The Impact of Emissions Trading on Industrial Air Pollution Emissions: Experimental Evidence from a New Pollution Market*

[Preliminary and Incomplete]

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

December 4, 2020

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 24% relative to control plants that remain in the command-and-control status quo regime. This preliminary draft estimates the effect of the market only on pollution. Ongoing work will also measure its effect on abatement costs.

*We thank the Gujarat Pollution Control Board for collaborating in this research, particularly Sanjeev Kumar, Rajiv Kumar Gupta, 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, Sanjana Gorti, Winston Hovekamp, Jared Stolove, and Aakash Bhalothia for extraordinary research assistance.

[†]Energy Policy Institute and Department of Economics, University of Chicago.

[‡]Department of Economics, Yale University.

[§]Energy Policy Institute, University of Chicago.

[¶]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 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 may be a promising 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, emissions markets appear mainly to be used in regulation in developed countries. A plausible reason is that the functioning of an emissions market depends on reliable emissions monitoring and the reliable, 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 instruments 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, we are presently in the field with an endline survey, designed to measure abatement costs for all sample plants. The present draft of this paper shows preliminary results from the analysis of the experiment that relies on only the first two data sources, that is, the baseline survey and administrative data on air pollution.

[1]Chile introduce what was nominally a market for particulates from point sources; however, due to the costs of monitoring, the market as constructed was based upon boiler capacity rather than measured pollution. Abatement was therefore arguably inelastic since it is difficult 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.

The main result of the analysis is that the emissions market caused a large drop in air pollution. After the introduction of the market, pollution in treatment plants declined by 24% (standard error 8%) relative to that in control plants, which remained under the status quo command and control regime. The decline in emissions was sustained during the nine-month life of the market, until operations ceased due to Covid-19 lockdown, which stopped industrial production in our sample almost entirely.

The main caveat to this result on pollution is that emissions reporting in the experiment was incomplete and rates of reporting were higher in treatment than control plants. The differential rate of reporting is itself a result of the market, since treatment plants were incentivized to maintain data reporting by data validation rules that assumed emissions were high if a plant did not report. We therefore consider the robustness of the treatment effect on pollution to alternative assumptions on emissions when a plant-week pollution observation is missing. We find that the point estimate for the effect of the emissions trading treatment on pollution is stable across a range of plausible imputation rules, though our estimates are much less precise when no missing pollution data is imputed or when the sample of plants is restricted to reduce the degree of differential reporting.

This analysis of pollution is part of a larger project to measure both the benefits and the costs of adopting an emissions market. Future revisions of this paper will aim to calculate the social surplus from the adoption of the market, including both the benefits of reducing pollution and the costs of pollution abatement. Given the analysis to date, it is ambiguous whether total abatement costs increased or decreased in the market: even if the market is more efficient in allocating abatement, it may cost more than status quo regulation in total because it achieved lower pollution levels. We are presently running a plant survey to collect data on abatement costs and enable a broader analysis of the effect of the market on pollution and costs together.

The remainder of the paper goes as follows. Section 2 introduces environmental regulation in India and the experimental design. Section 3 discusses the data and the balance of baseline covariates by treatment arm. Section 4 presents the main empirical result and Section 5 briefly concludes.

2 Context and experimental design

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

a 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 recent 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.

b 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. This part reviews the basic market rules that comprise the experimental treatment.

Plants. The regulated entities are 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 particulate matter emissions. Emissions are monitored by Continuous Emissions Monitoring Systems (CEMS) on the stacks of plants.

Compliance. The compliance period is the period over which permits for emissions are valid. 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.

Cap. The cap in the market was set by GPCB to achieve 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 mg/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. The cap was revised downwards in later compliance periods to 170 tons per month.

Permit allocation. Most permits, 80%, were allocated to plants for free. 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, NCDEX e-Markets Limited (NeML), a leading Indian commodity exchange. 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 high-frequency movements 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 around INR 20 to 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. 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, at the rate of twice the ceiling price for every unit of emissions in excess of their permit holdings.

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.

c Sample and experimental design

The experiment introduced the emissions market in phases for industrial plants in the city of Surat, Gujarat, and in 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 A1). 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 $\mu g/m^3$) are measured by passing CEMS emissions rates for the sample plants into a simplified Gaussian dispersion model, comparable to the SCREEN3 model used by the US Environmental Protection Agency.[2]

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 $\mu g/m^3$ to 180 $\mu g/m^3$ across space.[3] 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.

The treatment arm plants participated in the emissions market as described above. 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

[2]The 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.

[3]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 1.8 to 5.4.

India, in response to the Covid-19 pandemic, that stopped most industrial activity in the country. The duration of treatment is therefore roughly nine months inclusive of the mock trading period or seven months without. The control arm plants continued on the status quo command-and-control environmental regulation for air pollution.

3 Data and summary statistics

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

a Data sources

We use two primary sources of data. The first is a baseline survey of plant characteristics, abatement equipment, and economic variables such as sales and capital. The second is high-frequency pollution data from Continuous Emissions Monitoring Systems (CEMS).

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 A3. The main rule we use is to impute missing data for any plant-week at the mean level of pollution in the treatment group in that week. We expect this rule will tend to understate differences in emissions, to the extent that the treatment did in fact reduce emissions. We will test the robustness of our results to alternative rules that either do not impute missing data across plant-weeks at all, or that impute missing data at the 75th percentile of the treatment emissions load distribution. The latter assumption, to impute at a high rate, is supported by an ancillary finding, not shown, that plants that began reporting pollution later on in the experiment tended to have higher levels of pollution. This suggests that the pollution of plants that do not report should be expected to be higher than that of plants that do.

b Balance of baseline covariates by treatment arm

Table 1 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. The average control plant generates sales of USD 13 million per year (panel A). While many plants are formally classified as "small scale" (69%) this is a government classification, based on the 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 380 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, 60% of control (64% of treatment) plants have a scrubber installed, 86% of control (81% of treatment) plants have a bag filter 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 1, panel D 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 175 mg/Nm^3 in the control group and 181 mg/Nm^3 in the control group. Both of these *average* levels of emissions exceed the SPM *maximum* standard of 150 mg/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.

Finally, because the emissions trading experiment is part of a larger phase-in design of monitoring, it is possible to consider balance in our sample further back in time as well. Figure 2 shows the distribution of pollution by treatment arm, using data from a survey wave in late 2014, well prior to our experiment. The mean SPM readings in the two treatment arms are very close, 353 mg/Nm^3 in the control versus 366 mg/Nm^3 in the treatment, and even higher than those observed in our baseline survey. There is a considerable right tail of readings at concentrations three times the regulatory standard or beyond.

4 Empirical results

This section considers the average effect of the emissions trading treatment on plant pollution emissions. In ongoing work, we plan to extend our analysis to consider heterogeneity across plants in pollution as well as the effect of trading on abatement costs.

a Graphical analysis

Figure 3 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 15 April 2019 to

15 March 2020, by treatment arm. Panel A shows the mean imputation data series where missing plant-weeks are imputed at the treatment group mean for that week, in both the treatment and control arms. Panel B shows the raw data series where missing pollution readings are imputed within a plant-week, but not across plants or weeks. In both panels, treatment firms are represented by the solid (blue) line, control firms by the dashed (grey line). The vertical shaded regions mark the two mock trading periods and six compliance periods in the emissions market. The horizontal (red) lines denote the per-plant month market cap for each compliance period, calculated as the aggregate market cap divided by 162 Treatment plants.

In Panel A, 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 the first live trading period, in September (shaded grey), treatment plants are emitting roughly 500 kilograms per month less particulate matter than control plants. This difference is sustained throughout the rest of the sample, though it narrows somewhat in February and March of 2020.

In Panel B, without imputation across plant-weeks, the pattern is similar, but the balance of CEMS-reported pollution prior to the introduction of the market is somewwhat weaker. Mean treatment emissions appear somewhat higher than mean control emissions before the experiment started. Given the strong evidence of balance in several other measures of pollution at baseline, we expect this difference may be due to marginally higher control reporting via CEMS in the pre-baseline period (see Figure A1).

The emissions reductions reported in Figure 3, Panel A 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 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 binds.

b 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:

$$Pollution_{it} = \beta_1 Treatment_i \times Post_t + \beta_2 Treatment_i + \beta_3 Post_t + X_i \beta_4 + \alpha_t + \varepsilon_{it},$$

where $Treatment_i$ is a dummy variable equal to one for plants assigned to the emissions market treatment, $Post_t$ is a dummy variable equal to one after the start of the mock-trading period, X_i are plant characteristics observed in our baseline survey that we expect may predict pollution and α_t are month-of-sample fixed effects. In some specifications we drop the controls X_i in favor of plant fixed effects. Standard errors are clustered at the plant level.

Table 2 reports the results. Panel A has the level of pollution (mass of emissions, in kilograms per month) as the outcome variable and Panel B the logarithm. In column 1, without fixed effects, treatment plants are estimated to have pollution 342 kg per month (standard error 140 kg per month) lower than control plants. There is a small, statistically insignificant difference between treatment and control plants prior to the beginning of the treatment. However, the main effect of $Post_i$ is large, negative and precisely estimated, showing the large decline in pollution that is observed also in the control group. The remaining columns enrich the specification by adding control variables. Column 2 adds month fixed effects and column 3 baseline control variables. These controls increase both the magnitude of the estimated treatment effect and its precision. Column 4 additionally adds plant fixed effects. With plant fixed effects, the coefficient on $Treatment_i \times Post_t$ is -342 kg per month (standard error 147 kg per month).

These are substantial reductions in pollution relative to both the control group and the baseline status quo. Panel B shows the same specifications with the dependent variable in logs. The treatment effect relative to the control group is -0.27 log points, or a 24% reduction in pollution. The coefficient $\hat{\beta}_3$ on post is -0.61 log points, showing that pollution in the control group declined by 46% from before to after the experiment.

Table 3 considers the robustness of these estimates to alternative pollution imputation rules (see Section 3 and Annex Table A3 for a description of the rules). Our baseline estimates are shown in column 1. Column 2 uses no imputation across plant-weeks. Column 3 uses imputation at the 75th percentile of the treatment distribution. Column 4 restricts the sample of the first 100 plants that began reporting in the treatment and control group. The idea is that if this group of earlier reporting plants is similarly selected in both arms, the restricted sample will provide an estimate of the treatment effect within that group.

We find that the imputation rule has relatively small effects on the point estimates but large effects on estimated standard errors, by forcing plant-months with any missing data out of the sample. The point estimates for the effect of treatment on pollution range from -331 kg per month to -406 kg per month, similar to our baseline estimate of -342 kg per month. In the column 2 specification, without imputation across plant-weeks, the number of plant-month observations falls and the estimated standard error increases, so that the treatment effect on pollution is no longer statistically significant. In the column 4 specification, in the restricted sample, the estimated treatment effect is very close to our baseline estimate, but again, less precisely estimated.

We conclude that pollution emissions declined in the treatment relative to the control group in the period after the market was introduced. That relative decline is observed both in the raw data and under plausible alternative assumptions on missing pollution. However, the precision of the estimated treatment effects on pollution degrades when the sample is restricted to plants that are continuously reporting.

5 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 preliminary estimates show that the introduction of emissions trading reduced pollution by 24%, relative to a control group of plants that remained in the command-and-control regime. In ongoing work we will analyze how the introduction of the market affected plant abatement costs and efficiency.

References

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*.

Duflo, 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.

Duflo, 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.

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/.

Lall, Ramona, Michaela Kendall, Kazuhiko Ito, and George D Thurston. 2004. "Estimation of historical annual PM2. 5 exposures for health effects assessment." *Atmospheric Environment*, 38(31): 5217–5226.

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.

6 Figures

Figure 1: Modeled Contribution of Sample Plants to Ambient Pollution Levels at Baseline

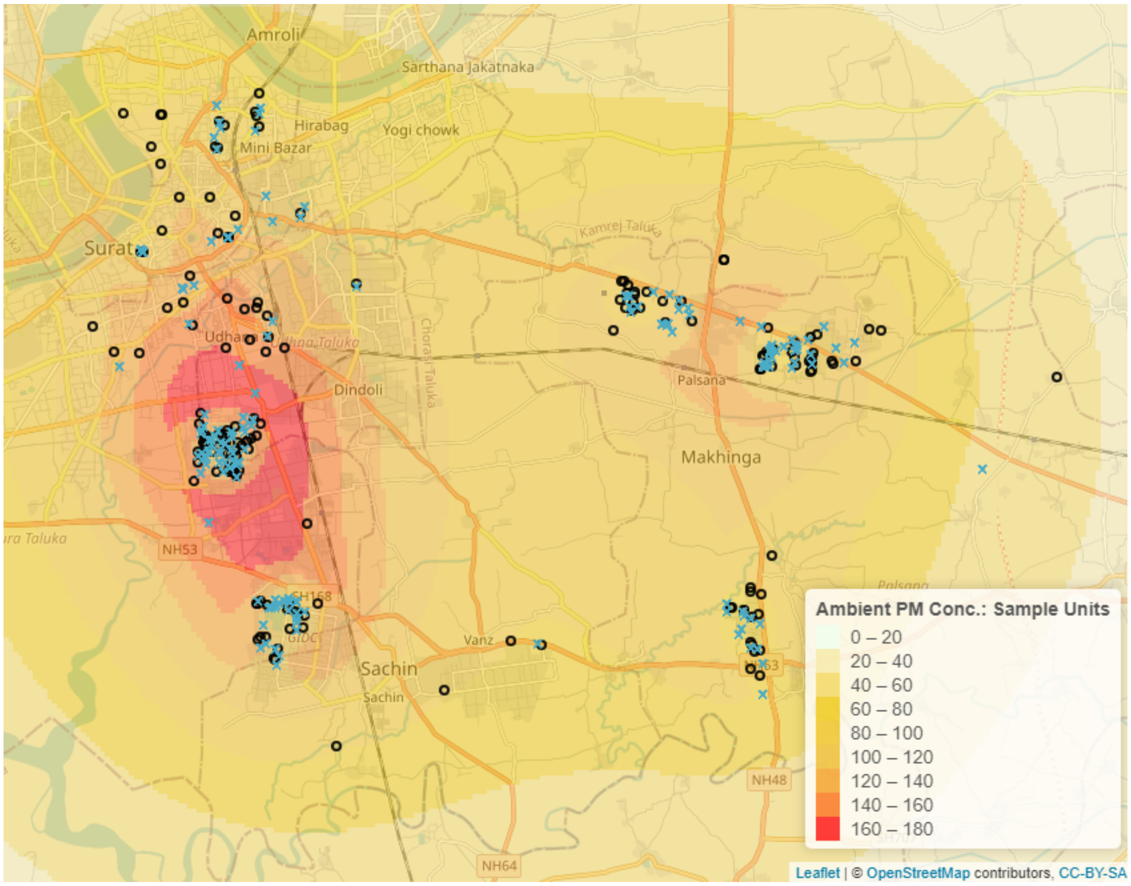

The figure shows the contribution of particulate matter emissions from sample plants to ambient particulate matter concentrations in Surat, Gujarat in June 2019. Treatment firms are represented by blue X markers and control firms by black O markers. 304 plants are shown, with 13 sample plants lying outside the bounds of the map. The ambient PM concentrations (in $\mu g/m^3$) are measured by passing CEMS emissions rates for the sample plants into a simplified Gaussian dispersion model, comparable to the SCREEN3 model used by the US Environmental Protection Agency. The 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.

Figure 2: Distribution of pollution before the experiment

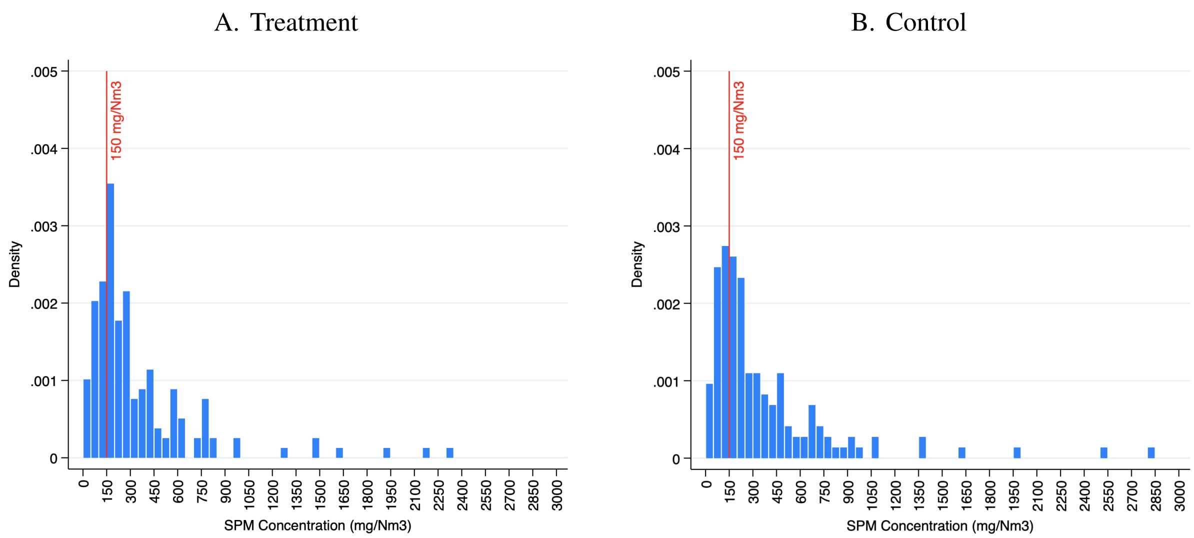

The figure shows distributions of SPM Concentration (mg/Nm^3) as measured by manual iso-kinetic stack sampling at CEMS Baseline (Sep−Dec 2014), by treatment arm. One SPM sample was collected from each industrial stack by a third-party laboratory. Left: 162 in-sample Treatment plants were sampled (mean SPM concentration 352.6 mg/Nm^3). Right: 146 in-sample Control plants were sampled (mean SPM concentration 366.2 mg/Nm^3). The red line indicates the regulatory concentration standard for SPM at 150 mg/Nm^3. An iso-kinetic spot measurement of SPM concentration above this threshold indicates non-compliance under the status quo regime. At CEMS Baseline, 69.2% of sampled Control plants and 73.4% of sampled Treatment plants were reported to be out of compliance with the 150 mg/Nm^3 SPM regulatory limit.

Figure 3: Pollution emissions by treatment arm over time

A. Mean imputation across plant-weeks

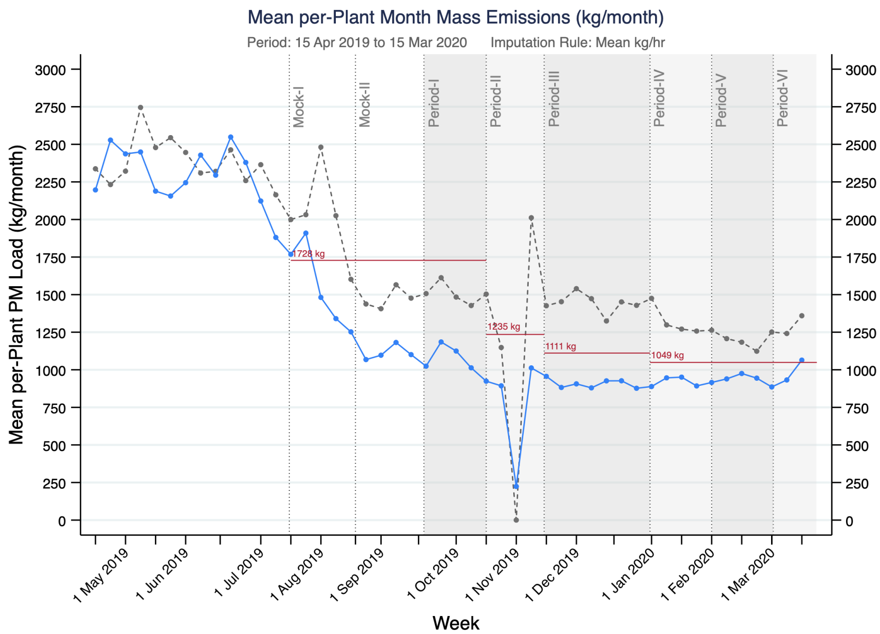

B. No imputation across plant-weeks

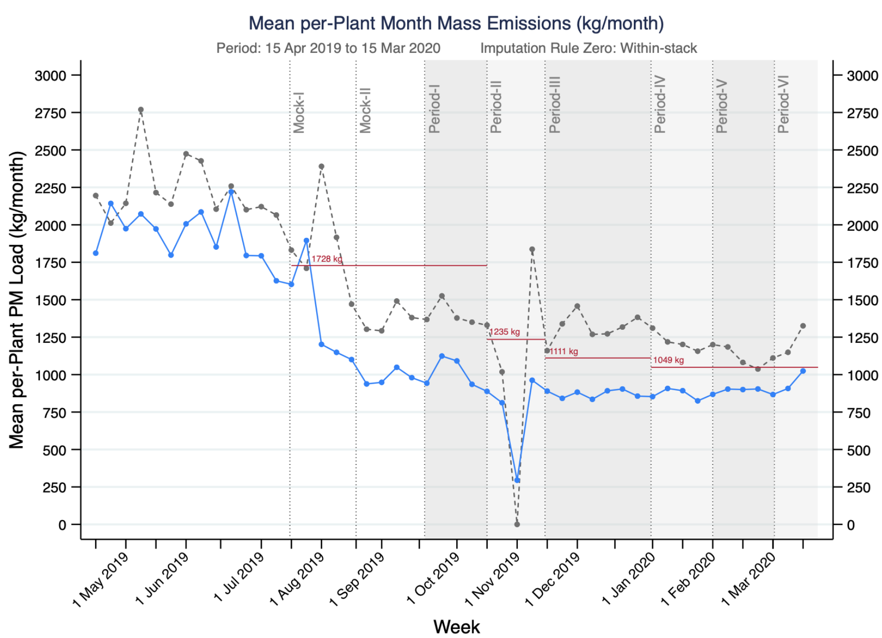

The figure shows the mean per-plant emissions in kilograms per month, at weekly frequency, over 11 months from 15 April 2019 to 15 March 2020. Panel A shows the mean imputation data series where missing plant-weeks are imputed at the treatment group mean for that week, in both the treatment and control arms. Panel B shows the raw data series where missing pollution readings are imputed within a plant-week, but not across plants or weeks. In both panels, treatment firms are represented by the solid (blue) line, control firms by the dashed (grey) line. The vertical shaded regions mark the two mock trading periods and six compliance periods in the emissions market. The aggregate market caps for each compliance period were: 280 tons/month (for mock trading and Period-I), 200 tons/month (for Period-II), 180 tons/month (for Period-III), and 170 tons/month (for Periods IV-VI) respectively. The horizontal (red) lines denote the per-plant month market cap for each compliance period.

7 Tables

Table 1: Balance of Plant Covariates at ETS Baseline

	(1) Control	(2) Treatment	(3) Difference
Panel A. Plant Characteristics			
2016 Sales Revenue (USD millions)[‡]	12.1	11.2	-0.90
	[48.7]	[37.2]	(5.10)
2017 Sales Revenue (USD millions)[‡]	13.4	12.4	-0.99
	[52.2]	[41.8]	(5.57)
2016 Electricity Cost (USD thousands)[‡‡]	357.4	432.3	74.9
	[543.3]	[829.6]	(82.2)
2017 Electricity Cost (USD thousands)[‡‡]	381.3	447.1	65.8
	[647.5]	[836.0]	(87.8)
Work Days, annual	300.2	302.7	2.52
	[20.6]	[26.3]	(2.70)
Boiler Capacity (TPH)	6.72	8.83	2.11
	[4.69]	[16.4]	(1.36)
TFH + HAG Capacity (U)	2182.7	2446.8	264.2
	[1415.6]	[3255.3]	(284.9)
Boiler Vintage (Installation year)[‡‡‡]	2008.1	2007.4	-0.75
	[6.16]	[7.28]	(0.79)
Small-scale (=1)	0.69	0.72	0.033
	[0.47]	[0.45]	(0.053)
Textile sector (=1)	0.85	0.85	-0.0032
	[0.36]	[0.36]	(0.041)
N_{plants}	147	157	304
Panel B. Boiler House Expenditures			
Operating Cost, annual (USD thousands)	108.8	135.3	26.5
	[83.2]	[198.5]	(17.3)
Electricity Cost, annual (USD thousands)	57.1	72.6	15.5
	[69.9]	[150.5]	(13.3)
Aggregate Power Rating (KW)	79.4	95.5	16.1
	[97.5]	[181.9]	(16.6)
Inputs Cost, annual (USD thousands)	9.35	11.9	2.58

Table 1: continued from previous page

	[10.8]	[20.8]	(1.88)
Water Cost, daily (USD)	23.1	29.5	6.41
	[34.1]	[55.4]	(5.24)
Chemical Cost, monthly (USD)	205.5	220.8	15.4
	[271.0]	[264.0]	(30.7)
Maintenance Cost, annual (USD thousands)	8.32	8.30	-0.023
	[6.31]	[6.59]	(0.74)
Labor Cost, annual (USD thousands)	42.3	50.8	8.44
	[26.7]	[55.8]	(4.97)
No. of Workers	31.7	36.8	5.08
	[30.0]	[32.5]	(3.59)
Capital Investment, annual (USD thousands)[‡‡‡‡]	161.0	194.1	33.1
	[187.1]	[390.7]	(36.0)
Repair Costs, annual (USD thousands)	4.20	4.35	0.15
	[4.99]	[5.92]	(0.63)
Major Modifications Cost, annual (USD thousands)	1.99	2.30	0.31
	[3.89]	[8.11]	(0.72)
N_{plants}	147	157	304
Panel C. Abatement Equipment			
At least 1 Cyclone (=1)	0.97	0.98	0.0081
	[0.16]	[0.14]	(0.017)
At least 1 Scrubber (=1)	0.60	0.64	0.038
	[0.49]	[0.48]	(0.056)
At least 1 Bag Filter (=1)	0.86	0.81	-0.048
	[0.35]	[0.39]	(0.043)
At least 1 ESP (=1)	0.082	0.11	0.033
	[0.27]	[0.32]	(0.034)
N_{plants}	147	157	304
Panel D. Baseline Pollution Emissions			
PM Concentration (mg/Nm^3)	175.2	181.3	6.13
	[157.3]	[153.1]	(17.8)
Volumetric Flow Rate (Nm^3/hr)	20976.3	22302.4	1326.0
	[13582.6]	[25977.6]	(2356.7)

Table 1: continued from previous page

Mean pre-Mock Ringelmann Score[‡‡‡‡‡‡]	1.36	1.36	0.0064
	[0.37]	[0.43]	(0.046)
N_{plants}	147	155	304

Table Notes: Columns (1) and (2) show means with standard deviations in brackets. Column (3) shows the coefficient on treatment from regressions of each characteristic on treatment with no fixed effects. Statistical significance indicated as: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The plant-level emission variables in Panel D (*Plant PM Concentration, Plant Volumetric Flow Rate, Plant PM Mass Rate*) are calculated as the sum of each emission quantity across all stacks at a given plant.

[‡] Due to surveyor typos, 6 Control and 7 Treatment observations omitted.

[‡‡] Due to missing values and/or surveyor typos, 7 Control and 9 Treatment observations omitted.

[‡‡‡] N_{plants} for *Mean Boiler Vintage* is 136 Control and 151 Treatment observations.

[‡‡‡‡] N_{plants} for *Boiler House Capital Investment* is 137 Control and 147 Treatment observations.

[‡‡‡‡‡‡] *Mean pre-Mock Ringelmann Score* is calculated as the mean plant Ringelmann score over the period April-June 2019 (Ringelmann rounds 7-10). N_{plants} for *Mean pre-Mock Ringelmann Score* is 143 Control and 156 Treatment observations.

Table 2: Estimates of Treatment Effect on Pollution Outcomes

	(1)	(2)	(3)	(4)
	No FE	Month FE	Controls+ Month FE	Controls+ Month FE+ Plant FE
Panel A. Y = Plant-month Emissions				
ETS Treatment=1 × Post=1	-342.2**	-431.0***	-519.4***	-346.1**
	(139.9)	(115.0)	(95.67)	(147.5)
ETS Treatment=1	-88.77			
	(172.0)			
Post=1	-914.3***			
	(92.86)			
Adjusted R^2	0.125	0.144	0.292	0.618
Control mean pre-Mock Y	2405.7	2405.7	2405.7	2405.7
Panel B. Y = Ln(Plant-month Emissions)				
ETS Treatment=1 × Post=1	-0.270***	-0.392***	-0.414***	-0.268***
	(0.0734)	(0.0667)	(0.0625)	(0.0767)
ETS Treatment=1	-0.121*			
	(0.0656)			
Post=1	-0.611***			
	(0.0461)			
Adjusted R^2	0.201	0.250	0.317	0.595
Control mean pre-Mock *Ln(Y)*	7.656	7.656	7.656	7.656
Month FE		Yes	Yes	Yes
Plant FE				Yes
Observations *(Plant-Month)*	3487	3487	3487	3487
N *(Plants)*	317	317	317	317

Table Notes: Standard errors in parentheses; clustered by *industry_id*. Statistical significance is indicated as: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Plant-Month observations are from a balanced panel of 11 months (over the period 16 April 2019 - 15 March 2020). The dependent variable in Panel A is Plant-month SPM Emissions in kg per month, constructed using the Mean Imputation Rule for missing values of mass emission rate. Columns (3) and (4) include the following baseline control variables, not reported in the table: *Boiler House CapEx, Boiler Capacity, Boiler Year, Month Permit Allocation, 2018 Boiler House Operating Cost.* Of these, *Month Permit Allocation* is time-variant.

Table 3: Robustness of Pollution Estimates to Imputation Rules

	(1)	(2)	(3)	(4)
	Mean Imp	Rule Zero Imp	P75 Imp	Mean Imp *N*=First 100
	Panel A. Y = Plant-month Emissions			
ETS Treatment=1 × Post=1	-342.2**	-406.5	-351.9**	-331.1
	(146.9)	(410.7)	(153.3)	(227.5)
Adjusted R^2	0.618	0.555	0.623	0.576
Control mean pre-Mock *Y*	2405.7	2236.4	2599.5	2453.8
	Panel B. Y = Ln(Plant-month Emissions)			
ETS Treatment=1 × Post=1	-0.270***	0.0147	-0.282***	-0.202*
	(0.0771)	(0.196)	(0.0778)	(0.112)
Adjusted R^2	0.594	0.528	0.600	0.556
Control mean pre-Mock *Ln(Y)*	7.656	7.224	7.733	7.609
Month FE	Yes	Yes	Yes	Yes
Plant FE	Yes	Yes	Yes	Yes
Observations *(Plant-Month)*	3487	2284	3487	2200
N *(Plants)*	317	287	317	200

Table Notes: Standard errors in parentheses; clustered by *industry_id*. Statistical significance is indicated as: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Plant-Month observations are from a balanced panel of 11 months (over the period 16 April 2019 - 15 March 2020). The dependent variable in Panel A is Plant-month SPM Emissions in kg per month, constructed using imputation rules for missing values of mass emission rate as noted in each specification name.

A Appendix: Data

Table A1: Sample Determination and Attrition by Treatment Arm

	Control	Treatment	Total
Plants that received treatment assignment	168	174	342
Closed/extinct plants with treatment assignment	*11*	*10*	*21*
Operational-at-baseline plants with treatment assignment	157	164	321
Plants removed from ETS sample	*2*	*2*	*4*
In-sample Plants	**155**	**162**	**317**
Plants incompletely treated due to closure	*8*	*6*	*14*
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 reporting CEMS data	137	156	293
Treatment plants with market bids data	–	154	154
In-sample plants with GPCB administrative data	155	162	317

Table Notes: A total of 342 operational plants received treatment assignment in May 2019 *(row 1)*. Of the 342 plants included in the ETS treatment randomization, 21 plants were extinct or permanently closed *(row 2)*. The permanent shutdown status of these 21 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 21 plants who received assignment while extinct or shutdown, yield 321 operational plants with treatment assignment at baseline *(row 3)*. Four of these 321 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. Therefore, this paper reports experimental results from a sample of 317 plants *(row 5)*. Of the 317 in-sample plants, 14 are known to have been ”incompletely treated” by the intervention, due to temporary financial closure before or after the treatment assignment was done *(row 6)*. These incompletely treated plants are included in all analysis reported in this paper; where they were not reporting pollution data from CEMS, their emissions have been imputed. The 304 plants surveyed at baseline are distinct from the 304 plants manually sampled at Baseline and are therefore reported separately *(rows 7, 8)*.

Table A2: Treatment Assignment and CEMS Reporting

	(1) Y=CEMS Reporting
ETS Treatment=1	0.0855***
	(0.0303)
N *(Plants)*	317

Table Notes: Standard errors in parentheses; clustered by Plant ID. Statistical significance is indicated as: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is a dummy indicating whether the plant has reported from CEMS during the 11-month period from 16 April 2019 - 15 March 2020.

Table A3: Imputation Rules for Missing Values of Stack day kg/hr

		(1)	(2)	(3)
Step	Consideration	Rule 0 **Within-stack**	Rule 1 **Mean kg/hr**	Rule 2 **P75 kg/hr**
Step 0	Imputation Level	Stack day PM Mass Rate (kg/hr)	Stack day PM Mass Rate (kg/hr)	Stack day PM Mass Rate (kg/hr)
Step 1	Stack day kg/hr truncated at p99?	Yes	Yes	Yes
Step 2	Impute for missing values of kg/hr:	Stack week mean kg/hr	Stack week mean kg/hr	Stack week p75 kg/hr
Step 3	If no week kg/hr or reported mass (kg), impute for missing values:	-	Treatment arm week mean kg/hr	Treatment arm week p75 kg/hr

Figure A1: Data Availability from CEMS

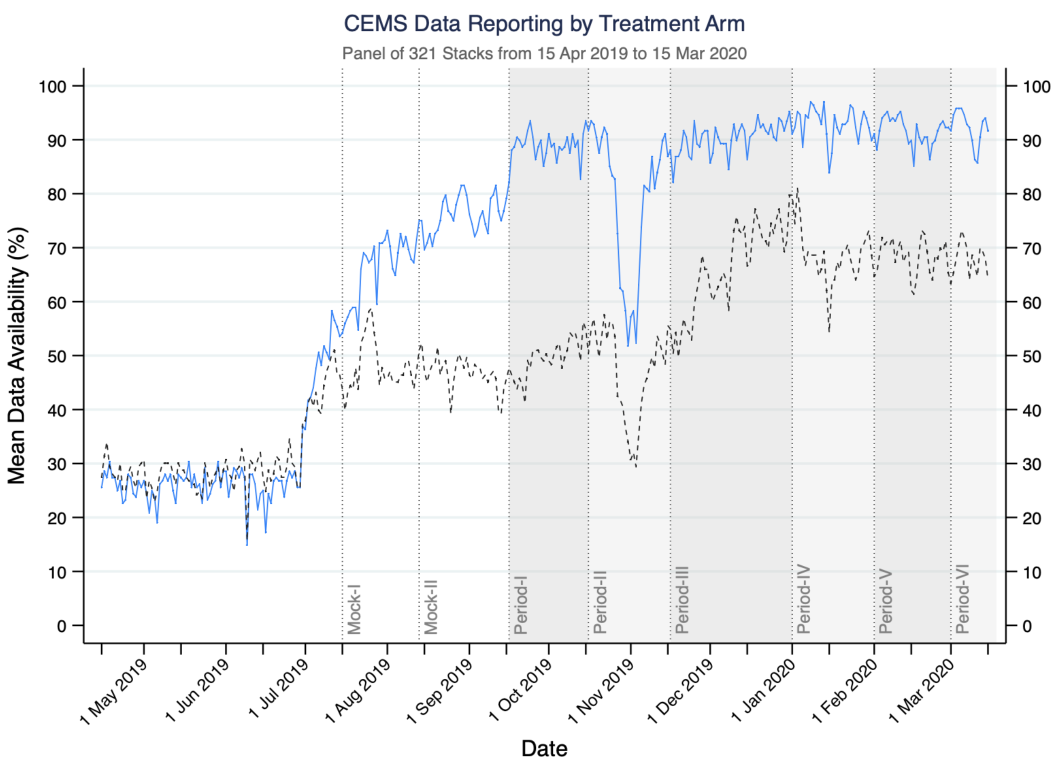

The figure shows the CEMS data availability, at daily frequency, from the sample of 317 plants (comprising 321 industrial stacks) over 11 months from 15 April 2019 to 15 March 2020. Treatment firms are represented by the solid (blue) line, control firms by the dashed (grey) line. The vertical shaded regions mark the two mock trading periods and six compliance periods in the emissions market. The dip in November 2019 reflects week-long plant shutdowns following the Diwali festival.