

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

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

January 27, 2023

Abstract

Market-based environmental regulations are seldom used in developing countries, where pollution is the highest but state capacity is often low. We experimentally evaluate a new particulate matter emissions market, the first in the world, covering industrial plants in a large Indian city. There are three main findings. First, the market functioned well: permit trade was active and plants obtained permits to meet their compliance obligations almost perfectly. Second, treatment plants, randomly assigned to the emissions market, reduced pollution emissions by 20% to 30%, relative to control plants. Third, the market, holding emissions constant, reduces abatement costs by 11% to 14%. These cost estimates are based on a model that estimates heterogeneous plant marginal abatement costs from plant bids for emissions permits. More broadly, we find that emissions can be reduced at small increases in abatement costs. The pollution market therefore has health benefits that exceed costs by at least twenty-five times.

*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. This project is registered in the AEA RCT registry as AEARCTR-0003860. We thank Gargee Goswami, Dipika Gawande, Prajval Jhunjhunwala, Bogdan Mukhametkaliev, Gargi Pal, Shruti Bhimsaria, Vineet Gupta, Winston Hovekamp, Jared Stolove, Jeanne Sorin, Fikremariam Gedefaw, Kaixin Wang and Neil Himwich for excellent research assistance and Sanjana Gorti, Jasdeep Mandia, and Amber Luong for project management. We thank Jairam Ramesh for encouraging this project.

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

1 Introduction

Air pollution harms people by shortening their lives and reducing the human capital they form as children (Ebenstein et al., 2017; Isen, Rossin-Slater and Walker, 2017). With these harms in mind, India may be considered the epicenter of a global air pollution crisis. Nearly all of India's 1.4 billion people breathe air that exceeds WHO standards for particulate matter, often by a factor of ten or more. Reducing air pollution throughout India to the WHO standard would, by some estimates, increase Indian life expectancy by an average of five years (Pande et al., 2015; Greenstone, Hasenkopf and Lee, 2022).

India has stringent environmental regulations on the books (Piette, 2018). In practice, however, enforcement of these regulations is uneven, and many plants choose to remain out of compliance with pollution standards (Duflo et al., 2013, 2018). It may be that it is costly for polluters to cut emissions in an industrializing economy. It may also be that abatement, privately, is not too costly, but that plants have little reason to abate, given gaps in the enforcement of existing command-and-control environmental regulations.¹

Can pollution markets work in developing countries? Markets, in theory, can abate pollution at the lowest possible cost (Dales, 1968). In the United States and the European Union, pollution markets have been hailed for turning sound economic principles into policy (Ellerman et al., 2000). Yet, there has been little similar progress in developing countries, which are presumably more cost sensitive, but still rarely use markets to regulate pollution (Stavins, 2003; Blackman, Li and Liu, 2018).

One plausible reason that developing countries do not regulate with markets is that the basic assumptions of an emissions market do not hold in low-capacity states. In pollution market models, emissions are known, plants always comply, and permit trade happens without any transaction costs (Coase, 1960). In developing countries, regulators struggle with monitoring and may lack

¹A complementary explanation for high pollution is that developing countries are unwilling to pay much for environmental quality, because their citizens are poor and demand other essentials first (Ito and Zhang, 2020; Baylis et al., 2023; Jack et al., 2022).

the credibility to force each polluter to hold a permit for every unit of emissions. Put differently, emissions markets trade in a commodity created by the state, and these markets cannot function if the state cannot uphold the value of that commodity.

This paper is an experimental evaluation of how a new pollution market in India impacts emissions and compliance costs. Our experiment established the world's first market for particulate matter emissions. The market is the result of a more than decade-long collaboration with the Gujarat Pollution Control Board (GPCB), the environmental regulator in Gujarat, India. In 2010, we authored a white paper on the potential of emissions trading for regulation in India (Duflo et al., 2010). Thereafter, we collaborated with GPCB as it assembled the building blocks for an emissions market. GPCB and its partners set standards for continuous emissions monitoring systems (CEMS), mandated that plants install CEMS, laid down the rules for the market and built a trading platform.

The market for particulate matter emissions from industrial plants was launched in 2019. GPCB mandated CEMS for large particulate sources, totaling 317 industrial plants, in the airshed of Surat, a city of 7 million people. All sample plants were mandated to install CEMS for particulate matter (PM) and did so before the start of the experiment. The status quo regulatory regime for these plants is a command-and-control regime in which plants are mandated to install abatement equipment and are subject to an intensity standard on the maximum concentration of particulate emissions. The status quo regulation is incompletely enforced (Duflo et al., 2013, 2018), and our baseline data show that nearly a third of plants are out of compliance with these standards.

A randomly assigned treatment group of 162 plants were then shifted into the new emissions market, which traded in permits worth one kilogram of particulate matter apiece. At the beginning of each compliance period, permits were allocated to treatment plants both for free and at auction, to aid price discovery. Treatment plants, thereafter, could only trade permits with other treatment plants. The control group remained in the status quo regime. The evaluation ran from September 2019 to April 2021, with an interruption due to a nationwide Covid-19 lockdown.

Conceptually, the treatment group regulation differs from that in the control in three respects.

First, the essence of market-based regulation: the compliance obligation for each treatment plant is tradeable, rather than a fixed standard, as in the control group. Second, treatment plants are subject to standards that limit total pollution load (i.e., mass), rather than concentration at one point in time. Third, the stringency of regulation in the two regimes may differ, because of differences in how load and concentration standards are set.

There are three main findings. First, the market functioned well. Nearly all plants held enough permits to cover emissions in all compliance periods. The regulator assured compliance by establishing a reputation for enforcement early on.² Qualitatively, permit trade was fluid. Plants traded often and trading volume was high, up to 20% of the market cap on some single days. Plant permit holdings at the end of each compliance period differed greatly from initial allocations and plants left little money on the table, in terms of unused and unsold permits at the end of compliance periods. The large and yet precise changes in permit holdings through trade show that any transaction costs in the permit market were small. The fidelity to the market rules, in contrast to widespread non-compliance in the status quo, suggests that the market facilitated enforcement.

Second, the treatment caused a 20% to 30% decline in particulate emissions, relative to the control plants in the command-and-control regime. The range of point estimates reflects differences in how emissions are imputed during periods when plants fail to report data through CEMS. The reduction in emissions is a function of the emissions cap: based on incomplete data, the cap had initially been set with the goal of making the load standard in the treatment as stringent as the concentration standard in the control. The regulator adjusted the cap downwards over the first several compliance periods. The emissions treatment effect therefore reflects both improved compliance and an endogenously greater stringency of regulation.

The promise of an emissions market is to reduce abatement costs. We begin our analysis of costs by estimating treatment effects in plant survey data, finding null effects of the market treatment on abatement capital and variable expenditures. Abatement expenditures are hard to

²Two firms failed to turn in enough permits at the end of the first compliance period. As specified in the market rules, the regulator responded by levying monetary damages at twice the level of the permit price ceiling for each unit of excess emissions. Compliance thereafter was essentially perfect.

measure in surveys, however, and conceptually the right measure of cost changes due to a market would hold emissions constant—which is not the case across our treatment arms. We therefore favor revealed-preference estimates of abatement cost that we derive from data on the universe of 8,433 permit bids in the treatment group. In a competitive market, a plant’s willingness-to-pay for permits equals their marginal abatement cost. We use variation in plant permit bids, at different levels of emissions within a plant and period, to estimate the shape of plant-period specific marginal abatement cost functions. We then apply these functions to compare costs under the market, at a range of regulatory stringencies, and under counterfactual command-and-control regimes, which we estimate to match the distribution of emissions in the control group.

Third, with this approach, we find that the market reduces variable abatement costs by 11% at the treatment level of emissions. This estimate comes from model counterfactuals comparing abatement costs in the market against costs if plants were required to meet the market emissions cap under a command-and-control regime. The cost savings are higher, at 14%, if emissions are held constant at the control level instead. More broadly, our estimates show that plant emissions can be reduced quite cheaply in Gujarat. Our benefit-cost analysis of a potential expansion of the emissions market shows that, under a range of plausible assumptions on the health damages from pollution, the benefits of the market exceed costs by at least twenty-five times.

A main theme of the literature at the intersection of environmental economics and development has been the Herculean difficulty of environmental regulation in low-state-capacity settings (Greenstone and Hanna, 2014; Jayachandran, 2022). Common findings are that poor or corrupted monitoring impedes regulation (Duflo et al., 2013; Oliva, 2015; Duflo et al., 2018; Zou, 2021), and that coarse regulations, themselves adopted in response to poor monitoring, are partly undercut by behavioral responses (Davis, 2008; He, Wang and Zhang, 2020).³ Our findings suggest that, if the

³The one prior example of a market targeting particulates with which we are familiar in fact regulated boiler capacity, not emissions, because of limitations in pollution monitoring at the time (Montero, Sanchez and Katz, 2002). The market in Santiago, Chile targeted particulate matter emissions indirectly and achieved little abatement, because it traded in boiler capacity, a proxy for emissions, rather than directly in emissions load. This coarse proxy removes the prospect of low-cost abatement because most particulate abatement happens after combustion, rather than by limiting the size of plants or how much fuel they can burn. Ultimately, the market disbanded after many covered sources switched fuels in response to a fall in natural gas prices.

enforcement problem can be addressed, the *private* costs of abatement may not be high, even in settings with high emissions and ambient pollution (see also Greenstone et al., 2022).

The literature on market-based instruments has focused on the landmark US environmental markets, including the RECLAIM program (introduced in 1994 to target SO_2 and NO_X), the Acid Rain program (1995 to target SO_2) and the NO_X Budget Trading Program (2003), and more recently on carbon markets such as the EU ETS (2005) and California’s AB32 (2013). Although it is conventional wisdom that emissions markets abate pollution at lower cost than extant command-and-control regulations (Ellerman et al., 2000; Fowlie, Holland and Mansur, 2012), most evaluations of emissions markets require strong assumptions to estimate what plant costs and emissions would have been without a market (Fowlie, Holland and Mansur, 2012; Martin, de Preux and Wagner, 2014; Borenstein et al., 2019).⁴ This study estimates the effect of emissions trading on both pollution and plant costs against a well-defined experimental counterfactual.⁵ The emissions market we study is a policy innovation in a setting where, like in many parts of the world, regulators with a constrained set of tools face crisis levels of air pollution.

The paper proceeds as follows. Section 2 introduces status quo environmental regulation in India and the design of the emissions market. Section 3 details the experimental design and data and presents some summary statistics on the sample and experimental balance. Section 4 presents the paper’s first main result, on how the market functioned in the treatment group. Section 5 presents estimates of experimental treatment effects on emissions and abatement costs. Section 6 describes the model and estimates marginal abatement costs from bids. Section 7 uses the model to compare the costs of abatement under the two regulatory regimes. Section 8 concludes.

⁴On the difficulty of developing a counterfactual, 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.”

⁵Martin, Muûls and Wagner (2016) write “An ideal evaluation of the EU ETS would combine a representative firm- or plant-level data set of sufficient detail with a study design that attributes to the EU ETS only those observed behavioral changes it has actually caused. It is difficult to solve this identification problem because there are so many factors that might simultaneously affect firm behavior, thus confounding the impact estimate. The state-of-the-art solution would be to conduct a randomized control trial or field experiment (e.g., Greenstone and Gayer (2009)). As in other real-world settings, however, randomizing participation in the EU ETS is neither desirable nor politically feasible.”

2 Status Quo Regulation and Emissions Market Design

This section outlines the status quo regulatory approach and its shortcomings. It then describes our collaboration with the Gujarat Pollution Control Board (GPCB) and the design of the emissions trading market. We postpone the experimental design to Section 3.

2.1 The Command and Control status quo

The Air Act of India (1981) established a command and control framework for regulating industrial pollution.⁶ The “command” is a mandate that plants, when they seek to open or to expand, must install equipment capable of limiting pollution. The “control” is an intensity standard, under which plants can be sanctioned if their pollution emissions, measured with a manual pollution sample during a spot check, exceed a concentration limit (typically $150 \text{ mg}/\text{Nm}^3$). Aspects of this framework, such as emission limits and measurement protocols, are set by India’s Central Pollution Control Board (CPCB), with enforcement delegated to State Pollution Control Boards (SPCBs).

Prior research, by some of us and in partnership with the GPCB, has shown that compliance with the status quo regime is far from complete. Imperfect compliance could be the result of weak enforcement or of high abatement costs. On enforcement, our prior research showed that the regulator has difficulty in obtaining reliable third-party reports of pollution (Duflo et al., 2013). While penalties are large, they are reserved, in practice, for the largest offenses, and therefore give weak incentives for plants with more routine violations to abate (Duflo et al., 2018). On abatement costs, the regulator conditions abatement equipment mandates on plant scale (boiler capacity). However, pollution standards are constant for all plants, regardless of age, scale, fuel, and the like, which may impose inefficiently high costs on some plants. For example, holding abatement equipment and pollution concentrations constant, small factories that operate short shifts do much less environmental damage than large plants operating 24 hours a day, 7 days a week. Status quo regulation does not differentiate between these two kinds of plants.

⁶The Water Act (1974) established State Pollution Control Boards as state environmental regulators, giving them authority to enforce water pollution standards. The Air Act (1981) expanded these powers to cover air pollution. SPCBs can also introduce additional regulation in highly polluted regions.

2.2 The Creation of Surat's Emissions Market

The GPCB became interested in piloting an emissions market due to the limitations of the control and command regulatory system. After proposing the use of emissions markets in a white paper (Duflo et al., 2010), we worked with the GPCB to develop the monitoring, trading and regulatory infrastructure for an emissions market.

Monitoring particulate emissions.—On monitoring, we proposed that the government introduce Continuous Emissions Monitoring Systems (CEMS). These devices are common in large plants in the United States and Europe but had never been mandated in India. Anant Sudarshan, one author of this paper, then sat on a Central Pollution Control Board panel that created the first technical standards for the use of CEMS in India (Central Pollution Control Board, 2013). Once these standards were published, GPCB mandated CEMS for a sample of plants in Surat in 2013. The Central Pollution Control Board (CPCB), in 2014, also mandated the installation of CEMS for a variety of pollutants in large plants in 17 manufacturing sectors.

Surat's CEMS rollout posed implementation challenges and was only completed in 2019. These challenges were caused by a combination of technical, institutional and economic factors (Sudarshan, 2023). The main issue, as seen with pollution abatement equipment, was not whether the CEMS machines could be installed, but whether they would be run often enough to measure PM emissions reliably. After testing with support from CEMS vendors, the GPCB concluded in 2019 that CEMS readings could form the foundation for a market.

Plants in the status quo (control arm) remained governed by the same command-and-control regulations as described above, despite being equipped with CEMS.⁷ The presence of CEMS in conjunction with pre-existing command-and-control rules reflects current practice across India. After the 2014 mandate, CEMS data is now transmitted from plants across the country to the CPCB, but it remains legally inadmissible in court and, to our knowledge, is not used widely by SPCBs in regulatory enforcement (one prominent exception is in the state of Tamil Nadu).

⁷Part of the reason is regulatory: the CEMS data was not legally admissible, as the basis for enforcing concentration standards, because national rules say compliance with these standards is determined by in-person readings. Another reason is more bureaucratic: the GPCB did not incorporate CEMS into its routine for selecting what plants to inspect.

Market logistics.—Starting an emissions market required that the regulator improve monitoring, issue new regulations, establish a trading platform and build market participants’ capacity. Our research team helped GPCB design the market rules and evaluate its implementation. NCDEX e-Markets Limited (NeML) helped design the market and hosted it on a custom-built trading platform. NeML is a subsidiary of the National Commodity and Derivatives Exchange, which develops and operates several commodity exchanges in India. GPCB, our research team, and NeML developed the market rules, which the Forest and Environment Department, Government of Gujarat notified (GVN-2019-17-GPCB-SFS-1-2019-ETS-T, dated June 4th, 2019).

GPCB disseminated information about the market through a series of capacity-building workshops with all stakeholders. Our research team trained firms on the market rules, data quality, emissions and permit accounting and other relevant topics. CEMS manufacturers trained firms on how to set up monitoring and reporting. NeML explained how to conduct trades and participate in auctions. The South Gujarat Textile Processors Association (SGTPA), the local industrial association for most sample plants, was instrumental in gathering feedback from firms and ensuring high participation.

2.3 Market Design

While the market we study is the world’s first for particulate emissions, its design is similar to that of other countries’ markets for other pollutants. Here we summarize the key design elements.

Regulated entities.—The market regulates industrial plants with a high potential for air pollution emissions, which, in Surat, are primarily in the textile processing industry. Section 3 describes the eligibility criteria and the characteristics of sample plants.

Permit definition and compliance periods.—Emissions are the total mass of particulate matter released during a compliance period. A permit lets plants emit 1 kg of particulate matter (PM). As is typical for regulations of particulate emissions, rather than ambient pollution levels, the permit definition does not differentiate by particle size. Permits are valid for one compliance

period only; there is no banking or borrowing. Market compliance periods ranged from one month to six weeks. Appendix Table A1 gives the timeline of ten compliance periods. The India-wide Covid-19 lockdown interrupted the market from April to December 2020.

Cap.—Status quo regulation limits the *concentration* of pollution in a plant’s stack gas emissions. The cap in the market instead limits the *load*, or mass, of pollution emitted. Load is pollution concentration multiplied by the volume of gas emitted, which depends on plant capacity utilization and the gas volume per hour. GPCB initially capped particulate emissions at 280 tons per month. The regulator aimed to set the cap to approximate the status quo regulatory stringency and tighten over time. The cap was set to match the existing concentration standard under ex ante assumptions on flow and capacity utilization, which turned out to be too high ex post. As more accurate data became available, the cap was lowered to 170 tons per month. Appendix Table A2 gives each period’s cap.

Permit allocation.—Plants are given 80% of permits for free in proportion to their share of total emissions capacity. Plant emissions capacity, in tons of steam per hour, is calculated by adding the capacities of the boiler and thermic fluid heater, the two main fuel-burning pieces of equipment. Plant capacity is measured with administrative data fixed before the market was announced. Plants therefore could not game or adjust capacity measures in response to the pro rata allocation rule. The balance of 20% of permits are allocated to plants via a uniform price, multi-unit auction. The auction is two-sided; both plants and the regulator can bid to buy permits or offer to sell permits. Each compliance period opens with an auction in which GPCB offers 20% of the cap in permits at the market floor price. If GPCB does not sell all these permits at the first auction, they are offered again in subsequent weekly auctions.

Permit trade.—Plants can trade permits at weekly auctions or over-the-counter (OTC) trades between auctions. NeML runs both markets. Plants can offer step functions from price to the quantity of permits they seek to buy or sell, but in practice bid single quantities. The auction clearing price is the lowest price such that net quantity demanded is weakly negative. OTC trades

occur between weekly auctions, in any quantity, but only at the price revealed by the most recent weekly auction. This OTC price restriction was adopted to encourage auction participation and limit price volatility.

Price collar.—Permit prices were limited to be between INR 5 and 100 per kg. The range of the price collar was informed by engineering estimates that particulate matter abatement by equipment commonly used in the sample, could occur at an average cost of INR 10 to 40, depending on the type of equipment installed and plant scale. The ceiling price is sufficiently high that plants would rather abate than pay the ceiling price for permits. The GPCB committed to trade to enforce the price collar. The floor price was supported by a GPCB commitment to buy back permits at the floor price at the end of each compliance period. 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.

Replacement rule for missing CEMS data.—Plants that did not report data for any length of time had emissions replaced with a rule that made increasingly punitive assumptions based on the extent of missing data (see Appendix Table C3). The purpose of the replacement rule was to incentivize complete and accurate reporting. Emissions with replacements for missing data are called validated emissions and used for the calculation of compliance. Our analysis of emissions levels uses different imputation rules, meant to be unbiased, not punitive (see Section 3).

Compliance and penalties for non-compliance.—At the end of each compliance period, when plants knew emissions with certainty, a one-week true-up period was held during which plants could trade to reach compliance. After the true-up period plants must surrender enough permits to cover their cumulative emissions during the compliance period. Plants with insufficient permits were fined twice the ceiling price for every unit of emissions in excess of their permit holdings.

To show the regulator's commitment to levying this penalty, and to ease implementation, plants at the start of the market post an environmental bond, called an Environmental Damage Compen-

sation Deposit (EDCD).⁸ The term Environmental Damage Compensation was chosen because enabling regulations empower SPCBs to levy fines as compensation for environmental damage.⁹ The size of the EDCD depends on plant scale. The EDCD for most plants is INR 200,000, well in excess of permit expenditures at realized market prices. Any later fine is then deducted from this deposit. The status quo concentration standard was not legally repealed for treatment plants. However, the GPCB did not pursue enforcement under the old PM standard if plants were in compliance with the market permit-holding requirement.

3 Experimental Design, Data, and Summary Statistics

3.1 Experimental Sample and Timeline

Surat, Gujarat is an industrial hub with a population of 7.5 million. Factories contribute an estimated 32% of ambient fine particle pollution in and around the city (Guttikunda, Nishadh and Jawahar, 2019).¹⁰ The industrial share of ambient fine particles is twice as large as that of the next most significant source, transportation, at 16%.

Our experimental sample consists of the plants in and around Surat with the highest air pollution potential. The initial sample included all 342 plants in the GPCB's records that met the following criteria: (i) the plant consumes solid fuel (coal or lignite, mainly); (ii) has boiler capacity of at least one ton of steam per hour; (iii) has a stack diameter of at least 24 cm (which is needed for CEMS). All of these plants were mandated to install CEMS. After CEMS installation, plants were assigned to the market treatment with probability one-half. Some plants were deemed ineligible, after being assigned but before the market started, because they were closed or operated only seasonally. The market started with 162 plants in the treatment group and 156 in the control

⁸While this particular EDCD mechanism is new to the market, GPCB already commonly used environmental bonds as a means of regulatory enforcement. In the status quo, such bonds are used only for plants that incur some violation, as a guarantee against future non-compliance.

⁹This power is derived from the "polluter pays principle" widely recognized in Indian environmental law (Piette, 2018). A relevant precedent is the ruling of the National Green Tribunal (NGT), India's environmental high court. The NGT has directed that the Central Pollution Control Board "may also assess and recover compensation for damage to the environment" (WP (CIVIL) No. 375/2012, Paryavaran Suraksha Samiti vs. Union of India & Others).

¹⁰A current assessment for Surat can be viewed at: <https://urbanemissions.info/india-apna/surat-india/>

group. Appendix Table A3 summarizes plant attrition by treatment arm.

Figure 1 maps the sample plants and the ambient concentration of fine particles in Surat. Sample plants, distinguished by \times (treatment) and \circ (control) markers, lie close to the center of population. Most sample plants are to the south-east of the bend in the river Tapi, which curves around the center of the city. Dense clustering of plants implies that they contribute to pollution in the same areas. The average pairwise distance between sample plants is 11 km, far less than the dispersion of particulate matter from a high plant stack (Guttikunda, Nishadh and Jawahar, 2019). Since plants are close by, relative to how far pollution spreads, trade among nearby plants is unlikely to generate areas of locally increasing pollution. The ambient shading in the background shows that mean fine particulate (PM_{2.5}) concentrations in Surat are ten to twenty times the WHO standard of $5 \mu\text{g}/\text{m}^3$.

The experiment ran for ten compliance periods over about one-and-a-half calendar years. Appendix Table A1 shows the timeline. The first of two mock compliance periods, meant to familiarize plants with trading, began on July 16th, 2019. During mock periods, the market rules were the same as described above, but plants were endowed with fake money and there were no real sanctions for non-compliance. Six real compliance periods followed from September 16th, 2019 to March 22nd, 2020. Market operations were then suspended, as most of India’s economy, including our sample plants, shut down in the nationwide Covid-19 lockdown. The market restarted in December, 2020 and we have data from four additional compliance periods up to the second (Delta) wave of Covid hit in April, 2021. The treatment thus spans roughly one year of market operations spread over one-and-a-half calendar years.

3.2 Data sources

The paper relies on three main sources of data. First, two waves of a plant survey conducted prior to market launch and roughly one year after the market started, respectively. Second, manual measurements of pollution before the start of the market and high-frequency pollution data from Continuous Emissions Monitoring Systems (CEMS) spanning the market operation period. Third, trading data from the market operator covering the universe of plant bids and trades.

Plant surveys.—An in-person baseline survey of plants was conducted from December 2018 to January 2019. A phone-based survey wave was conducted in November 2020. The survey had two parts, general and technical. The general part was administered to the plant owner or manager, and covered plant economic characteristics like inputs, outputs, sales and energy consumption. The technical part directly observed the abatement equipment installed on every point source of emissions in the plant. Our team recorded the characteristics of all emissions sources and abatement equipment and interviewed plant staff about costs of equipment operation.

Pollution measurement.—CEMS installed in treatment and control plants provide high-frequency air pollution emissions data. CEMS devices are calibrated by comparing their readings to physical pollution samples taken simultaneously in the same stack. CEMS readings measure particulate matter (PM), which includes particles of all sizes and is the basis of status quo concentration standards.

The main data concern with CEMS is incomplete reporting during the experiment. The data handling system stores data locally during brief internet outages. However, longer outages, temporary closures, and device malfunctions leave data gaps. The mean rate of weekly data reporting rose from low levels to over 85% of plants by the end of the sample. CEMS data availability is higher in treatment plants than in the control, especially at the start of the market (see Appendix Figure C1). All sample plants were mandated to report CEMS data, but treatment plants had stronger incentives to report, in practice, because market rules that replace missing data at a high value essentially charge treatment plants for non-reporting (see Section 2). Thus, an unanticipated benefit of the market was to speed up plants' willingness to provide reliable CEMS readings. CEMS data availability for control plants largely caught up in later compliance periods. To account for differential reporting, we analyze the CEMS data using several different imputation rules for emissions in plants that do not report pollution readings in a given week. Appendix C describes our treatment of missing data.

We also hired independent environmental labs to manually sample plant emissions concentra-

tions at one point in time before the emissions trading experiment started. These readings were taken with the same protocol GPCB uses to judge compliance, but our data was neither legally actionable nor shared with the regulator.

Trading data.—Our third main data source is the universe of permit bids and trades, from the market operator (NeML). The data include a ledger of permit purchases and sales for all compliance periods, both during the weekly auctions and in the over-the-counter trading period between auctions. We have the complete order book of all bids and offers for permits regardless of whether a bid resulted in a trade. Our data covers all bids, because permit trades could only be registered on the market operator’s platform. We will use this dataset to describe the market and to estimate plant marginal abatement costs.

3.3 Summary Statistics

Survey data and the state of status quo regulation.—The survey data allow us to characterize status quo regulation. Table 1 summarizes plant covariates by treatment arm at baseline. The sample is balanced across a wide range of measures of inputs, outputs, equipment and pollution. Sample plants are large factories with large expenditures on energy and related inputs.¹¹ Average control plant electricity costs are USD 350,000 (Table 1, panel A). The boiler, the plant’s main pollution source, costs USD 112,000 annually to run, excluding fuel expenditures.

The “command” portion of regulation works well, as plants universally have installed some kind of air pollution control device (APCD). Table 1, panel B describes plant APCD holdings. In our sample, 97% of control plants (98% of treatment) have a cyclone installed, 86% (81%) plants have a bag filter installed, 60% (64%) plants have a scrubber installed and 8% (11%) of plants have an electrostatic precipitator. Installation rates are inversely proportional to the cost, and also the efficacy, of abatement equipment. Cyclones are inexpensive but have a low efficacy, reducing PM emissions by 60-90% but $PM_{2.5}$ by only 0-40%. Larger plants with multiple emissions sources must install more expensive APCDs like scrubbers, which remove more than 95% of PM.

¹¹Many plants are, nonetheless, formally classified by the government as “small scale” (71%), based on the reported capital stock at the plant’s establishment.

The “control” portion works less well, as there is widespread non-compliance with pollution standards. Pollution concentrations and mass rates are balanced across treatment arms at baseline (Table 1, panel C). Roughly 30% of plants in both arms have pollution concentrations above the standard. 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 *average* levels of emissions exceed the SPM *maximum* standard of $150 \text{ mg}/\text{Nm}^3$. Appendix Figure B1 shows the distribution of concentrations prior to the market’s launch by treatment arm.

Summary statistics on bidding data.—The cost analysis will use both direct measures of costs from the survey and revealed preference measures from permit bids.

Table 2 shows the evolution of bid counts, prices and quantities. Table rows show statistics for different compliance periods. The columns show the number, quantity and price of bids submitted (not trades cleared). The permit market was active from the start, with 1,525 bids, nearly ten bids per treatment plant, submitted in the first compliance period, at a mean price of INR 12.70 per kg (standard deviation INR 16.65 per kg) (column 6). Both the level and the dispersion of bid prices fell after the first compliance period, up to the Covid lockdown after period six, before rising again when the market reopened from period seven onwards. Plants were active on both the buy and sell sides of the market. The average bid size of 412 kg, across all compliance periods (column 3), can be compared to average emissions of roughly 1,000 kg per plant-month. The volume of bidding activity appears large, relative to the level of emissions, especially as plants were allocated permits totaling 80% of the overall market cap.

Are permit bids a reasonable proxy for abatement costs? We observe mean bids of INR 11 per kg across both sides of the market (column 6), with bids ranging from the floor of INR 5 per kg to INR 45 per kg. To serve as a benchmark, we gathered data from the manufacturers of air pollution control devices on the capital and operating costs of their equipment, which imply variable abatement costs in the range of INR 2 per kg to INR 20 per kg, or sometimes higher. We therefore find that the bidding data in the market overlaps heavily with the wide range of abatement costs implied by engineering estimates under different operating conditions (Appendix Table D1).

4 The Functioning of the Emissions Market in the Treatment

This section provides a descriptive analysis of how the emissions market functioned. First, we describe the evolution of permit prices and trading activity. Second, we examine whether the market's most essential rule, that all plants hold permits for each unit of emissions, was followed. A failure to uphold this rule could cause the market to unravel by undermining the value of permits. Third, we describe how final permit holdings differ from initial allocations, to gauge whether transaction costs impede trade. We find that the market functioned well on all counts, with low permit prices, high compliance and robust trade.

4.1 Permit prices and quantities

Figure 2 shows the weekly time series of permit prices (panel A) and permit quantities traded (panel B). The scattered data points show the mean permit bid each week. The weekly clearing prices are reflected in the solid line in Panel A, which alternates between blue and black to indicate the change of permit vintage with each compliance period. The market rules deliberately reduced price volatility by constraining over-the-counter trades to occur at prices revealed by weekly auctions (Section 2.3). The price floor is indicated by the dashed red line at INR 5 per kg.

Market-clearing prices are generally low, ranging from the price floor to INR 16 per kg depending on the compliance period and week range. Prices were generally lower in the pre-Covid-interruption compliance periods (1-6), 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 plummet during the true-up period, when emissions are known with certainty.¹² Mean bid prices were substantially higher than the market clearing price in the early periods, but this difference declined over time, consistent with market participants learning that the costs of emissions reductions were lower than initially expected.¹³

¹²This price behavior is consistent with uncertainty, prior to the end of the compliance period, as to whether the market would be short or long on permits in aggregate. When the market closes and this uncertainty is resolved, we expect prices to converge to the ceiling or floor, respectively. The market-clearing mechanism of having a single auction after the close of the compliance period may mute this price volatility at the end of each period.

¹³A similar pattern of declining trading prices was observed at the start of the US Acid Rain program market for

Panel B plots daily permit quantities traded as a fraction of the cap per compliance period. The double-sided auction held on the first Tuesday of a compliance period typically causes a spike in quantity. Overall trade volume is significant, with volumes as high as 20% of the monthly cap, or more, on many single days. Trade volumes are higher during the first part of a compliance period as plants buy or sell permits to align permit holdings with expected emissions. As plants' uncertainty about total emissions for the period diminish, toward the end of the period, so do trade volumes.

4.2 Compliance with Market Rules

An existential question for any pollution market is whether participants will respect the mandate to hold permits for all of their emissions. Figure 3 plots the distribution of emissions across plants as a fraction of permit holdings, separately at the end of each of the ten true-up periods that followed, respectively, the ten compliance periods. Any plant with emissions in excess of permit holdings (i.e., more than 100% in the figure) is non-compliant. By contrast, plants with emissions less than their permit holdings (i.e., less than 100%) “left money on the table” by not selling their excess permits to other plants or back to the regulator at the floor price.

There are two findings. The first is that compliance is nearly perfect, defined as emissions during the compliance period being equal or less than permit holdings at the end of the true-up period. There are only a few, scattered non-compliant plants (see periods 1, 3 and 8). Typically, no plants emit more than their permit holdings. This finding should by not be regarded as a foregone conclusion, because in the status quo regime non-compliance has been widespread (Duflo et al., 2018).

We believe that compliance was high because the regulator quickly and credibly established that violations would be penalized. Two plants had emissions that exceeded their permit holdings during the first compliance period. Plant A had emissions of 3928 kg against permit holdings of 3456 kg and Plant B emissions of 4716 kg against permit holdings of a mere 1456 kg. These plants were levied Environmental Damage Compensation (EDC) in accordance with the market sulfur dioxide (Schmalensee et al., 1998).

rules. Plant A paid the EDC and then topped up their environmental bond. Plant B had failed to post the required bond in the first place. The regulator ordered the closure of plant B. Plant B then posted a bond and paid a penalty of INR 652,000, more than ten times what it would have cost to buy permits on the market to cover all plant emissions in the period. The regulator revoked their closure and allowed the plant to reopen after two weeks.

The second finding is that plants leave little money on the table. The mass in the histograms is stacked up at 100%; the vast majority of plants hold permits that almost exactly equal their total emissions at the end of each period. Looking down the first column of distributions, and then down the second, we see that more plants left money on the table in early compliance periods, when market participants had limited experience and the clearing price was relatively low. By later compliance periods almost all plants hold only the permits they need to cover their emissions. The precision of permit holdings suggests that plants understood the incentives for permit trade and that transaction costs in the market were low.

4.3 Permit allocations and plant emissions

The purpose of an emissions market is to take advantage of heterogeneity in abatement costs that might not be easily observed by the regulator. In the Surat market, permits totaling 80% of the cap were allocated *pro rata* to plants based on a plant's total heat output, the regulator's best ex ante measure of emissions capacity. However, to the degree that unobserved heterogeneity in costs exists, we should expect that this allocation rule, or any rule based on observables, would not dictate actual ex post emissions. Rather, plants would trade away from their initial allocations based on their underlying abatement costs (unless, improbably, all plants had the same capacity utilization, rate of emissions per unit heat output and marginal costs).

To test the idea of trade motivated by heterogeneity, Figure 4 plots the distribution of plant emissions as a percentage of their initial permit allocation in each compliance period (rather than final holdings, as in Figure 3). Plants that emit exactly what they were allocated would appear at 100%, while plants that emit twice what they were allocated would appear at 200%. Because only 80% of the total cap is freely allocated in each period, with the rest auctioned, we expect mean

emissions as a percentage of the permit allocation to exceed 100%.

The main finding is that plant emissions are widely dispersed with respect to initial permit allocations in all compliance periods. Most of the mass of the distribution falls between 50% and 200% of the initial allocation, with a relatively modest share between 100% to 125%. This dispersion suggests that the market satisfies two conditions. First, low transaction costs: plants are not constrained by their initial permit allocations, but trade freely to set permit holdings equal to their emissions. Second, unobserved heterogeneity: boiler capacity is a weak proxy for ultimate emissions. Differences in plant capacity utilization, emissions rates, and marginal abatement costs may all contribute to dispersion relative to the capacity-based measure used for permit allocation.

The over-arching finding of this section is that the market functioned well. A large volume of trade at low prices enabled plants to move from their initial permit allocations to permit holdings that met their emissions. Emissions rarely exceed holdings, nor do plants leave money on the table by holding extra permits. The wide spread in emissions relative to initial allocations suggest there is scope for gains from trade based on unobserved heterogeneity in costs.

5 Experimental Results on Pollution Emissions and Costs

This section reports experimental treatment effects on pollution emissions and costs.

5.1 Pollution emissions

Graphical analysis.—Figure 5 shows weekly mean per plant emissions in kilograms per month, from April 2019 to April 2021, by treatment arm. Treatment firms are represented by the solid (blue) line and control firms by the dashed (grey) line. Vertical lines separate market compliance periods. The Covid-19 lockdown, denoted interregnum on the figure, is shaded in light blue. The interregnum divides the sample into early (1 to 6) and late (7 to 10) compliance periods.

As discussed earlier, pollution reporting in the experiment was incomplete. Appendix Figure C1 shows the rates of reporting by treatment arm over time. Initially, reporting is low, and treatment plants report much more than control plants, but this gap narrowed to a few percentage

points by the end of the experimental period. The main pollution series in Figure 5 imputes missing plant emissions using observations for the same plant from the same week or month.

There are two findings from the pollution figure. First, pollution emissions are consistently lower in the treatment group than in the control. By the beginning of compliance period 1, in September, 2019, treatment plants emit roughly 300 kilograms per month less particulate matter than control plants. The difference between treatment and control average emissions is maintained throughout the early compliance periods and into the late periods, after the market reopens. The size of the gap is fairly steady despite marked increases in control plant reporting over time.

The second finding of Figure 5 is that emissions met the cap in all compliance periods. We plot the mean emissions per plant required to meet the cap exactly with red horizontal lines. In later periods, the cap is roughly 1,000 kg (1 metric ton) of SPM per plant-month. All compliance periods have mean treatment emissions, shown by the solid blue line, below this level, sometimes sharply below (around the Diwali holiday, in November, many plants cease operations for a week and emissions fall). The seeming over-compliance in early compliance periods is due to the market replacement rule for missing data (Appendix Table C3).¹⁴ We explore the sensitivity of the finding of lower pollution in the treatment to alternate imputation rules in Appendix C and as part of the regression analysis below.

Regression analysis.—We estimate the size of the treatment effect on pollution with the following specification, at the plant-month level:

$$\log(Pollution_{it}) = \beta_1 Treatment_i + \alpha_t + \varepsilon_{it}. \quad (1)$$

Here $Pollution_{it}$ is the mass of plant-month PM emissions in kg, $Treatment_i$ is an indicator variable equal to one for plants assigned to the emissions market treatment, and α_t are year-month fixed

¹⁴We are plotting mean emissions with imputations at the plant mean in nearby periods; for the purpose of market operations, missing emissions are replaced with a rule that fills in punitively high values, meant to deter non-reporting. With these higher imputations for missing data the cap binds more or less exactly (as implied by Figure 3).

effects.¹⁵ Standard errors are clustered at the plant level.

Table 3 reports the results. The columns differ in the treatment of missing data. Columns 1 to 4 use pollution series without plant-month imputation, so missing observations are dropped. Columns 1 and 2 are unweighted. Columns 3 and 4 are reweighted by the inverse probability of a plant reporting emissions, in order to capture treatment effects on the full sample of plants (DiNardo, Fortin and Lemieux, 1996). We predict the probability of a plant reporting using baseline observable characteristics from before the experiment. Columns 5 to 8 report specifications that impute missing pollution observations using two different imputation rules. The rules are detailed in Annex C. Briefly, Rule A, in columns 5 and 6, imputes a stack missing emissions in a given month at its mean emissions from other times in the experiment. Rule B, in columns 7 and 8, imputes a stack at the monthly mean emissions load of its own treatment group for the same month. Odd-numbered columns are again unweighted and even-numbered columns are reweighted by the inverse probability of reporting.

We find that the market treatment significantly reduced PM emissions. In column 2, with no imputation and no reweighting, the estimated treatment effect on log emissions is -0.193 log points (standard error 0.076 log points). Re-weighting with the inverse reporting probability gives very similar estimates (column 4). The treatment effect on pollution is larger when using either imputation rule than in the raw data (columns 5 to 8 compared to columns 1 to 4). The treatment effect on pollution is -0.282 log points (standard error 0.074 log points) for Rule A and -0.316 log points (standard error 0.057 log points) for Rule B. The higher treatment effects reflect that imputations tend to replace missing control group observations for log particulate emissions load with values higher than the mean among control plants that did report. When control plant emissions are imputed under these rules, therefore, control emissions go up, as does the estimated difference between treatment and control emissions.

We interpret the decline in emissions in the treatment as due to a binding market cap and not a reporting effect due to the greater availability of CEMS data exposing plants to more penalties

¹⁵We favor this simple specification over difference-in-differences specifications because CEMS data reporting was sparse in the period before the market started (Figure C1).

in the status quo regime. The GPCB did not use CEMS data in targeting or penalizing plants in the sample, aside from its role as the basis for the market. Prior to market launch, the CEMS deployment included a randomized control trial designed to estimate the impact of CEMS alone on plant emissions. This trial found that CEMS by itself had no effect on plant pollution emissions (see Appendix C).

5.2 Plant abatement costs

A goal of the experiment is to measure the impact of the market on pollution abatement costs. We present estimates of the market's effect on costs in this subsection, using survey data, and revealed preference estimates, from bidding data, in Sections 6 and 7. With survey data alone, it is difficult to isolate abatement expenditures, because many inputs, like fuel, electricity, labor, and some kinds of capital, may be used for both abatement and non-abatement purposes.¹⁶

Table 4 reports plant-level treatment effect regressions of cost measures on a treatment status indicator. Each column's dependent variable is a different measure of abatement costs. Columns 1 to 5 consider only capital expenditures for air pollution control devices (APCDs). As context, our sample plants have median sales of \$4.2 million and spend roughly \$600,000 on expenditures in the boiler house, the part of the plant responsible for air pollution. Abatement capital is a small part of total plant costs, with an installed APCD stock of \$44,000 (column 1). Columns 6 to 10 zoom out to show all boiler house input costs including capital, labor, electricity, and fuel. All specifications control for the baseline value of the dependent variable.

There is no evidence that the treatment changed abatement costs. The treatment has a small, negative, and statistically insignificant effect on abatement capital of $-\$3,467$ (standard error $\$3,089$). The lack of an effect on abatement capital is consistent with our prior observation that sample plants had high rates of APCD installation at baseline (Table 1). Appendix Table D2 demonstrates that the treatment had no effect on the rate of abatement equipment installation. We interpret these null results as indicating that the existing APCD installation mandate was enforced

¹⁶Analysts have noted similar difficulties with the US Pollution Abatement Costs and Expenditures (PACE) survey, which was discontinued in 2005 (Becker, 2005; Keller and Levinson, 2002).

and so plants already had enough abatement capital to reduce emissions.¹⁷

The treatment also did not cause an increase in a broader measure of input costs. The boiler house is the section of the plant that houses the boiler and other fuel-consuming equipment. Column 6 reports a treatment effect on total boiler house costs of \$11,000 (standard error \$26,000), on a base of \$580,000 in the control group.

These null treatment effects are difficult to interpret. In addition to the problem of measuring abatement costs through surveys, the simple treatment effects comparisons here do not hold constant the level of emissions. The treatment group has both a different regime and a lower level of emissions, and this cut in emissions would be expected, within any regime, to raise costs.

With these limitations in mind, we introduce in Section 6 an alternative, revealed preference approach to inferring abatement costs from plant permit bids. The basic idea is to use information on plants' willingness to buy or sell permits at different prices to estimate plant-specific marginal abatement cost (MAC) functions. With these MAC functions we can compare the cost to plants of reaching *any* given pollution level, under either the market or the status quo regulatory regime.

6 Model of Pollution Abatement

This section lays out how we use plant permit bids to estimate the differences in total abatement costs between the market regulatory regime (treatment) and a command-and-control regime with fixed emissions intensity standards (control).

There are four steps needed to make this cost comparison: (i) we set out a model for plants' pollution abatement decisions that connects permit bids to marginal abatement costs (subsection 6.1); (ii) we calculate total abatement costs, given plant MAC functions, under the market and intensity standard regimes (subsection 6.1); (iii) we estimate plants' MAC functions using variation in plant

¹⁷Our setting, which layers a market on top of an existing equipment mandate, is basically the norm for environmental markets. For example, in the US, the Clean Air Act Amendments (1990) required stationary NO_x sources to install abatement equipment by 1995. In 1999, these sources became part of a regional NO_x cap-and-trade scheme. Schmalensee and Stavins (2013) discusses the interaction of the US SO₂ trading markets with other concurrent policy instruments such as equipment mandates. All carbon markets (e.g., the EU ETS, AB32 and RGGI) coexist with other policy instruments, like renewable purchase obligations, that indirectly regulate carbon.

permit bids at different emissions levels within a compliance period (subsection 6.2); and (iv) we predict the emissions intensity standards the treatment plants would face if they were subject to status quo regulation, as in the control (subsection 6.2). The last step is necessary because permit bids are only observed in the treatment group, so we are only able to recover MACs for the treatment plants. To compare costs across regimes, we need to estimate the fixed standards that treatment plants would have faced in the control group. Since compliance is imperfect, this standard is not a uniform limit, but rather a distribution of emissions rates across plants.

With the complete cost structure of abatement we can then compare abatement costs across plants and regimes. Subsection 6.3 reports estimates of the MAC functions and uses them to illustrate the gains from an emissions market. Section 7 does a formal counterfactual analysis of abatement costs by regime.

6.1 Model specification and abatement costs by regime

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

We let $Z_{it}(E_{it})$ be the level of expenditures as a function of emissions. Assume that $Z' < 0$ and $Z'' > 0$; expenditures are decreasing as a function of emissions but at a rate that decreases in magnitude as emissions grow. Further, there is some high, uncontrolled level of emissions \bar{E}_i such that $Z_{it}(\bar{E}_i) = 0$. The plant spends an added fixed cost Z_{i0} to maintain its abatement capital. We treat this cost as sunk given the finding that abatement capital did not change in the experiment.

Emissions market regulation and total variable abatement costs.—An emissions market is a regulation that sets a market-level cap Q_t on emissions in period t and allows plants to trade

permits so they collectively meet that limit. The regulator allocates permits A_{it} to each plant and may retain or sell the balance. In the Surat market, the allocation rule gave plants permits totaling 80% of the market cap in proportion to their heat output capacity, $A_{it} \propto H_i$. Let P_t be the equilibrium price of permits, known to the plant. Also assume that the plant is small:

Assumption 1. *Plants make abatement decisions assuming $P_t \perp E_{it}$.*

The plant then chooses emissions to minimize the total cost of compliance:

$$\min_{E_{it}} Z_{i0} + Z_{it}(E_{it}) + P_t(E_{it} - A_{it}). \quad (2)$$

The first-order condition for the plant's problem under Assumption 1 is

$$-\frac{\partial Z_{it}(E_{it})}{\partial E_{it}} = MAC(E_{it}) = P_t. \quad (3)$$

This condition is the familiar one that the marginal abatement costs of the plant at the chosen emissions level equal the permit price. This equation has a unique solution for $E_{it}^* = MAC^{-1}(P_t)$ under our assumptions on the $Z(\cdot)$ function.

All plants choose emissions to set their marginal abatement cost equal to the permit price, and therefore the marginal costs of all other plants. When all plants equalize their marginal abatement costs, the market as a whole reduces emissions at the lowest possible aggregate cost. The level of emissions depends on neither the plant's fixed costs of abatement Z_{i0} nor the initial permit allocation A_{it} .

Permit market equilibrium requires that aggregate emissions equal the market cap Q_t . Writing emissions as a function of the price, the equilibrium price is the P_t^* that solves

$$E_t(P_t^*) = \sum_i E_{it}(P_t^*) = Q_t. \quad (4)$$

The equilibrium price is unique given that emissions for each plant monotonically decrease in price. At the equilibrium allocation, the total variable costs of abatement in the market can be

written $Z_t^{ETS} = \sum_i Z_{it}(E_{it}^*)$, with plant emissions given by E_{it}^* .

Command-and-control regulation and total variable abatement costs .—A command-and-control regime is any rule that dictates emissions $\{E_{it}\}$ for each plant, rather than setting a limit across all plants. The current regime, *de jure*, sets a maximal concentration limit on pollution emissions. However, both in the control group and our prior work (Duflo et al., 2018), we observe non-compliance with the intensity standard and fairly wide dispersion in emissions rates, rather than a point mass at the standard \bar{R} .

We therefore estimate the command-and-control intensity standards that plants face based on observed emission rates in the control group. We represent this limit as a plant-specific emissions rate $\bar{R}_{it} = E_{it}/H_i$. We assume that a plant's observed emissions rate is the *de facto* intensity standard they were required to meet.

With this form of regulation, it is straightforward to develop expressions for total emissions and total variable abatement costs. Total status quo emissions $E_t^{CC} = \sum_i H_i \bar{R}_{it}$ depend on the stringency of the plant-period specific intensity standards. Plant abatement costs are then the plant-period abatement cost function evaluated at this emissions level, $Z_{it}(H_i \bar{R}_{it})$. Summing across plants, total variable abatement costs under command-and-control are $Z_t^{CC} = \sum_i Z_{it}(H_i \bar{R}_{it})$.

In contrast to the outcome under an emissions market, there is no reason to expect that costs must be minimized by the command-and-control allocation of emissions. In setting each plant's emissions rate, the regulator may not even seek to minimize compliance costs to plants, but rather to minimize pollution given its own costs and means of regulation. Even if the regulator sought to minimize costs, GPCB is unlikely to have the necessary information about each plant's MAC function. While our past work found that the regulator has some, albeit very noisy, information on pollution (Duflo et al., 2018), we expect marginal abatement costs are more difficult to estimate, since they cannot be observed directly on a plant visit.

6.2 Model Estimation

Abatement cost functions.—The main objects of estimation are plant marginal abatement cost functions, which we estimate using bidding data for treatment plants. We assume the log of marginal abatement cost is isoelastic in emissions

$$\log MAC(E_{it}) = \beta_0 + \beta_1 \log E_{it} + \beta_2 \log H_i + \tilde{\xi}_{it}. \quad (5)$$

The parameter β_1 is the elasticity of marginal abatement costs with respect to emissions. We expect $\beta_1 < 0$ such that marginal abatement costs are decreasing in emissions (increasing in abatement).

There are two difficulties in the estimation of (5). First, the marginal abatement cost of a plant is not observed. Second, emissions are endogenous to abatement costs. Plants with a high abatement cost shock $\tilde{\xi}_{it}$ will choose high levels of emissions E_{it} . We therefore expect ordinary least squares estimates of equation (5) would yield positively-biased estimates of β_1 .

Our approach to estimation is to exploit that firms often made multiple bids within the same period at different levels of emissions (see Table 2, column 2 and Appendix Figure B2). We use this within plant-period variation in bids to estimate the marginal abatement cost function. We estimate a version of (5) as

$$\log b_{itk} = \beta_1 \log E_{itk} + \xi_{it} + \varepsilon_{itk}. \quad (6)$$

The dependent variable is the log of plant i 's bid number k in period t , which we substitute in for marginal abatement costs. The plant-period effects $\xi_{it} = \beta_0 + \beta_2 \log H_i + \tilde{\xi}_{it}$ subsume the effect of heat output. Here E_{itk} is the emissions level (permit holdings) the plant would have if bid k in period t were executed, equivalent to the level of emissions at which the MAC function is being evaluated. For example, if a plant is first allocated $A_{it} = 1,500$ kg of permits, and then with bid $k = 1$ seeks to buy 500 kg of permits, $E_{it,k=1} = 1,500 + 500 = 2,000$.¹⁸

¹⁸Let $\mathcal{K}(k) = \{k' : k' < k, k' \text{ executed}\}$ be the set of bids already executed at the time k is offered. Then generically $E_{itk} = A_{it} + \sum_{k' \in \mathcal{K}(k)} B_{itk'} + B_{itk}$ where buy bids are represented as positive quantities B and sell bids with negative

Equation (6) is our main estimating equation. To consistently estimate the elasticity β_1 of marginal cost with respect to emissions, we assume mean conditional exogeneity

Assumption 2. $\mathbb{E}[\varepsilon_{itk} | E_{itk}, \xi_{it}] = 0$.

This assumption is economically justified if plants form rational, unbiased expectations of their emissions, and therefore marginal costs, at the time of bidding, but have uncertainty about the exact emissions level.¹⁹ We find this assumption plausible in this setting, because plants had CEMS devices, so they could know their emissions, but validated emissions for the market were released at a lag. Moreover, plants may not know at the start exactly how much they will operate in a compliance period, so they likely have some uncertainty about their total emissions level.

This specification allows that marginal abatement costs are unobservably higher for some plants and in some periods, but rules out that innovations in marginal costs of abatement *within* a period are related to emissions choices within a period. Variation in $\log E_{itk}$ comes from different bids that plants submit within the same compliance period at different points along their MAC curve.

We restrict the sample of bids to those from the first half of each compliance period when fitting (6). The rationale is that our model is static and assumes that plants can always make a choice of whether to comply by reducing emissions or buying permits. In the actual market, plants have this choice at the start of each compliance period. As the end of a period approaches, however, within-period plant emissions are largely sunk, so this choice is no longer possible. At the extreme, during the true-up period when the compliance period has ended, emissions are fixed. We probe the rationale for this sample restriction when discussing the estimation results.

To calculate total costs in the market we will also need each plant's total variable cost function. The marginal abatement cost function (5) we assume is consistent with a simple representation of quantities.

¹⁹For example, assume firms anticipate emissions of $\tilde{E}_{itk} = E_{it} v_{itk}$ with $v_{itk} \perp E_{itk}, \xi_{it}$ and $\mathbb{E}[\log v_{itk}] = 0$. Then firms bidding their expected marginal costs yields the above specification (6) with a residual $\varepsilon_{itk} = \beta_1 \log v_{itk}$ based on the forecast error.

total variable abatement costs. We parameterize the abatement cost function as

$$Z_{it}(E_{it}) = e^{\beta_0 + \xi_{it}} H^{\beta_2} \left(\frac{1}{\beta_1 + 1} \right) \left(\bar{E}_i^{\beta_1 + 1} - E_{it}^{\beta_1 + 1} \right), \quad \beta_1 \in (-1, 0). \quad (7)$$

where the parameters are common with (5). Moving from marginal to total variable abatement costs introduces a constant of integration. In (7), the constant \bar{E}_i has a physical interpretation as the high level of uncontrolled emissions for a plant of size $H = 1$ when no variable abatement expenditures are made.

Stringency of command-and-control regulation.—Although by rule all plants face the same concentration standard in the control group, in practice concentrations vary substantially across plants (Appendix Figure B1). We therefore estimate, rather than assume, the stringency of regulation in the command-and-control regime, and then apply the results to the treatment plants. We use five different representations of the status quo to capture the distribution of emissions rates in the command-and-control regime. The regimes differ in whether emissions rates are constant or dispersed across plants and whether they are conditioned on plant characteristics.

The first two regimes we consider are: (i) *constant emissions rate* $R_{it} = \bar{R}$; (ii) *constant emissions rate with error* $\log R_{it} \sim \mathcal{N}(\mu_t, \sigma_t)$, fit separately in each period. These regimes are too simple to represent the status quo, because the data make clear that the emissions rate is declining in heat output capacity. This fact is consistent with a regulatory regime that inspects large plants more often and so imposes greater expected penalties on them for high emissions rates.

We therefore favor regimes where the emissions rate depends on plant heat capacity. We fit the following regression in the control group separately for each compliance period:

$$\log R_{it} = \beta_{0t} + \beta_{1t} \log H_i + \varepsilon_{it}. \quad (8)$$

The remaining three regimes we consider follow this approach: (iii) *capacity-based emissions rate* $R_{it} = \exp(\widehat{\log R_{it}})$; (iv) *capacity-based emissions rate with error* $R_{it} = \exp(\widehat{\log R_{it}} \varepsilon_{it}^s)$ for draws

$\varepsilon_{it}^s \perp \hat{\xi}_{it}$ from the residuals of (8); (v) *capacity-based emissions rate with correlated error*, similar to (iv), but with draws ε_{it}^s that are slightly negatively correlated ($\rho = -0.1$) with marginal abatement cost shocks $\hat{\xi}_{it}$.²⁰ We draw the emissions rate shocks from a log normal distribution fit to the variance of $\hat{\varepsilon}_{it}$ in each period. We include regime (iii) as a basis of comparison, though it will be biased due to the exponentiation of a predicted value fitted in logs.

We use these regimes to set counterfactual emissions rates, our proxy for intensity standards, for the treatment group plants, had they been regulated like control group plants. We then evaluate treatment plants' MAC functions at the simulated emissions rates to calculate the treatment plants' total abatement costs if they had been assigned to the control group.

6.3 Model estimates of marginal abatement costs

Figure 6 illustrates our approach to estimation of marginal abatement cost curves with raw bid data from compliance period 8. The figure plots bid prices b_{itk} against emissions E_{itk} , if each respective bid were to be executed, for two example plants. In panel A, we show the bids for ‘‘Surat Polyfilms,’’ a pseudonym, which offers buy bids (triangles) between INR 5 and INR 10 per kg, with prices that do not vary much by the level of emissions. In panel B, we show the bids for ‘‘Mahadev Textiles,’’ which offers bids that range from INR 5 up to INR 20 per kg. Mahadev Textiles offers higher sell bids (circles) at lower potential emissions levels and lower buy bids (triangles) at higher emissions levels. The plant would have to be paid more, per unit, to abate to a very low level of emissions, consistent with the idea of increasing marginal costs of abatement as emissions fall (abatement increases).

Estimation results.—Table 5 summarizes the bid data statistically by reporting estimates of variants of the specification (6). The data set is at the plant-period-bid level for all bids offered in the treatment group in the first halves of all compliance periods. The first column controls for heat output. The second and third columns add period fixed effects and plant and period fixed

²⁰This implies that high-cost plants will have somewhat lower emissions rates. We introduce this correlation to capture, in a simple way, the observation that the regulator does have some information about plant emissions and targets more polluting plants more aggressively (Duflo et al., 2018).

effects, respectively. The specification in column (4) includes plant-by-period fixed effects that non-parametrically control for all plant-period abatement cost shocks. The coefficient of interest is the elasticity of marginal abatement costs (MAC) with respect to plant emissions. This elasticity is estimated from plants that submit multiple bids in the first half of a period.²¹ The final column (5) additionally allows the elasticity of MAC with respect to emissions to vary by what air pollution control devices a plant has installed at baseline.

Our preferred estimate in Table 5, column 4 is that the elasticity of bid prices (marginal abatement costs) with respect to emissions is -0.609 (standard error 0.087). As emissions increase, the marginal cost of reducing emissions decreases. This estimate satisfies the parameter restriction in (7) that $\beta_1 \in (-1, 0)$, which was not imposed in estimation; hence, as we would expect, the marginal costs of abatement rise as abatement increases (emissions decline).

The pattern of results across columns 1 through 4 shows the importance of using within-period data to estimate this elasticity. Estimates with basic controls or only period fixed effects are positively biased (columns 1 and 2). Even the model that includes plant and period fixed effects (column 3) yields an elasticity less than half as great in magnitude as the preferred estimate. The positive bias of these estimates suggests that plant bid quantities (emissions) are endogenous to cost and likely to be higher when a plant has a high plant-period shock to marginal abatement costs. Appendix Figure D1 estimates the MAC elasticity with different samples of bids, and finds justification for our sample restriction to use bids from only the first half of each period.²²

A main motivation for emissions markets is to take advantage of unobserved differences in marginal abatement costs across plants. In Table 5, column (4), we test for the joint significance of the plant-period fixed effects, for the level of marginal abatement costs, relative to a model with plant-period random effects. We reject the random effects model in favor of our plant-period fixed

²¹The average plant has eight bids in a period. Our data has 8,433 bids, of which 3,120 were submitted in the first half of a period and 2,775 were offered by plants that submitted multiple bids in that time.

²²The rationale for this restriction is that the problem (2) assumes plants can make a choice of whether to comply by reducing emissions or buying permits. As the end of a compliance period approaches, plant emissions within a period are largely sunk, so this choice is no longer possible. Appendix Figure D1 shows that, if we move the sample from early to late within each compliance period, the estimated elasticity of bids with respect to emissions quantity goes from being negative and significant to close to zero. This finding is consistent with the idea that plants can no longer choose emissions flexibly as the end of a compliance period approaches.

effects model (p -value < 0.001 , column 4). In column (5), we also allow the MAC elasticity to differ by what air pollution control devices (APCDs) the plant has installed at baseline. We find that the MAC curve is slightly more elastic (in absolute terms) when plants have only the less expensive APCDs (cyclones and bag filters) installed, but that this difference is not statistically significant (p -value 0.47, column 5). We proceed with column 4 as our preferred specification.

MAC curves and the potential for gains from trade.—Figure 6 plots the fitted values from the column 4 results. In panels A and B, the MAC curves are plotted on the same axes as the bidding data for the two example plants. Some of the strengths and limitations of this approach should be kept in mind. The flexible plant-period effects mean that our model will always fit the mean of log bids and therefore get the level of MAC right. By contrast, it is a strong assumption that there is a single elasticity β_1 across all plants; we judged the alternative of imprecisely estimating separate elasticities for each plant to be inferior. Our fitted model therefore recovers the average elasticity of MAC curves but will overfit (panel A) or underfit (panel B) this elasticity for individual plants and periods.

The figures illustrate the mechanism for gains from trade under the emissions market. Consider the scenario under the status quo, which we represent here, for simplicity, as a regime where all plants have standards set to the average per plant emissions load (1,090 kg, shown by the dashed vertical lines in panels A and B). The two plants have different marginal abatement costs at this common level of emissions load.

Now consider, in the same panels, the scenario under the emissions market. The horizontal dashed line represents the clearing price of INR 10 per kg in week 2 of period 8. Each plant emits to the point where its MAC curve intersects the market-clearing price. The total variable abatement cost of each plant, beneath the MAC curve from the chosen emissions upward, is shaded in pale red. Surat Polyfilms (panel A), in this scenario, would sell permits by reducing emissions from 1,090 kg down to 441 kg. Although the additional abatement, relative to the command-and-control regime, is costly, the plant is paid a price of INR 10 per kg that exceeds its marginal cost. The profit from this additional abatement is given by the shaded blue area above the MAC curve and

below the price. In contrast, Mahadev Textiles (panel B) finds it more costly to abate than Surat Polyfilms does. At the market price, Mahadev emits more than under command-and-control (2,223 kg versus 1,090 kg) and purchases the difference (i.e., 1,133 permits) at a price of INR 10 per kg. Mahadev's profits increase, again by the area in blue, because they find it cheaper to buy permits than it would have been to reduce emissions to 1,090 kg on their own. The opportunity to trade has changed each firm's abatement choices and increased their profits.

This example includes only two plants, whereas we recover heterogeneous MAC curves for every plant and period. Figure 6, panel C summarizes the results in the same period eight data from which the examples in panels A and B were drawn. At emissions levels observed in the market, around 1000 kg per plant-month, there is significant heterogeneity in marginal abatement costs. These heterogeneous curves are the basis of the counterfactual analysis of the gains from trade presented in the next section.

7 Counterfactual Abatement Costs by Regulatory Regime

This section uses the estimated marginal abatement cost functions to compare abatement costs under the command-and-control status quo and the emissions market. Section 7.1 describes how we use the fitted model to calculate counterfactual abatement costs under the two regimes. Section 7.2 shows the fit of our estimated model. Section 7.3 examines how the emissions market affects abatement costs at various counterfactual regulatory stringencies, relative to the status quo. Collecting all of our estimates, Section 7.4 conducts a benefit-cost analysis of the introduction of the emissions market.

7.1 Abatement costs by regime

The treatment involved a bundle of changes in the type (market vs. command-and-control), basis (load vs. concentration) and stringency (170 tons vs. 240 tons, in aggregate) of regulation. In the model we can compare abatement costs in the market and command regimes to isolate the effect of changing only the type of regulation, or of varying the stringency of regulation within

type. To allow this comparison we describe how we aggregate costs in each regime.

Emissions market total variable abatement cost function.—Our aim is to aggregate variable abatement costs across all the plants in the treatment as a function of the emissions cap, $Z_t^{ETS}(Q_t)$. The first step for any given cap Q_t is to solve (4) to find the equilibrium price P_t^* . With the plant’s first order condition (3) and the estimated marginal abatement cost functions, the empirical inverse MAC function is:

$$E_{it}(P_t) = P_t^{1/\hat{\beta}_1} e^{-\hat{\xi}_{it}/\hat{\beta}_1}. \quad (9)$$

With this equation, we solve for plant emissions as a function of the clearing price. With each plant’s emissions, we evaluate their variable abatement costs and sum across plants to find aggregate costs Z_t^{ETS} , as described in Section 6.1. The result of these steps is that we can write aggregate costs as a function of the aggregate emissions cap, $Z_t^{ETS}(Q_t)$.

Command and control total variable abatement cost function.—Here we develop an analogous aggregate cost function under the command-and-control regime. Section 6.2 described how we estimate the command-and-control emissions rates for each treatment plant using (8) at the status quo level of stringency observed in the control group. Take as given this set of fitted emissions rates across plants $\{\hat{R}_{it}\}$, which implies total emissions E_t^{CC} .

Counterfactually, we assume that a differently stringent command-and-control regime would scale up or down all emissions rates by a common factor δ . Then $\delta(Q_t) = Q_t/E_t^{CC}$ for any proposed emissions cap Q_t . We evaluate plant-specific costs at alternate stringencies to calculate aggregate costs $Z_t^{CC}(\delta(Q_t)) = \sum_i Z_{it}(\delta(Q_t)H_i\hat{R}_{it})$.

The idea of this approach is to preserve the dispersion in compliance, as observed in the current regime, while scaling emissions upwards or downwards to meet different possible caps. This assumes that the range of compliance at any new stringency would be the same, in proportional terms, as is observed in the control group. Since plant abatement costs are convex, this approach of evaluating costs as we shift the distribution of emissions rates will produce higher aggregate

abatement costs than would simply evaluating all plants at the new mean emissions rate.

7.2 Model fit

The model delivers a predicted equilibrium price that would result from efficient abatement, given the marginal abatement cost functions. In practice, the equilibrium price will differ from the model's prediction for a variety of reasons, including a failure by plants to minimize costs, transactions costs in permit trade, misspecification or estimation error in the marginal abatement cost functions, and so forth.

With these potential sources of difference in mind, Figure 7 shows the fit of the model to market-clearing prices by compliance period. The dashed black line shows data on mean bids; the dotted black line shows data on mean clearing prices; the solid blue line shows the model simulation of market-clearing prices.

There are two main findings. First, the model's predicted market prices have a very good fit to mean bid prices' level and fluctuations across periods. Bids and simulated prices are relatively high in the first period, fall to about INR 8 – 10 for periods two through six, and then rise in the final four compliance periods to about INR 10 – 12. Second, the model's predicted prices are also similar to the market-clearing price, but the fit is not quite as good. The model predictions for market prices in the first through sixth periods are consistently above actual clearing prices. The predicted and actual clearing prices are much closer in periods seven through ten.

We believe the initial gap in model fit and later convergence may be due to plants' learning about compliance costs. The model is fit to bids offered in the first half of each compliance period at the start of a brand-new market. The model can therefore be thought of as representing plants' expected abatement costs at the beginning of each compliance period, which may have been too high, in the early going. The model of expected abatement costs therefore fits firm bids well in all periods but actual market-clearing prices better during later compliance periods.

We conclude that the model provides a good approximation to actual market outcomes. In the next subsection we apply the model to counterfactual analysis.

7.3 Counterfactual Analysis

Example aggregate abatement cost curves.—Aggregating across all plants, Figure 8 depicts total variable abatement cost curves as a function of aggregate emissions under both the emissions market (lower, blue line) and command-and-control (higher, black line) regimes using the estimates from compliance period eight (as in Figure 6). The vertical dashed line at left depicts emissions at the treatment level of 170 tons, which was the cap from compliance period four onward. An estimated treatment effect of 30% on emissions (Table 3) then implies control emissions of approximately 240 tons, shown by the vertical dashed line at right. On each vertical line we place a filled-in marker at the intersection with the abatement cost curve of each regime.

At any given level of pollution, Figure 8 shows that the emissions market achieves lower variable abatement costs than the command-and-control regime. The model enables a comparative static comparison of the cost savings holding constant emissions levels. The filled black circle represents emissions under the command-and-control regime. The corresponding cost under the emissions market, shown for period 8, is 10% lower (moving down the vertical dashed line). Because total abatement costs are not very elastic with respect to emissions, the cost curves imply that the emissions market, at the same variable abatement costs as in the status quo, would cut total emissions by 48% (moving left along the horizontal dashed line). Alternately, a range of outcomes with both lower emissions and lower costs are available along the arc of the emissions trading cost curve between the horizontal and vertical dashed lines. The treatment bundle of emissions and costs selects one point along this arc.

Perhaps the most striking finding is that abatement costs rise only slowly in response to reductions in pollution, under either regime. For example, under the emissions market, the arc elasticity of total variable abatement costs with respect to total emissions is -0.23 at the status quo level of pollution (240 tons) and -0.19 at the treatment level of pollution (170 tons). These elasticities depend on, but are naturally lower than, the elasticity of *marginal* abatement costs estimated in Table 5, because total abatement costs include inframarginal abatement costs and are therefore less elastic to emissions than marginal costs. The cost estimates suggest that substantial improvements

to Gujarat's air quality are available at relatively small increases in industries' abatement costs. We explore this idea further in Subsection 7.4.

Summary of abatement cost differences by regime.—Table 6 expands on Figure 8 by comparing costs under the two regimes across all compliance periods. We use the same two reference levels of emissions, 170 tons (columns 1 to 3) and 240 tons (columns 4 to 6). Panel A reports on the equilibrium market price and total variable abatement costs under the market. Panel B presents the total variable abatement costs under the command-and-control regime and its percentage difference, relative to costs in the emissions market. The rows under the command-and-control regime differ in how exactly they model the distribution of emissions; see Section 6.2, on stringency in the command-and-control regime, for details.

There are two main findings from Table 6. First, total variable abatement costs are lower under the emissions market than under the command-and-control regime. At the treatment emissions level, 170 tons per month, total variable abatement costs are 12% higher under the status quo (column 3, row B4) than under emissions trading (column 3, row 1). The cost difference between regimes is great enough that costs are 6% lower under the emissions market—with a 30% cut in emissions—than in the command-and-control regime at the status quo emissions level (column 2, panel A versus column 5, row 4).

Second, the cost differences among the alternative representations of the command and control regime in Panel B are small and indeed smaller than the difference in cost between the market and command-and-control regimes. The differences in costs in the command and control regime are due to two forces: (i) heterogeneity in emissions rates interacting with convex abatement costs and (ii) scale effects.²³ We find that abatement costs are 8 to 13% higher under command-and-control at the lower level of emissions (column 3) and 10 to 16% higher at the higher level of emissions,

²³On heterogeneity, command and control regimes that allow idiosyncratic shocks across plants have higher costs than regimes that do not because abatement costs are convex. This convexity pushes up marginal abatement costs for plants that are assigned lower rates of emissions more than it reduces them for plants with higher rates (compare rows B1 and B2). On scale, we find that larger plants tend to have higher marginal abatement costs because the scale efficiencies in marginal abatement costs are outweighed by the higher marginal abatement costs associated with the more stringent emissions standards that large plants face (compare row B1 with rows B3 to B5).

with our preferred estimates lying near the upper end of these ranges. For our preferred estimates, in row B4, the level of costs under the command-and-control regime are 12% and 15% higher, at the respective treatment and control levels of emissions, implying that the market cuts costs by 11% and 14% for these emissions levels.

7.4 Benefit-cost Analysis of Emissions Market Expansion

Our analysis to this point has quantified the costs of using a market to reduce pollution emissions. The low cost of abatement in the market greatly expands the set of feasible policies, and this subsection conducts a benefit-cost analysis for expanding the emission market to all of the 906 industrial plants that burn solid fuel in Surat, at different potential market caps. We compare the benefits of lower ambient concentrations of particulate matter increasing life expectancy against the emission market's operational and abatement costs. All of these comparisons take as given the existing stock of abatement capital in Surat.

Table 7 presents this analysis and Appendix E discusses the inputs in more detail. The columns (1) to (3) give the benefit-cost analysis for emissions reductions of 10%, 30% and 50%. The first two reductions are within the range of our experimental data and the third extrapolates outside that range using the emission market's total abatement cost function.

On the cost side, we include both abatement costs and the capital and operating costs of improved CEMS monitoring (panel A). Improved monitoring is a necessary condition to start an emissions market. Given our administrative data on CEMS, it costs roughly \$5,000 per plant-year to set up a monitoring system. This increase in costs is counterbalanced by an estimated reduction in abatement costs of \$645 per plant-year, which is the abatement cost savings from moving from the status quo regime to a market with emissions 30% lower. On net, then, the new market costs plants roughly \$4,000 per plant year. Scaling this cost to cover all eligible plants in Surat yields an aggregate expense of about \$4 million per year.

On the benefit side, we consider only the health benefit of lower pollution lengthening people's lives. We use the reduction in pollution in the experimental treatment and the best available source apportionment study, Guttikunda, Nishadh and Jawahar (2019), to calculate that a scaled-up market

covering the 906 eligible plants in Surat would reduce ambient PM_{2.5} levels by $8.5 \mu\text{g}/\text{m}^3$, or 10% (panel B), with a per plant cap as in the treatment (column 2). We calculate estimates of the gain in life-years using four different studies, representing the range of estimates of the mortality impacts of exposure to particulate matter air pollution. Panel D illustrates this calculation with a change in life expectancy with respect to PM_{2.5} estimated at air pollution concentrations similar to those in Surat (Ebenstein et al., 2017). We value these gains with an India-specific VSL of \$665,000 which implies a value of a statistical life year (VSLY) of \$9,500 (Nair et al., 2021).

Putting the numbers together, using an emissions market to reduce air pollution from industrial sources has benefits that *greatly* exceed the costs, across a range of targeted emissions reductions. If the 30% reduction achieved in the treatment group were replicated for all polluting industrial plants in the area, the estimates imply a benefit-to-cost ratio that ranges from 26/1, for low estimates of particulate matter damages (panel E, row 4), up to 215/1, for high estimates of particulate matter damages (panel E, row 1). Regardless of the preferred estimate of PM damages, the social benefits of the market exceed costs by at least twenty-five times. Looking across columns 1 to 3 in panel E, the benefit-cost ratio rises steeply with greater reductions in emissions. The increasing ratio is due to the large fixed component of monitoring costs (panel A) and the gently increasing abatement cost curve (Figure 6).

8 Conclusion

This paper evaluates the world's first emissions market for particulate matter, which we designed in collaboration with the Gujarat Pollution Control Board. There are three main findings. First, the market functioned well: permit trade was active, and plants obtained permits to meet their compliance obligations almost perfectly. Second, the emissions market treatment caused a 20% to 30% reduction in particulate matter emissions, relative to the status quo command-and-control regulatory regime. Third, the market reduced abatement costs by 11% to 14%, holding emissions constant. More broadly, we estimate that emissions can be reduced at seemingly small costs in

Gujarat. The benefits of pollution reductions under an emissions market therefore exceed the costs of the market by at least twenty-five times.

In conjunction with prior work, these results suggest that most of the costs of environmental regulation are not due to plants' abatement costs, which are the primary focus of theory. Rather, the fixed costs of monitoring and enforcing regulation appear high in India. On a similar theme, Dufflo et al. (2018) describe and model the high costs of visiting plants, taking observations and levying penalties to induce plants to comply under the status quo regime.

Often state capacity limits the use of sophisticated regulatory instruments. Yet, as this paper demonstrates, state capacity is not a universal constant. While the establishment of the emissions market in Gujarat took years, the results from investments in new monitoring and forms of regulation are extraordinary in terms of reducing both pollution *and* abatement costs. Gujarat is now expanding the emissions market to cover additional plants in Surat and creating a new, separate market for plants in Ahmedabad. Market mechanisms have the potential to transform environmental quality in India and should be more widely considered and adopted as a policy tool.

References

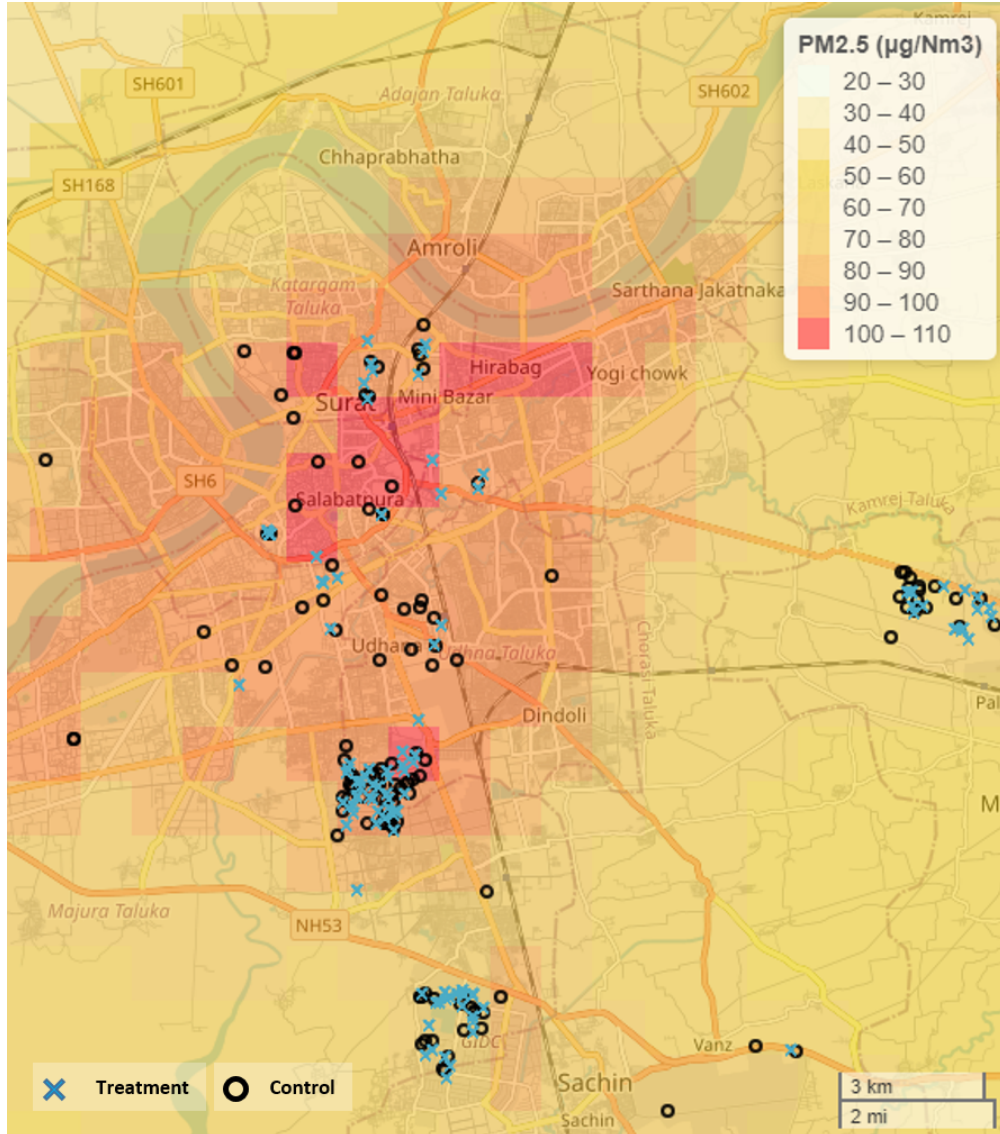
- Apte, Joshua S., Michael Brauer, Aaron J. Cohen, Majid Ezzati, and C. Arden Pope III.** 2018. "Ambient PM_{2.5} Reduces Global and Regional Life Expectancy." *Environmental Science and Technology Letters*, 5(9): 546–551.
- Baylis, Patrick, Michael Greenstone, Kenneth Lee, and Harshil Sahai.** 2023. "Pollution Masks and the Demand for Clean Air: Experimental Evidence from Delhi." *Mimeo*.
- Becker, Randy A.** 2005. "Air pollution abatement costs under the Clean Air Act: evidence from the PACE survey." *Journal of environmental economics and management*, 50(1): 144–169.
- 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.
- Central Pollution Control Board.** 2013. "Specifications and Guidelines for Continuous Emissions Monitoring Systems (CEMS) for PM Measurement with Special Reference to Emission

- Trading Programs.” *CPCB/e-PUBLICATION/2013-14*.
- Coase, Ronald H.** 1960. “The Problem of Social Cost.” *The Journal of Law & Economics*, 3: 1–44.
- Dales, John Harkness.** 1968. *Pollution, property & prices: an essay in policy-making and economics*. Toronto: University Press.
- Davis, Lucas W.** 2008. “The Effect of Driving Restrictions on Air Quality in Mexico City.” *Journal of Political Economy*, 116(1): 38–81.
- DiNardo, John, Nicole M. Fortin, and Thomas Lemieux.** 1996. “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach.” *Econometrica*, 64(5): 1001–1044.
- Dufo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan.** 2010. “Towards an Emissions Trading Scheme for Air Pollutants in India.” *MIT Center for Energy and Environmental Policy Research*, 10(11).
- Dufo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan.** 2013. “Truth-telling by Third-party Auditors and the Response of Polluting Firms: Experimental Evidence from India.” *The Quarterly Journal of Economics*, 128(4): 1499–1545.
- Dufo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan.** 2018. “The Value of Regulatory Discretion: Estimates From Environmental Inspections in India.” *Econometrica*, 86(6): 2123–2160.
- Ebenstein, Avraham, Maoyong Fan, Michael Greenstone, Guojun He, and Maigeng Zhou.** 2017. “New evidence on the impact of sustained exposure to air pollution on life expectancy from China’s Huai River Policy.” *Proceedings of the National Academy of Sciences of the United States of America*, 114(39): 10384–10389.
- Ellerman, A. Denny, Paul L. Joskow, Richard Schmalensee, Juan-Pablo Montero, and Elizabeth M. Bailey.** 2000. *Markets for Clean Air: The U.S. Acid Rain Program*. Cambridge University Press.
- 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, Christa Hasenkopf, and Kenneth Lee.** 2022. “Air Quality Life Index Annual Update.”
- Greenstone, Michael, Guojun He, Ruixue Jia, and Tong Liu.** 2022. “Can Technology Solve the Principal-Agent Problem? Evidence from China’s War on Air Pollution.” *American Economic Review: Insights*, 4(1): 54–70.
- Guttikunda, Sarath K., K. A. Nishadh, and Puja Jawahar.** 2019. “Air pollution knowledge assessments (APnA) for 20 Indian cities.” *Urban Climate*, 27: 124–141.
- He, Guojun, Shaoda Wang, and Bing Zhang.** 2020. “Watering Down Environmental Regulation in China.” *The Quarterly Journal of Economics*, 135(4): 2135–2185.
- Isen, Adam, Maya Rossin-Slater, and W. Reed Walker.** 2017. “Every Breath You Take—Every Dollar You’ll Make: The Long-Term Consequences of the Clean Air Act of 1970.” *Journal of Political Economy*, 125(3): 848–902.

- Ito, Koichiro, and Shuang Zhang.** 2020. “Willingness to Pay for Clean Air: Evidence from Air Purifier Markets in China.” *Journal of Political Economy*, 128(5): 1627–1672.
- Jack, B. Kelsey, Seema Jayachandran, Namrata Kala, and Rohini Pande.** 2022. “Money (Not) to Burn: Payments for Ecosystem Services to Reduce Crop Residue Burning.” National Bureau of Economic Research Working Paper 30690.
- Jayachandran, Seema.** 2022. “How Economic Development Influences the Environment.” *Annual Review of Economics*, 14(1): 229–252.
- Keller, Wolfgang, and Arik Levinson.** 2002. “Pollution Abatement Costs and Foreign Direct Investment Inflows to U.S. States.” *Review of Economics and Statistics*, 84(4): 691–703.
- 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.** 2016. “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*, 10(1): 129–148.
- 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.
- Nair, Moorthy, Hemant Bherwani, Shahid Mirza, Saima Anjum, and Rakesh Kumar.** 2021. “Valuing burden of premature mortality attributable to air pollution in major million-plus non-attainment cities of India.” *Scientific Reports*, 11: 1–15.
- Oliva, Paulina.** 2015. “Environmental Regulations and Corruption: Automobile Emissions in Mexico City.” *Journal of Political Economy*, 123(3): 686–724.
- Pande, Rohini, Michael Greenstone, Janhavi Nilekani, Nicholas Ryan, Anant Sudarshan, and Anish Sugathan.** 2015. “Lower Pollution, Longer Lives: Life Expectancy Gains if India Reduced Particulate Matter Pollution.” *Economic and Political Weekly*, 50: 40–46.
- Piette, Lauren.** 2018. “Improving environmental enforcement in India.” *Mimeo, University of Chicago Law School*.
- Pope, C. Arden III, Majid Ezzati, and Douglas W. Dockery.** 2009. “Fine-Particulate Air Pollution and Life Expectancy in the United States.” *New England Journal of Medicine*, 360(4): 376–386.
- Schmalensee, Richard, and Robert N. Stavins.** 2013. “The SO₂ Allowance Trading System: The Ironic History of a Grand Policy Experiment.” *Journal of Economic Perspectives*, 27(1): 103–22.
- Schmalensee, Richard, Paul L. Joskow, A. Denny Ellerman, Juan Pablo Montero, and Elizabeth M. Bailey.** 1998. “An Interim Evaluation of Sulfur Dioxide Emissions Trading.” *Journal of Economic Perspectives*, 12(3): 53–68.
- Stavins, Robert N.** 2003. “Experience with Market-Based Environmental Policy Instruments.” Elsevier Handbook of Environmental Economics.
- Sudarshan, Anant.** 2023. “Monitoring Industrial Pollution in India.”
- Zhou, Xuehua, Zhaoyu Cao, Yujie Ma, Linpeng Wang, Ruidong Wu, and Wenxing Wang.** 2016. “Concentrations, correlations and chemical species of PM_{2.5}/PM₁₀ based on published data in China: potential implications for the revised particulate standard.” *Chemosphere*, 144: 518–526.
- Zou, Eric Yongchen.** 2021. “Unwatched Pollution: The Effect of Intermittent Monitoring on Air Quality.” *American Economic Review*, 111(7): 2101–26.


9 Figures

Figure 1: Ambient Pollution Levels and the Location of Plants in Surat



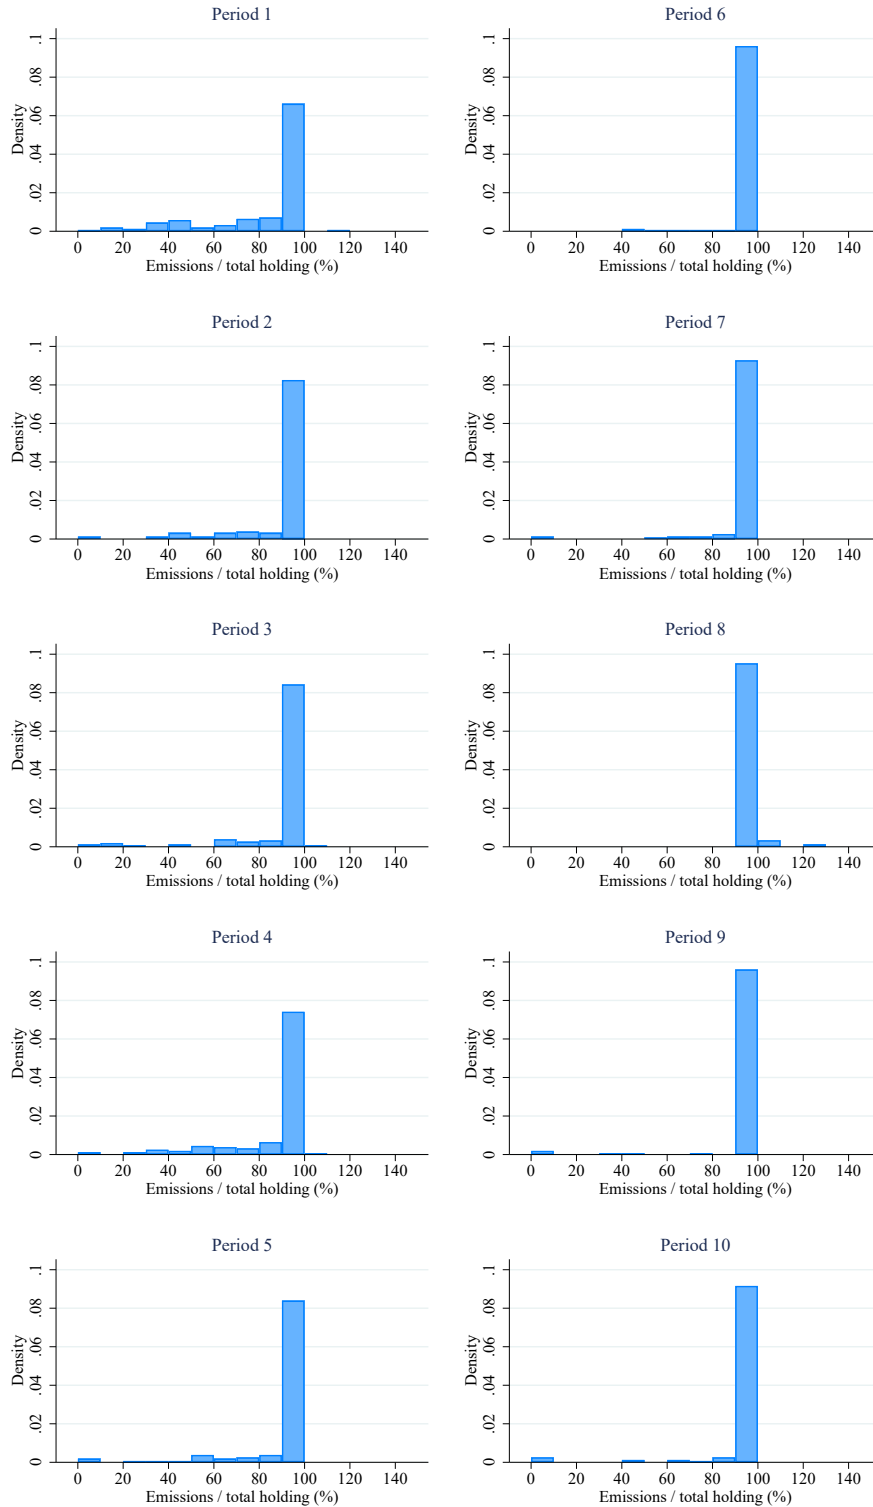
The figure shows ambient PM_{2.5} $\mu\text{g}/\text{m}^3$ concentrations in Surat, Gujarat averaged over the year 2018, overlaid with the locations of sample plants. The ambient pollution data is from Guttikunda, Nishadh and Jawahar (2019). As a basis for comparison, India's National Ambient Air Quality Standard for PM_{2.5} is $40 \mu\text{g}/\text{m}^3$ and the WHO standard is $5 \mu\text{g}/\text{m}^3$. The plant locations are geolocations from our plant survey. Treatment plants are represented by \times markers and control plants by \circ circles.

Figure 2: Permit Prices and Quantities by Compliance Period



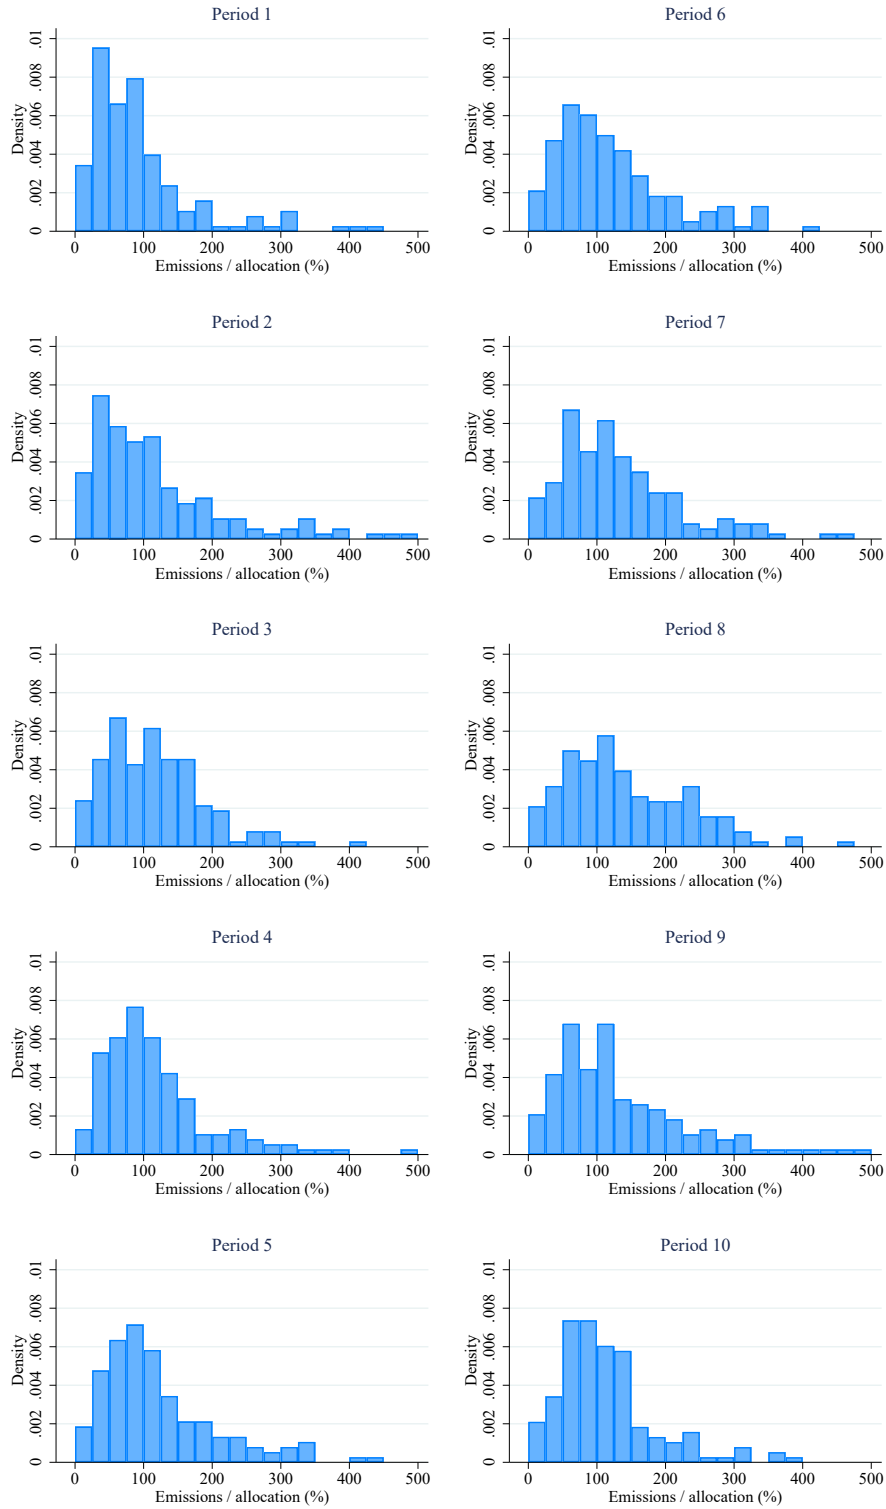
This figure shows weekly permit prices (panel A) and quantities (panel B) from September 2019 to April 2021. In panel A, the scattered points are the mean bid prices (both sale and purchase) and the solid line the market-clearing price. Since permits of different vintages, from two consecutive compliance periods, are traded simultaneously on some days, the market-clearing price line alternates between black and blue colors to differentiate them. The dashed red horizontal line shows the price floor at Rs 5 per kg. In panel B, quantities are expressed as a percentage of the period emissions cap. The large spike near the start of each compliance period is the weekly auction held on the first Tuesday of the compliance period.

Figure 3: Distribution of Emissions over Final Permit Holdings by Compliance Period




This figure plots the distributions of $(\text{emissions} / \text{final permit holdings} \times 100\%)$ across treated plants ($N = 156$) by compliance period, truncated at the 99.5th percentile. Final permit holdings are the total number of permits a plant held at the end of the true-up period after each compliance period. Emissions data and permit holdings are from the administrative records of the market operator. Emissions are the validated emissions for each plant, which include any imputed emissions filled-in for periods of missing data. These validated emissions are used to determine compliance.

Figure 4: Distribution of Emissions over Initial Permit Allocation by Compliance Period



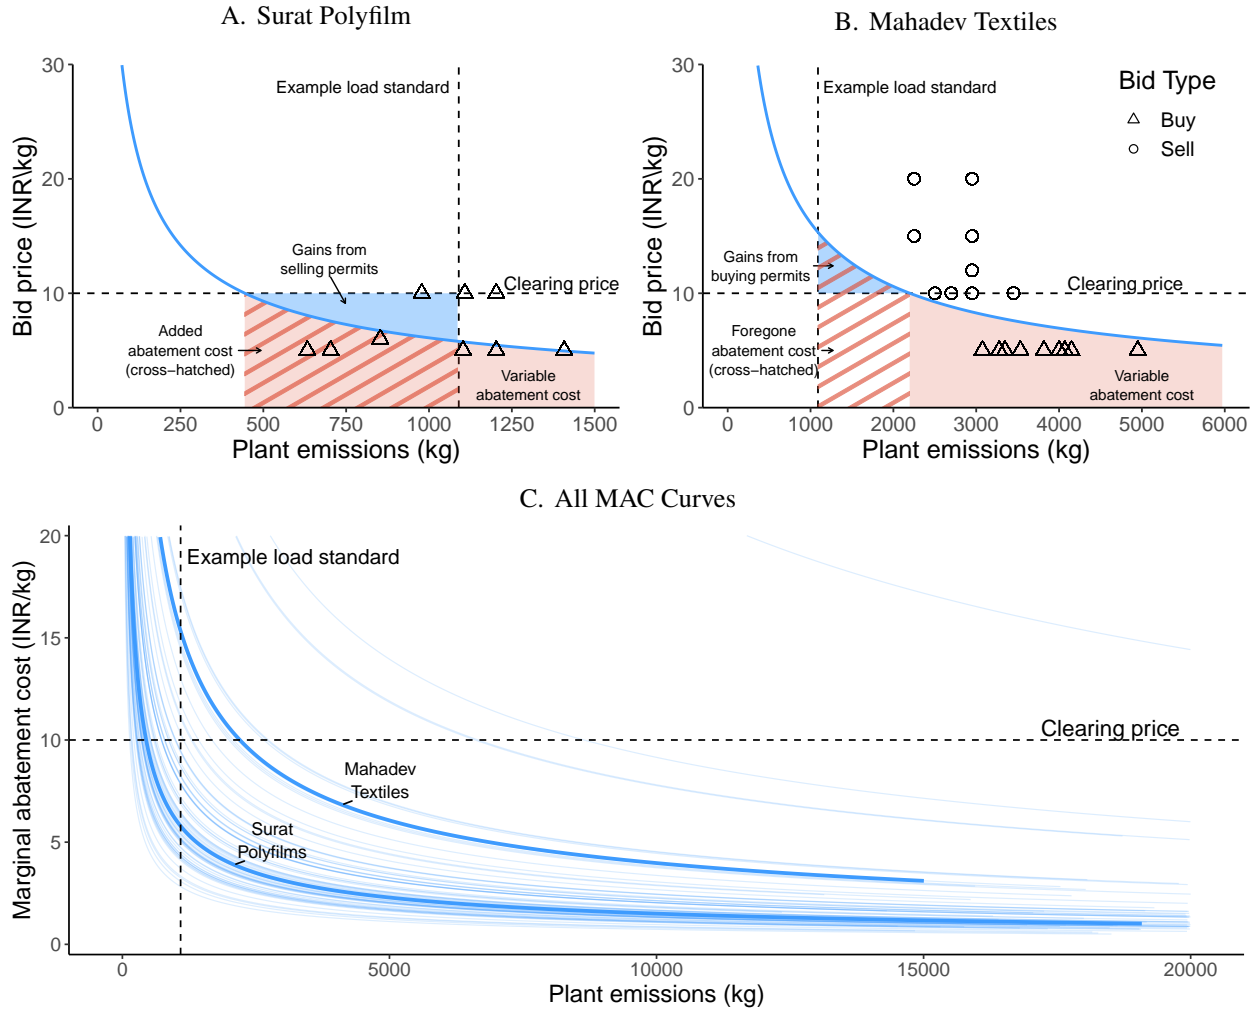
This figure plots the distributions of $(\text{emissions} / \text{initial permit allocation} \times 100\%)$ across treated plants ($N = 156$) by compliance period, truncated at the 97.5th percentile. Emissions data and permit holdings are from the administrative records of the market operator. Emissions are the validated emissions for each plant, which include any imputed emissions filled-in for periods of missing data. These validated emissions are used to determine compliance.

Figure 5: PM Emissions by Treatment Status




The figure shows the weekly mean plant PM emissions in kilograms (calculated at a monthly rate equivalent) from April 2019 to March 2021 by treatment status. The treatment group is represented by the solid (blue) line and the control group by the dashed (grey) line. The grey regions mark the ten compliance periods in the emissions market. The light blue regions mark interregnum periods when the emissions market was closed. The horizontal (red) lines denote the market cap for each period expressed per plant-month. 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. Pollution reporting over this period was incomplete and rising from early to late compliance periods (see Appendix Figure C1). Missing pollution readings are imputed within a stack-week and then within a stack-month (Appendix C.1). The sample consists of 292 plants that had at least one day of PM data from CEMS devices during the ETS experiment.

Figure 6: Examples of Marginal Abatement Cost Curves




The figure illustrates the estimation of marginal abatement cost curves from raw permit bidding data. The top two panels show bids placed by two plants, “Surat Polyfilm” (panel A) and “Mahadev Textiles” (panel B), both pseudonyms, during the first two weeks of compliance period eight. The scattered points mark bids by plotting the bid price (b_{itk}) against the emissions level that would result for each plant if each given bid was executed (E_{itk}). We overlay marginal abatement cost curves of the form (5), as fit in Table 5, column 4, on top of each scatter plot. The horizontal dashed line shows the market-clearing price in week 2 of period 8. The vertical dashed line gives a hypothetical load standard set at the average emissions per plant-month. Panel C shows the same fitted marginal abatement cost curves for all plants that bid in period 8. The domain of each curve extends upward to the uncontrolled emissions level for each plant.

Figure 7: Model Fit to Market-Clearing Prices



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

Figure 8: Variable abatement costs by regime



The figure shows the total (not marginal) variable abatement costs by regulatory regime as estimated for compliance period 8. The dotted (black) curve shows the total variable abatement cost curve under command-and-control and the solid (blue) curve under the emissions market. The command-and-control regime uses a capacity-based emissions rate with error to set emissions targets for each plant, as described in Section 6. The emissions market regime sets an emissions cap at each level of emissions on the horizontal axis. The dashed vertical lines show the approximate emissions levels in the treatment and control groups. The control costs are therefore represented by the upper-right shaded circle and the treatment costs by the lower-left shaded circle.

10 Tables

Table 1: Balance of plant characteristics by treatment status

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

	[0.43]	[0.37]	(0.047)
Above regulatory standard at ETS baseline (=1)	0.34	0.28	0.054
	[0.47]	[0.45]	(0.055)
Number of plants	156	136	

This table shows differences in plant scale (panel A), plant abatement and investment costs (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 same balance table in the full survey sample). In panel B, cyclone, bag filter, scrubber, and electrostatic precipitator (ESP) are different air pollution control devices (APCDs). Some plants did not respond to some questions in the survey and so certain variable rows have fewer observations than the full sample size. 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 2: Summary of permit bid data

Period	Count (1)	Bids per plant (2)	Bid quantities (kg)			Bid prices (INR/kg)		
			All bids (3)	Buy bids (4)	Sell bids (5)	All bids (6)	Buy bids (7)	Sell bids (8)
1	1,525	9.8 (10.4)	557 (1,172)	428 (604)	620 (1,361)	12.70 (16.65)	10.03 (14.47)	14.00 (17.47)
2	600	3.9 (4.1)	474 (747)	475 (535)	474 (810)	7.90 (10.72)	6.39 (9.16)	8.43 (11.18)
3	1,084	7.0 (6.7)	329 (533)	344 (600)	323 (502)	10.89 (11.36)	9.02 (12.24)	11.67 (10.89)
4	806	5.2 (6.1)	323 (551)	332 (483)	319 (575)	10.88 (9.37)	7.90 (6.77)	12.03 (9.96)
5	767	4.9 (7.0)	376 (515)	449 (559)	350 (496)	9.60 (10.80)	6.49 (1.74)	10.72 (12.36)
6	296	1.9 (3.2)	463 (558)	533 (559)	426 (556)	7.50 (6.33)	5.84 (3.11)	8.38 (7.34)
7	646	4.1 (4.5)	400 (533)	468 (580)	325 (466)	12.76 (6.55)	10.29 (4.81)	15.46 (7.11)
8	783	5.0 (6.3)	418 (588)	501 (671)	249 (298)	13.51 (13.92)	12.56 (16.53)	15.47 (4.96)
9	962	6.2 (9.7)	353 (423)	397 (458)	257 (314)	11.41 (6.65)	9.89 (7.00)	14.78 (4.15)
10	964	6.2 (8.8)	383 (533)	428 (532)	336 (531)	11.18 (9.10)	8.40 (5.65)	14.04 (10.92)
Total	8,433	54.1 (51.0)	412 (708)	430 (565)	399 (795)	11.25 (11.56)	9.47 (10.50)	12.52 (12.10)

The table shows summary statistics on plant permits bids across all ten compliance periods. The source of the data is the market operator NeML. Each row shows statistics for a separate compliance period. Each cell has the mean with the standard deviation below in parentheses. The columns show, respectively: (1) the total number of bids in each period, (2) the mean number of bids placed per plant ($N = 156$), (3) - (5) mean quantities for all bids, buy bids and sell bids, (6) - (8) mean prices for all bids, buy bids and sell bids.

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** (0.076)	-0.193** (0.076)	-0.177** (0.075)	-0.194** (0.075)	-0.282*** (0.074)	-0.282*** (0.075)	-0.316*** (0.057)	-0.316*** (0.057)
Month FE		Yes		Yes		Yes		Yes
Imputation rule					Rule A	Rule A	Rule B	Rule B
Reweighted			Yes	Yes				
Mean dep. var (control)	6.67	6.67	6.66	6.66	6.80	6.80	6.88	6.88
R ²	0.13	0.17	0.14	0.17	0.18	0.22	0.16	0.25
Plants	292	292	292	292	292	292	292	292
Observations	3235	3235	3235	3235	3796	3796	3796	3796

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 C.1. Columns 5 and 6 impute data with Imputation Rule A: *Stack-Experiment*. Under this rule, 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). Columns 7 and 8 impute data with Imputation Rule B: *Treatment-Month*. Under this rule, 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. 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 for the treatment status in a prior experiment that randomized CEMS installation timing. 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 abatement costs in survey data

	Abatement capital costs (\$1000s)					Boiler house input costs (\$1000s)					
	All (1)	Cyclone (2)	Bag (3)	Scrubber (4)	ESP (5)	Total (6)	Capital (7)	Labor (8)	Electricity (9)	Fuel (10)	Materials (11)
ETS Treatment (=1)	-3.467 (3.089)	0.602** (0.266)	0.530* (0.318)	-0.222 (0.407)	-4.281 (3.344)	11.26 (26.31)	-7.178 (19.05)	1.561 (3.332)	25.21* (13.53)	26.87* (15.35)	-0.142 (0.596)
R ²	0.90	0.85	0.83	0.84	0.89	0.93	0.63	0.05	0.65	0.98	0.19
Control mean	44.04	7.80	9.85	9.69	16.70	578.48	190.88	47.86	162.13	299.50	4.33
Plants	276	276	276	276	276	185	218	262	247	225	283

This table reports the effects of treatment assignment on the capital cost of APCDs (columns 1-5) and boiler house input costs (columns 6-11). In columns 1-5, the abatement capital cost is the product of the number of abatement devices at a plant and the industry-standard cost for that device for the plant's given boiler house capacity. In columns 6-11, specifications use our best estimates for boiler house costs from the endline survey (FY 2019-20). All specifications control for a corresponding baseline value (FY 2017-18) but in some cases the components of the input cost aggregate differ slightly within a category between the baseline and endline survey. Electricity costs are only reported at the plant level so are not only for the boiler house. Robust standard errors are given in parentheses with statistical significance indicated by * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 5: Elasticity of marginal cost with respect to emissions

	log(Bid price)				
	(1)	(2)	(3)	(4)	(5)
log(Emissions as bid)	-0.100 (0.061)	-0.143** (0.063)	-0.269*** (0.084)	-0.609*** (0.087)	
log(Emissions as bid) × cyclone / bag filter					-0.707*** (0.169)
log(Emissions as bid) × scrubber / ESP					-0.566*** (0.095)
log(Plant total heat output)	0.087** (0.041)	0.138*** (0.048)			
Period FE		Yes	Yes		
Plant FE			Yes		
Plant × Period FE				Yes	Yes
p-val: H_0 : No unobserved heterogeneity				0.000	0.000
p-val: H_0 : No observed heterogeneity					0.468
R ²	0.01	0.07	0.26	0.46	0.46
Plants	146	146	138	127	123
Observations	3120	3120	3112	2775	2753

This table reports the results of regressing log(bid price) on log(emissions as bid). Emissions as bid is defined as the permit holdings that will result if the bid is executed. We run regressions using bids placed in the first half of a compliance period. We include compliance period fixed effects in columns 2 and 3, plant fixed effects in column 3, and plant × period fixed effects in columns 4 and 5. In column 5, the interacted variables “cyclone/bag filter” and “scrubber/ESP” are indicators of the “maximal” (most effective) abatement technology. If a plant has only cyclones or bag filters, then cyclone/bag filter = 1 and scrubber/ESP = 0. If a plant has scrubbers or ESPs, then scrubber/ESP = 1 and cyclone/bag filter = 0. The footer of the table reports p -values for two tests of heterogeneity in marginal abatement costs. The first p -value is for a Hausman test comparing the plant-by-period fixed effects model against a model with plant-by-period random effects instead. The second p -value is for a test that the coefficient of log(Emissions as bid) × cyclone / bag filter is equal to that of log(Emissions as bid) × scrubber / ESP. 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 6: Variable abatement costs under alternative regulatory regimes

	Emissions = 170 tons			Emissions = 240 tons		
	Price (INR/kg) (1)	Total Variable Abatement Cost (INR m) (2)	% Δ Cost Command over ETS (%) (3)	Price (INR/kg) (4)	Total Variable Abatement Cost (INR m) (5)	% Δ Cost Command over ETS (%) (6)
A. Emissions market	12.23	10.08	–	9.91	9.31	–
B. Command-and-control						
1. Constant emissions rate		10.89	8.0%		10.24	10.0%
2. Constant emissions rate, with error		11.23	11.4%		10.62	14.1%
3. Capacity-based rate		10.91	8.2%		10.27	10.3%
4. Capacity-based rate, with error		11.27	11.8%		10.67	14.6%
5. Capacity-based rate, correlated error		11.39	13.0%		10.80	16.0%

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

Table 7: Benefit-Cost Analysis of Scaled-up ETS in Surat

	Emissions reduction			Units	Source
	10%	30%	50%		
	(1)	(2)	(3)	(4)	(5)
<i>Panel A. Annual costs from Scale-up of Emissions Trading</i>					
1. Monitoring costs (per plant)		5,000		\$/year	Author estimates
2. Abatement cost Δ (per plant)	-1,256	-645	147	\$/year	Author estimates
3. Total costs	3.36	3.91	4.62	\$m/year	= (A1 + A2) \times 906
<i>Panel B. Reduction in pollution</i>					
1. Ambient PM _{2.5} conc.		88.5		$\mu\text{g}/\text{m}^3$	Guttikunda et al. (2019)
2. Industry share		0.32			Guttikunda et al. (2019)
3. Reduction in industry PM _{2.5}	10	30	50	%	
4. Reduction in ambient PM _{2.5}	2.83	8.50	14.16	$\mu\text{g}/\text{m}^3$	= B1 \times B2 \times B3
<i>Panel C. Population and value of statistical life</i>					
1. Population		7.5		m	World Pop. Review
2. Life expectancy		70		years	World Bank
3. Value of statistical life		665,000		\$	Nair et al. (2021)
4. Value of 1 year of life		9,500		\$/year	= C3 / C2
<i>Panel D. Gain in life-years and valuation of benefits</i>					
1. Mortality impact		0.098		years/ $(\mu\text{g}/\text{m}^3)$	Ebenstein et al. (2017)
2. Gain in life expectancy	0.278	0.833	1.388	years	= D1 \times B4
3. Per year of ETS	0.004	0.012	0.020	years	= D2 / C2
4. Total gain in life-years	29,736	89,208	148,680	years	= D3 \times C1
5. Value of gain in life-years	282	847	1,412	\$m	= D4 \times C4
<i>Panel E. Benefit-cost ratio</i>					
1. Ebenstein et al. (2017)	83:1	215:1	303:1		= D5 / A3
2. Pope et al. (2009)	52:1	134:1	189:1		
3. Correia et al. (2013)	30:1	77:1	108:1		
4. Apte et al. (2018)	10:1	26:1	37:1		

The table presents the benefit-cost analysis of extending the ETS to the entirety of Surat. We compare the private costs of introducing an emissions market (monitoring and changes in abatement cost, in panel A) to the social benefits of cleaner air (gains in life-years, calculated and monetized across panels B to D). There are 906 plants according to the GPCB consent records. Abatement cost savings are based on capacity-rate estimate (with error) from Table 6. World Population Review estimates for the Surat population are from 2021. The number of plants is based on 2022 GPCB consent records. The annualized CEMS costs are based on an assumed system cost of \$ 12,000 with a 4-year equipment life and no discounting. This equipment life describes the realized experience of some plants in our sample but is lower than typical manufacturer claims. On the benefits side, we assume that the reduction in ambient pollution comes solely from primary particles. Ebenstein et al. (2017) estimate mortality effects of pollution based on PM₁₀ concentrations. We convert their findings using a 0.65 PM_{2.5}-to-PM₁₀ ratio (Zhou et al., 2016). Columns 1 to 3 show the benefit-cost analysis for reductions in pollution of 10%, 30% and 50%. The cost of abatement is calculated using our model at each emissions level. The benefit of emissions reductions are assumed to be linear and given by the estimate from each respective study in panel E.

Online Appendix

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

Michael Greenstone, Rohini Pande, Nicholas Ryan and Anant Sudarshan

This online appendix contains five different parts. Appendix A gives more information on the experimental design. Appendix B describes the data sources and cleaning. Appendix C covers the Continuous Emissions Monitoring System (CEMS) data on pollution, specifically, including imputation rules for missing pollution data. Appendix D gives additional empirical results to support those in the main text. Appendix E provides our benefit-cost analysis.

A Appendix: Experimental Design

This section gives more information about the experimental design. Table A1 gives the timeline of the experimental intervention and data collection. Table A2 describes the duration and market cap for each compliance period of the emissions market. Table A3 describes attrition in the sample by treatment arm and with respect to each source of data. Table A4 duplicates the balance table in the main text but without the sample restriction to plants that report CEMS data.

Table A1: Intervention timeline

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

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 A2: 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

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 varies across compliance periods, due to the 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 the treatment arm. The market was closed during Interregnum-I due to the COVID-19 pandemic and during Interregnum-II following the Divali festival.

Table A3: Sample determination and attrition by treatment status

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

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 that received assignment while extinct or shutdown, yield 322 operational plants with treatment assignment at baseline (row 3). Four of these 322 operational-at-baseline plants were officially removed from the ETS sample by GPCB after the treatment assignment (row 4). Three of the removed plants (2 in control, 1 in treatment) are seasonal sugar cooperatives, operational for only four months of the year; the fourth treatment plant is a particle-board producing plant which uses bagasse, rather than coal, as fuel. Of the 318 in-sample plants, 13 are known to have been incompletely treated by the intervention, due to temporary financial closure before or after the treatment assignment was done (row 6). The 304 plants surveyed at baseline are distinct from the 304 plants manually sampled, and are therefore reported separately (rows 8, 9). This paper reports experimental results from the sample of 292 plants reported at least one day of CEMS data from April 16, 2019 to April 3rd, 2021 (row 11). Of the 162 in-sample plants in the treatment group, 153 plants have market trading data (row 13).

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

	Treatment	Control	Difference
<i>Panel A: Plant Measures</i>			
Total electricity cost (1,000 USD)	456.2	389.1	67.1
	[853.1]	[660.7]	(89.6)
Log(plant total heat output)	15.6	15.5	0.085
	[0.61]	[0.59]	(0.067)
Size as recorded on environment consent (1 to 3)	1.36	1.40	-0.038

	[0.63]	[0.65]	(0.073)
Small-scale (size=1)	0.72	0.69	0.033
	[0.45]	[0.47]	(0.053)
Large-scale (size=3)	0.083	0.088	-0.0056
	[0.28]	[0.28]	(0.032)
Number of stacks	1.08	1.05	0.035
	[0.41]	[0.21]	(0.037)
Textiles sector (=1)	0.85	0.85	-0.0032
	[0.36]	[0.36]	(0.041)

Panel B: Plant Abatement and Investment Cost

Boiler house employment	36.8	31.7	5.13
	[32.5]	[30.0]	(3.59)
Boiler house capital expenditure (1,000 USD)	198.3	164.2	34.0
	[398.6]	[190.9]	(36.7)
Boiler house operating cost (1,000 USD)	138.1	111.0	27.1
	[202.6]	[84.9]	(17.6)
APCD: Cyclone present	0.98	0.97	0.0081
	[0.14]	[0.16]	(0.017)
APCD: Bag filter present	0.80	0.86	-0.055
	[0.40]	[0.35]	(0.043)
APCD: Scrubber present	0.64	0.61	0.032
	[0.48]	[0.49]	(0.056)
APCD: ESP present	0.11	0.082	0.033
	[0.32]	[0.27]	(0.034)

Panel C: Plant Pollution Measures

Plant total PM mass rate (kg/hr)	3.62	3.51	0.11
	[4.86]	[3.76]	(0.50)
Plant mean PM concentration (mg/Nm ³)	177.9	168.5	9.37
	[153.6]	[151.5]	(17.5)
Plant mean Ringelmann score (1 to 5)	1.36	1.35	0.0090
	[0.42]	[0.37]	(0.045)
Above regulatory standard at ETS baseline (=1)	0.33	0.28	0.052
	[0.47]	[0.45]	(0.053)
Number of plants	162	156	

This table shows differences in plant measures (panel A), plant abatement and investment cost (panel B), and plant pollution (panel C) between the treatment and control groups of plants in the baseline survey conducted from December 2018 to January 2019. This sample consists of 318 plants in the ETS experiment. In panel B, cyclone, bag filter, scrubber, and electrostatic precipitator (ESP) are different devices used to reduce emissions. Some plants did not respond to some questions in the survey. For the control group, the numbers of observations are 137 for boiler house capital expenditure, 141 for gross sales revenue, 148 for the 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 the 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. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

B Appendix: Data

This Appendix B discusses data from our plant survey and from administrative records on permit trade. The following Appendix C covers data from Continuous Emissions Monitoring Systems (CEMS).

B.1 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.

As part of the technical survey environmental labs collected samples from the stack(s) (i.e., chimney) attached to the boiler and/or thermopack to measure the PM concentration and PM mass rate. Most plants have only a single stack (Table B1). 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.


Figure B1 shows the distribution of emissions concentrations in the baseline survey by treatment arm. The red vertical lines at $150 \text{ mg}/\text{Nm}^3$ indicate the regulatory standard. Many plants are out of compliance with the standard, sometimes by a factor of two or more. Table 1, panel C shows the plant's mean PM concentration from sampling and an indicator for non-compliance. About 30% of plants are out of compliance at baseline.

Table B1: Distribution of the 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

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

Figure B1: Distribution of Pollution Before the Experiment



This figure shows the distributions of the plant PM concentration by treatment status as measured by manual isokinetic stack sampling at the ETS baseline (December 2018 to January 2019). One PM sample was collected from each industrial stack by a third-party laboratory. The histograms are truncated at the 95th percentile (520 mg/Nm³). The red, vertical lines indicate the regulatory concentration standard of 150 mg/Nm³. At the ETS baseline, 28% of sampled plants in the control group and 34% of sampled plants in the treatment group had readings above this standard.

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 density of smoke as it appears to the naked eye. 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 regulatory purpose.

In Table 1 variables in panel A are from the general section of the ETS baseline survey, and those in Panel B are from the technical section. In panel B, cyclones, bag filters, scrubbers, and electrostatic precipitators (ESPs) are air pollution control devices (APCDs) used to abate PM emissions. In panel C, the plant's total PM mass rate is the sum of the plant's stacks' PM mass rates measured from stack sampling. 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.

B.2 Trading data


The paper uses administrative data on permit bids and offers from NeML, the market operator. Table B2 shows summary statistics on permit bids (panel A) and executed trades (panel B). Figure B2 shows the distribution of the number of bids placed per plant in each compliance period.

Table B2: Trading data summary statistics

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

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 B2: Distribution of number of bids placed per plant by compliance period



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

C Appendix: Pollution Monitoring

C.1 Measuring plant emissions

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).

Imputation for outliers and missing data.—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 ($N = 222,890$).

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

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

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 C1.

Table C1: 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 observation 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 ($N = 7,008$).

C.2 Treatment effect on pollution under alternate imputation rules


Figure 5 shows the time series of pollution by treatment status. Here we show the same series under alternate imputation rules for missing data. Figure C1 shows the data availability of CEMS data on pollution by treatment arm. Figure C2 replicates Figure 5, from the main text, but with alternate imputation rules for missing data. Table C2 summarizes the level of log PM emissions by imputation rule and Figure C3 compares the distribution of pollution under different rules.

Figure C1: Data availability from CEMS by treatment status



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 C2: PM emissions by treatment status




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 C.1 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 interregnum periods when the emissions market was closed. The horizontal (red) lines denote the per-plant month market cap for each period. 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.

Table C2: Mean of the log(PM emissions) by imputation rules

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


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

Figure C3: Kernel density of PM emissions by treatment status



This figure plots the kernel density of PM emissions (kg/month) by treatment status in different stages of imputation described in Table C1. 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 C4: Kernel density of log(PM emissions) by treatment status



This figure plots the kernel density of log[PM emissions (kg/month)] by treatment status in different stages of imputation described in Table C1. Stack-Week corresponds to the emissions variable after step 2. Stack-Month, Stack-Experiment, and Treatment-Month correspond to the variables constructed based on the No Imputation Rule, the Imputation Rule A, and the Imputation Rule B, respectively. Imputing the treatment group mean causes values to converge to the group mean, so the distribution of PM emissions and that of log(PM emissions) should have less dispersion under Rule B. As the distribution of PM emissions is more clustered near the mean under Rule B, the mean of log(PM emissions) should be closer to the log of mean PM emissions for Rule B. By the concavity of log function, the log of mean is no less than the mean of log values. Hence, the mean of log(PM emissions) should be higher for Rule B than others.

C.3 Market Replacement Rule for Missing CEMS Data

Subsection C.1 gives the data imputation rule for pollution we use for the purposes of our analysis. The goal of this rule is to estimate mean emissions as accurately as possible for plants that are missing some observations on pollution. In this subsection we show the data replacement rule that was used in real time by the market. This replacement rule has two purposes: not only to estimate emissions, but also to penalize plants for non-reporting.

Table C3 shows the data replacement rule used in the market. The rule assumes that emissions are higher when data is missing for a longer period of time, in order to incentivize plants to report emissions reliably. By construction, the replacement rule used in the market will be upward biased relative to mean emissions during the time a plant is reporting.

Table C3: Imputation Rules for Missing CEMS Data

% Data Available During Week	Imputation for Missing Data Values (kg/hr)
> 95%	Plant mean operating emissions load during the compliance period
80-95%	Plant 75 th percentile emissions load during the compliance period
50-80%	Plant 90 th percentile emissions load during the compliance period
1-50%	Plant 90 th percentile emissions load during the compliance period and prior three months of valid CEMS data, up to the start of the compliance period
< 1%	Flat rate of population emissions load (8.08 kg/hr)

The table gives the data replacement rule used in the emissions market. The left column shows the raw data availability during the week. The right column shows the imputation rule for each level of data availability.

C.4 Absence of Direct Effects of Monitoring on Emissions

We have interpreted the control group in our evaluation as informative about outcomes under command-and-control regulation. One difference between the control and the status-quo in Gujarat prior to the introduction of the market was that the control group also reports real-time emissions data using CEMS technology. This data underpins our experimental evaluation but could not be used to penalize plants since the legal notifications governing the status-quo regime required that regulatory action be based on manual audits.

Here we ask whether CEMS monitoring, even without regulatory teeth, might itself change plant behavior. We worked with GPCB to rollout CEMS as a randomized experiment in order to test for such monitoring effects. Plants were randomly assigned to one of three phases. The random assignment means that plants receiving a late CEMS mandate form a valid control group for those with an early mandate.

We present results from a simple specification regressing measures of plant pollution obtained from manual measurements on CEMS treatment status. CEMS obviously cannot form the outcome measure for a treatment mandating CEMS installation. The pollution data comes from the result of manual audits carried out by the environmental regulator as part of their inspection schedule. These measurements were also actionable from a regulatory point of view under the command-and-control regime. We run the following regression:

$$y_{it} = \beta_0 + \beta_1 \text{treat}_i \times \text{post}_t + \alpha_i + \gamma_t + \varepsilon_{it}$$

where α_i is a plant fixed effect and γ_t is a month by year fixed effect. The dummy treat_i is 1 for plants in Phase 1 and 0 for plants in Phase 3. The outcome variable y is a measure of plant pollution from manual readings taken by the GPCB as part of their regular schedule of testing. The variable post_t is 0 for all control observations. For treatment (phase 1) units, it takes the value 1 once a plant has installed CEMS. β_1 is the treatment effect of CEMS on pollution.

Table C4 reports results from this regression. The main conclusion is that there is little evidence of differences in pollution between plants that had already installed CEMS relative to those that had not. Sudarshan (2023) provides related results including additional information on the rollout of CEMS in Gujarat, a description of different technologies, and practical considerations associated with using this data for regulation.

Table C4: Effects of CEMS Installation on Plant Emissions

	(1)	(2)
	PM, mg/Nm ³	Log(PM)
Treatment Effect	0.432 (23.84)	-0.0601 (0.0912)
Observations	796	796
R^2	0.3384	0.4276
Month FE	YES	YES
Unit FE	Plant	Plant
Plants	197	197
Mean	142.8	4.757

Standard errors are clustered at the plant-level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

D Appendix: Additional Results

D.1 Engineering Estimates of Abatement Costs

This section compares market prices for pollution permits to engineering estimates of the costs of running abatement equipment. In theory, the price of permits should reflect the marginal abatement costs to each plant. We check this assumption, in broad terms, by comparing permit prices to engineering measures of abatement costs.

To probe the validity of the assumption that bids can be used to infer marginal abatement costs, we compare the bids against engineering estimates of abatement costs from Indian air pollution control device vendors. As described in Section 4, the market cleared at prices between the floor of INR 5 per kg and INR 15 per kg, though average bid prices ranged as high as INR 45 per kg. Appendix Table D1 presents estimates of abatement costs under *ideal* operating conditions for four kinds of air pollution control devices under four hypothetical plant configurations. This table assumes, as is likely the case in our data, that plants are already operating a single 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. If a plant is already running a cyclone, then average (marginal) abatement costs for a mid-size plant (6 ton per hour boiler) to operate an additional cyclone are 7 (2) INR per kg and an additional bag filter 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 (21) INR per kg. Variable abatement costs therefore range from INR 2 per kg to INR 20 per kg, depending on what piece of equipment is used, under the assumed, ideal operating efficiency. If operating efficiency is actually lower, as seems likely, and the reduction in emissions therefore smaller, then the abatement cost per kg of emissions reduction would increase inversely with the decline in efficiency.

Overall, this exercise supports the assumption that the bidding data can be used to infer marginal abatement costs. We find that the market clearing permit prices overlap with engineering estimates of the marginal abatement costs associated with operating abatement equipment.

Table D1: 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

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

D.2 Treatment effects on capital installation

Table D2 shows that the ETS treatment is estimated to have no effect on the presence of air pollution control devices (APCDs), overall, since all plants already have APCDs of some kind installed. There is suggestive evidence of a small shift toward less expensive APCDs such as cyclones and bag filters (columns 1 and 2).

Table D2: Treatment effects on the presence of abatement devices

	All APCDs (1)	Components			
		Cyclone (2)	Bag (3)	Scrubber (4)	ESP (5)
ETS Treatment (=1)	0 (.)	0.0233* (0.0134)	0.0650*** (0.0231)	-0.0151 (0.0310)	-0.0311 (0.0207)
R ²	.	0.66	0.68	0.71	0.75
Control mean	1.00	0.95	0.85	0.67	0.12
Plants	276	276	276	276	276


This table reports the effects of treatment assignment on the presence of APCDs. All specifications control for the corresponding baseline value. Robust standard errors are given in parentheses with statistical significance indicated by * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

D.3 Model robustness checks

Heterogeneity in estimated elasticities by time of bid.—Section 6 estimates the elasticity of marginal abatement costs with respect to emissions using data from the first half of each compliance period. The argument is that plants only have a choice between abatement and the purchase of permits during the first half of the period, because by the end of a period, emissions are sunk and plant willingness-to-pay for permits should not depend on their abatement costs.

Figure D1 tests this idea by estimating the same elasticity separately in each week of the compliance period. We find that the elasticity of marginal abatement costs with respect to emissions is negative and economically and statistically significant during the first several weeks of the compliance period. When there are two weeks or less remaining in the compliance, by contrast, the same elasticity is estimated to be close to zero. As expected, plants' bids are not sensitive to abatement costs when there is little time left in a period in which to abate.

Figure D1: Elasticity estimate by weeks remaining in the order period



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

Table D3: Variable abatement costs under alternative regulatory regimes (with Heterogeneity by APCD)

	Emissions = 170 tons			Emissions = 240 tons		
	Price	Cost	Δ Cost	Price	Cost	Δ Cost
	(INR/kg)	(INR m)	(%)	(INR/kg)	(INR m)	(%)
	(1)	(2)	(3)	(4)	(5)	(6)
Emissions market	12.2	10.1	0	9.97	9.33	0
Constant emissions rate		10.97	8.61		10.33	10.72
Constant emissions rate, with error		11.3	11.88		10.7	14.68
Capacity-based rate		11	8.91		10.37	11.15
Capacity-based rate, with error		11.32	12.08		10.75	15.22
Capacity-based rate, correlated error		11.41	12.97		10.83	16.08

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

E Appendix: Benefit-Cost Analysis

We conduct a benefit-cost analysis of introducing an expanded ETS in Surat covering all plants that burn solid fuel. The analysis compares the social benefits of cleaner air, as measured by the valuation of the additional life-years that would be gained from pollution abatement, against the costs of emissions abatement and monitoring. For this exercise we assume that the ETS is expanded with the cap proportionately scaled to maintain the same regulatory stringency per plant as in the experiment. Table 7 summarises the analysis we describe below.

E.1 Costs of monitoring and abatement

The costs of the ETS include both the monitoring infrastructure necessary for the market and the abatement costs, or cost savings, induced by the market. In the experiment, both treatment and control groups purchased CEMS but these devices were not used under the status quo.

We estimate the annual costs of operating a CEMS system at approximately USD 5000 per plant. We arrive at this number by assuming an annualized capital cost of CEMS of INR 200,000, annual device calibration costs of INR 30,000, annual fees for software licenses and maintenance contracts of INR 60,000, and miscellaneous costs (replacement parts, labor etc) at INR 50,000. The annualized CEMS costs are based on an assumed system cost of INR 800,000 with a 4 year equipment life and no discounting. This equipment life describes the realized experience of some plants in our sample but is lower than typical manufacturer claims. License fee and contract costs are based on conversations with vendors and industry. Calibration costs assume three visits a year.

Partly offsetting this monitoring cost, our estimates imply a reduction in abatement costs of roughly USD 650 per plant-year, despite that treatment plants are operating at a sharply lower level of emissions than control plants (row A2). The net per plant costs of monitoring are therefore reduced to closer to USD 4,000. There were a total of 906 registered solid fuel burning plants in Surat during the period of the market and thus in a hypothetical scale-up to cover all plants, we estimate the total private costs, inclusive of both monitoring and abatement, to be USD 3.91

million per year.

E.2 Benefits of lower pollution

The benefit of the ETS is cleaner air. We monetize the benefit of cleaner air by using estimates of the damage from particulates, in terms of life-years gained, and valuing these life-years using estimates of the value of statistical life.

The first step is to estimate how much ETS would reduce ambient pollution (as opposed to industrial pollution emissions). This step is non-trivial because there are many sources of PM_{2.5}. A simple estimate of the impact of the ETS is that ambient pollution would fall by an amount equal to the the percentage reduction in emissions due to the regulation multiplied by the total contribution of these sources to ambient concentrations.

The first term is simply the assumed reduction in emissions, either 10%, 30% or 50%, across columns 1 to 3. For the second term we turn to an estimate from the atmospheric science literature that industrial sources in Surat raise ambient fine particle concentrations by $28.32 \mu\text{g}/\text{m}^3$. Gut-tikunda, Nishadh and Jawahar (2019) use pollution inventories combined with an atmospheric dispersion model to apportion ambient particulate concentrations in Indian cities to different sources.²⁴ The authors estimate annual average ambient PM 2.5 concentrations in the city at $88.5 \mu\text{g}/\text{m}^3$, with 32% (or $28.32 \mu\text{g}/\text{m}^3$) coming from local industry. Then the Surat ETS applied to all plants in the city would reduce pollution by $0.30 \times 28.32 = 8.5 \mu\text{g}/\text{m}^3$ (panel B, column 2).

The second step is to estimate the life-years gained from lower pollution. A large literature has attempted to quantify the impact of air pollution on life expectancy. Ebenstein et al. (2017) use a spatial regression discontinuity, at high levels of pollution in China, to estimate that a $10 \mu\text{g}/\text{m}^3$ reduction in pollution results in a 0.98 year increase in life expectancy. Other estimates in the literature include 0.61 years (Pope, Ezzati and Dockery, 2009) and 0.12 years (calculated from Table S2 in Apte et al. (2018)).

These estimates should be interpreted as the benefits of long-run changes in pollution. If we were to assume that an ETS were implemented in Surat for 70 years (roughly the current life

²⁴Their updated assessment for Surat is available at: <https://urbanemissions.info/india-apna/surat-india/>.

expectancy in India), reducing pollution each year by $8.5 \mu\text{g}/\text{m}^3$, then the health benefits from Ebenstein et al. (2017) would suggest life expectancy gains of $0.98 \times 8.5/10 = 0.83$ years per person. The population of Surat in 2021 as estimated 7.5 million people. Thus the total gain in life years would be these per-person estimates multiplied by the city population, or 6.24 million years. Assuming these accrue gradually over the 70 year period of the ETS, the gain from a single year of the program would be 89,208 years.

The third step is to value the life-years gained. We use a VSL estimate for India of USD 665,000 (Nair et al., 2021) and apply this equally to every year of an assumed 70 year life yielding a dollar value of USD 9,500 per life-year gained. This number, combined with the life years gained from a year of the ETS, would imply a single year health benefit of USD 847 million and thus a benefit to cost ratio as high as 215 to 1 (panel E, row 1, column 2). Using the lower estimates of health benefits from Apte et al. (2018) yields a benefit to cost ratio of 26 (panel E, row 4, column 2). By either estimate, the benefits of the expanded ETS greatly exceed the total of monitoring and abatement costs.