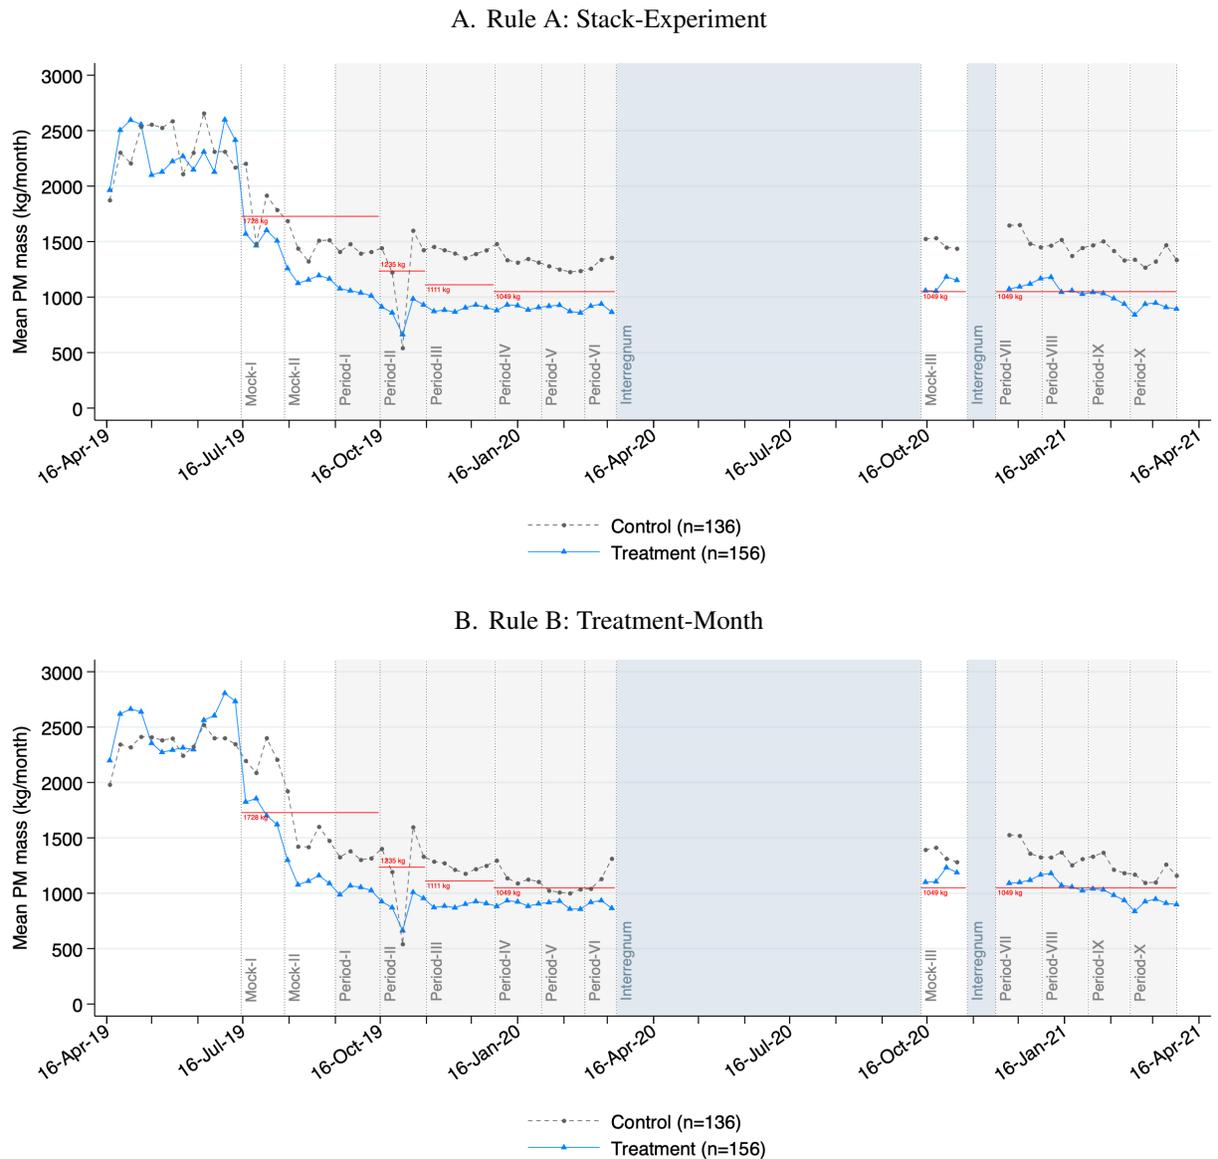
We first aggregate the validated stack-level daily PM mass to the stack-level monthly PM mass. We set the stack-level monthly PM mass as missing if there is one (or more) missing observations in that month. We then aggregate the stack-level monthly PM mass to the plant-level monthly PM mass. For a plant 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

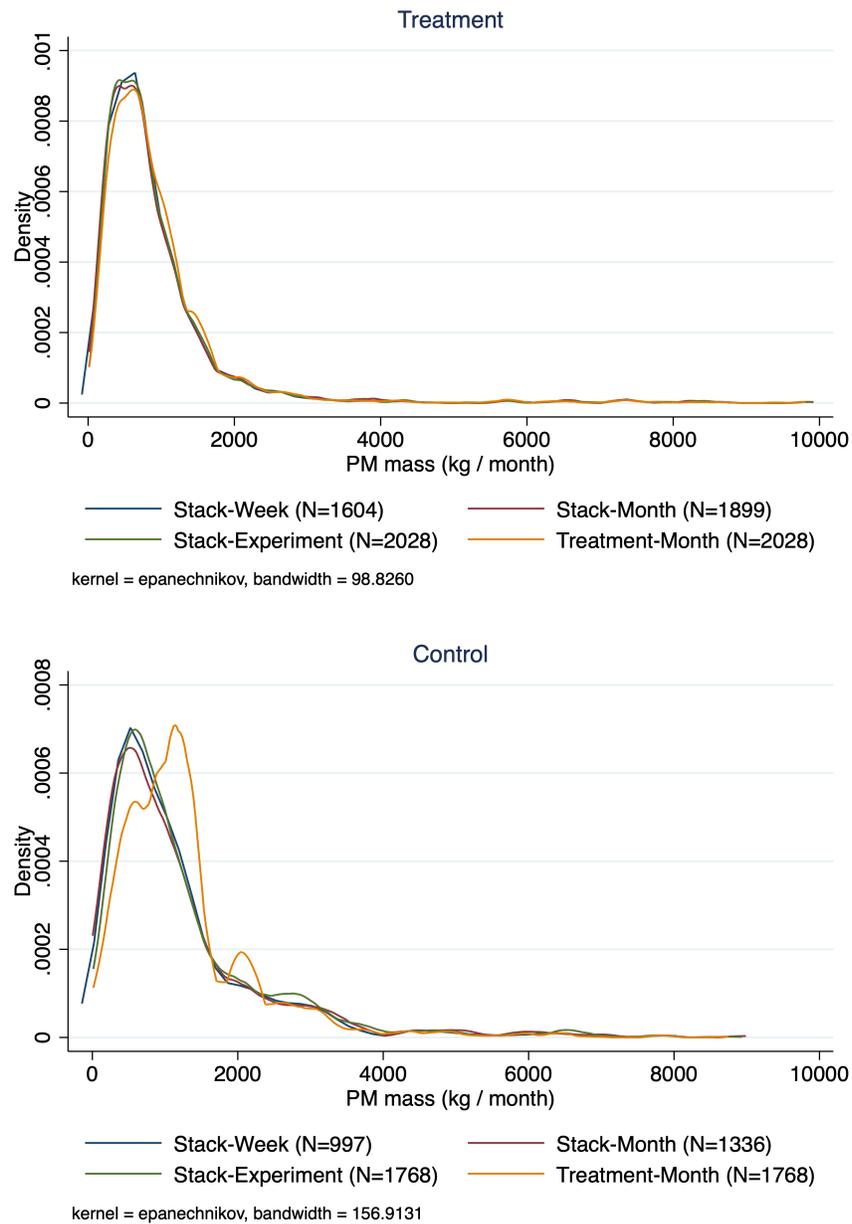
Note. This table shows the distribution of number of stacks by plant for 304 in-sample plants surveyed at ETS baseline.

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

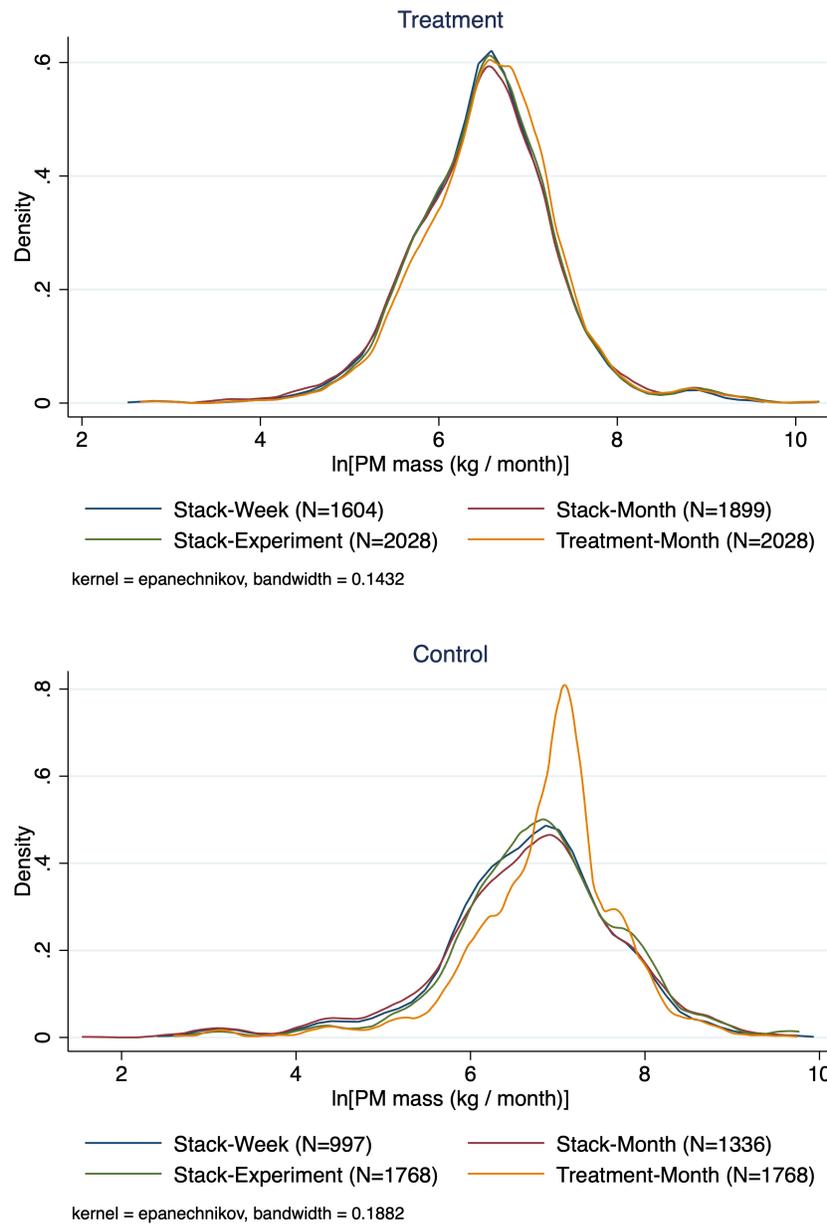
	Control	Treatment	All
No Imputation	6.67 [1336]	6.52 [1899]	6.58 [3235]
Rule A: Stack-Experiment	6.80 [1768]	6.54 [2028]	6.66 [3796]
Rule B: Treatment-Month	6.88 [1768]	6.59 [2028]	6.72 [3796]

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 $\ln(\text{PM emissions})$ by treatment status

Note. This figure plots the kernel density of $\ln[\text{PM emissions (kg/month)}]$ 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 $\ln(\text{PM emissions})$ 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 $\ln(\text{PM emissions})$ 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 $\ln(\text{PM emissions})$ 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, scrubbers, bag filters, 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

	Control	Treatment	Difference
<i>Panel A: Plant Measures</i>			
Total electricity cost (1,000 USD)	389.11 [660.68]	456.23 [853.08]	67.12 (89.64)
Ln(plant total heat output)	15.50 [0.59]	15.59 [0.61]	0.08 (0.07)
Size as recorded on environment consent (1 to 3)	1.40 [0.65]	1.36 [0.63]	-0.04 (0.07)
Small-scale (size=1)	0.69 [0.47]	0.72 [0.45]	0.03 (0.05)
Large-scale (size=3)	0.09 [0.28]	0.08 [0.28]	-0.01 (0.03)
Number of stacks	1.05 [0.21]	1.08 [0.41]	0.04 (0.04)
Textiles sector (=1)	0.85 [0.36]	0.85 [0.36]	-0.00 (0.04)
<i>Panel B: Plant Abatement and Investment Cost</i>			
Boiler house employment	31.71 [30.02]	36.85 [32.51]	5.13 (3.59)
Boiler house capital expenditure (1,000 USD)	164.24 [190.88]	198.27 [398.62]	34.02 (36.70)
Boiler house operating cost (1,000 USD)	111.00 [84.86]	138.13 [202.59]	27.14 (17.62)
APCD: Cyclone present	0.97 [0.16]	0.98 [0.14]	0.01 (0.02)
APCD: Bag filter present	0.86 [0.35]	0.80 [0.40]	-0.05 (0.04)
APCD: Scrubber present	0.61 [0.49]	0.64 [0.48]	0.03 (0.06)
APCD: ESP present	0.08 [0.27]	0.11 [0.32]	0.03 (0.03)
<i>Panel C: Plant Pollution Measures</i>			
Plant total PM mass rate (kg/hr)	3.51 [3.76]	3.62 [4.86]	0.11 (0.50)
Plant mean PM concentration (mg/Nm ³)	168.54 [151.48]	177.92 [153.55]	9.37 (17.50)
Plant mean Ringelmann score (1 to 5)	1.35 [0.37]	1.36 [0.42]	0.01 (0.05)
Above regulatory standard at ETS baseline (=1)	0.28 [0.45]	0.33 [0.47]	0.05 (0.05)
Number of plants	156	162	

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, scrubber, bag filter, 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 EST 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	-	153	153

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 nancial 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.

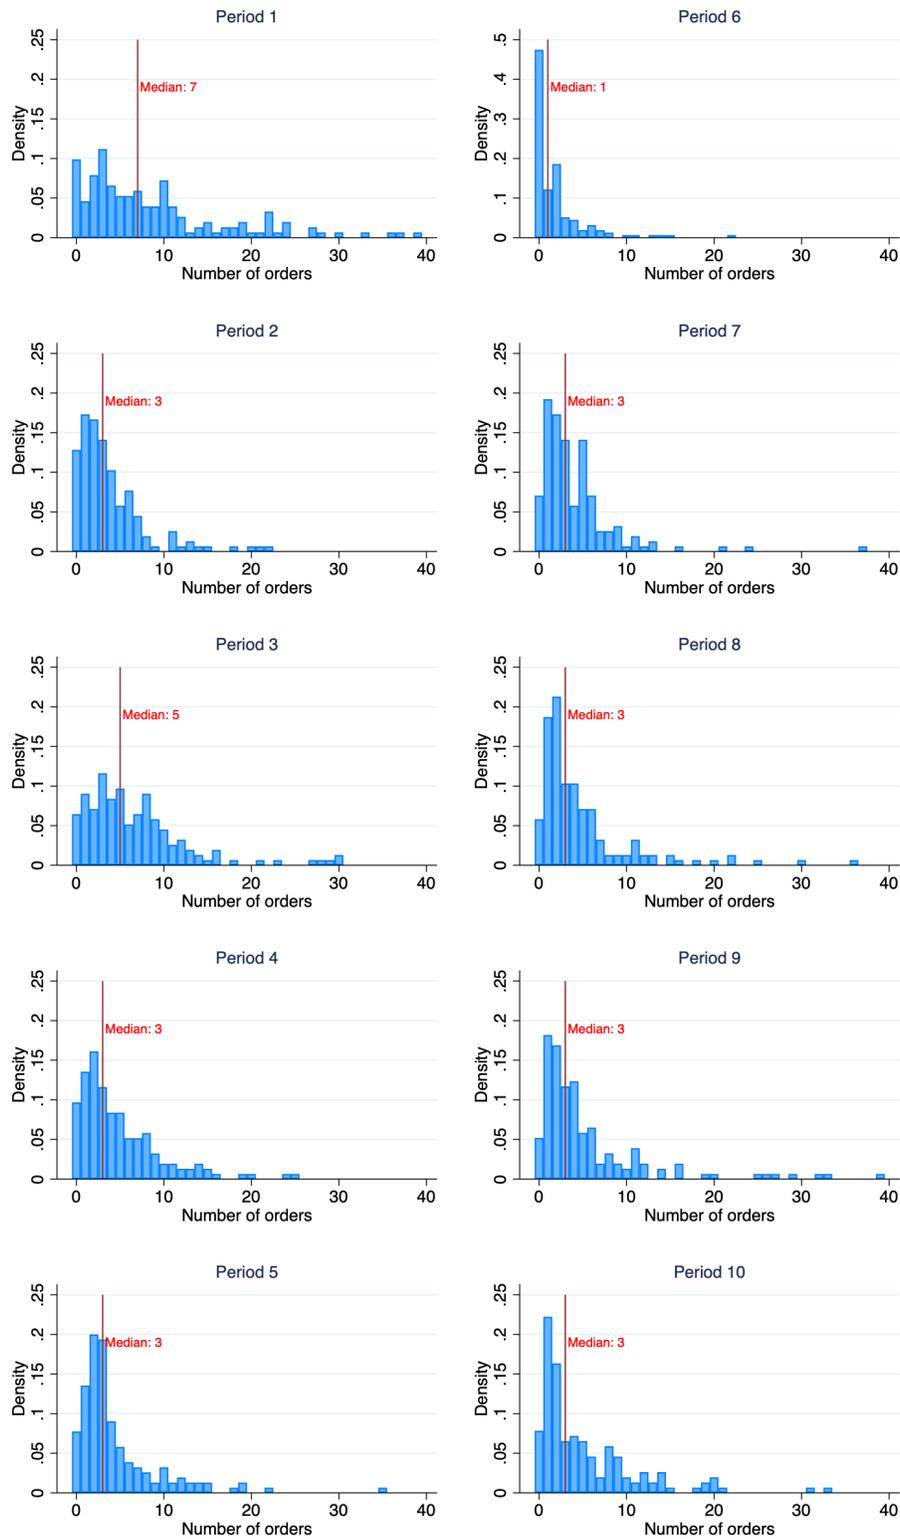
A.3 Trading Data from the Emissions Market

Table A8: Trading data summary statistics

	All	Purchase	Sale
<i>Panel A: Order</i>			
Order quantity	411.61 (707.98)	429.50 (565.09)	398.78 (794.52)
Order price	11.25 (11.56)	9.47 (10.50)	12.52 (12.10)
Order price (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	360.23 (563.25)	389.58 (543.76)	326.64 (583.10)
Trade price	9.32 (7.38)	9.21 (9.30)	9.45 (4.21)
Trade price (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 orders placed per plant by compliance period



Note. This figure presents the distributions of number of orders 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 orders placed.

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.