WARWICK
ECONOMICS

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

Michael Greenstone, Rohini Pande, Anant Sudarshan and Nicholas Ryan

February 2023(Updated November 2024) No: 1453
& accepted for publication in Quarterly Journal of Economics

Warwick Economics Research Papers

ISSN 2059-4283 (online)
ISSN 0083-7350 (print)

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

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

November 10, 2024

Abstract

Market-based environmental regulations are seldom used in low-income countries, where pollution is highest but state capacity is often low. We collaborated with the Gujarat Pollution Control Board (GPCB) to design and experimentally evaluate the world's first particulate matter emissions market, which covered industrial plants in a large Indian city. There are three main findings. First, the market functioned well. Treatment plants, randomly assigned to the emissions market, traded permits to become significant net sellers or buyers. After trading, treatment plants held enough permits to cover their emissions 99% of the time, compared to just 66% compliance with standards under the command-and-control status quo. Second, treatment plants reduced pollution emissions, relative to control plants, by 20% to 30%. Third, the market reduced abatement costs by an estimated 11%, holding constant emissions. This cost-savings estimate is based on plant-specific marginal cost curves that we estimate from the universe of bids to buy and sell permits in the market. The combination of pollution reductions and low costs imply that the emissions market has mortality benefits that exceed its 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, Tata Center for Development at Chicago and Yale Economic Growth Center for financial support. Ryan acknowledges financial support from NSF CAREER #2143098. 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

Many low-income countries today suffer from extraordinarily high air pollution. In India, for example, nearly the entire population of 1.4 billion people breathes air more polluted than World Health Organization standards for particulate matter, often by a factor of ten or more.[1] To face this air pollution crisis, India relies on command-and-control environmental regulations modeled on those in the United States fifty years ago (Piette, 2018). These regulations are stringent on paper, but weakly enforced in practice (Duflo et al., 2013, 2018), perhaps because strict enforcement would be too costly for firms.

A powerful alternative to standards is to regulate pollution with markets. Foundational theoretical work has shown that markets can abate pollution at the lowest possible cost (Coase, 1960; Dales, 1968). Further, the United States and the European Union have had great success in building markets to reduce air pollution (Ellerman et al., 2000; Martin, Muûls and Wagner, 2016; Dechezleprêtre, Nachtigall and Venmans, 2023). Yet, in spite of high pollution and a high concern for cost, low-income countries have not followed these examples (Stavins, 2003; Blackman, Li and Liu, 2018).[2] As a result, there is no evidence on whether pollution markets can work in countries with low state capacity. The absence of pollution markets in these countries may reflect a correct judgment that they cannot work, or a self-fulfilling policy—pollution markets cannot work only because they have not been tried before.

This paper examines how a new pollution market in India impacted plant compliance, pollution emissions and abatement costs. We study the world's first market for particulate matter emissions. This new market was introduced through a randomized-control trial that brought treatment plants into the market while keeping similar control plants under existing command-and-control regulations. We designed the market and accompanying experiment in collaboration with the Gujarat

[1] 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). Reducing air pollution throughout India to the World Health Organization standard would, by some estimates, increase Indian life expectancy by an average of five years (Pande et al., 2015; Greenstone and Hasenkopf, 2023).

[2] Markets are also rare or absent for common pool resources, like fisheries or groundwater, for which they could bring large efficiency gains (Chu, 2009; Ryan and Sudarshan, 2022)

Pollution Control Board (GPCB), the environmental regulator in Gujarat, India. Our decade-long collaboration produced the institutions to support an emissions market: monitoring infrastructure to measure pollution on a continuous basis, new market regulations and a platform to enable trade. The experiment allows us to measure causal effects of market-based regulation, which have been elusive in the rich literature on emissions trading (Fowlie, Holland and Mansur, 2012; Martin, Muûls and Wagner, 2016).

GPCB launched the market for industrial plants in and around Surat, Gujarat, a rapidly growing city of 8 million people, in 2019. Under the command-and-control status quo, plants are mandated to install abatement equipment and are then sporadically inspected in-person, by government regulators and auditors, to check that they meet limits on the concentration of pollution emissions. Our prior research has shown that enforcement in this regime is undercut by poor information and that many plants violate pollution standards as a result (Duflo et al., 2013, 2018). For the present experiment, GPCB mandated a sample of 318 large, coal-burning plants to install Continuous Emissions Monitoring Systems (CEMS) to measure the total mass of particulate matter (PM) emitted, as compared to the measurement, under the status quo, of pollution concentrations during spot visits.[3] The emissions market experiment then randomly assigned 162 out of 318 plants to the market while 156 control plants stayed under the command-and-control regime.

Treatment plants were shifted into a newly-built emissions market. GPCB set a cap on the total mass of particulates that could be collectively emitted by *all* treatment plants over a compliance period. They allocated permits to treatment plants, with permits summing to 80% of the cap distributed for free, in proportion to plant emissions potential, and 20% sold off in weekly auctions. Thereafter, treatment plants could trade permits with one another. At the conclusion of each compliance period, any treatment plant that did not hold enough permits to cover their emissions was subject to fines based on the size of the shortfall. The evaluation ran from September 2019 to

[3]The change from measuring pollution *concentration* to pollution *mass* (also called *load*) is important in its own right. Concentrations of air pollution, measured in mg/Nm^3, are the mass of pollution per volume of gas emitted from the plant's chimney during a half-hour sample on a spot visit. Load, measured in kg, is the total mass of pollution emitted—think of the weight of a giant pile of very fine dust. Measurement of load is preferable both because load is what determines a plants' contribution to ambient pollution (thus human health) and because load is measured with CEMS monitors that record pollution continuously, instead of only during spot visits.

April 2021, with an interruption due to a nationwide Covid-19 lockdown. There were a total of 10 compliance periods each of four to six weeks' duration.

There are three main findings from the analysis of the experiment. First, the market functioned well, with respect to both plant compliance and permit trade. Treatment plants complied—held enough permits to cover their emissions—in 99% of plant-periods. By contrast, the compliance rate with concentration standards at baseline was 66%. The regulator established a reputation for enforcement in the market regime early on, by levying fines specified in the market rules on a couple of plants that did not buy enough permits in the first period. We therefore find that the market caused greater compliance; high compliance is a benefit of the market treatment, as opposed to low compliance being an immutable feature of the regulatory context. The permit market, additionally, appeared to have low transactions costs. Plants traded often, with daily trading volume reaching up to 20% of the market cap. At the end of each compliance period, plant permit holdings differed greatly from initial allocations, as some became net sellers and others net buyers. Plants consumed 95% of their final permit holdings, on average, and therefore left little money on the table in unused and unsold permits.

Second, the treatment reduced treatment particulate emissions by 20% to 30%, relative to control plant emissions in the command-and-control regime.[4] The regulator set the cap, with incomplete initial data on emissions, to try to match the stringency of control regulation. The relative reduction in emissions for the treatment plants was due to a combination of the initial cap turning out to be more strict than control regulations, the higher rates of compliance in the treatment, and the regulator reducing the cap over the first several compliance periods. The regulator's reduction of the cap was an endogenous tightening of regulation after seeing the relatively low costs of compliance, reflected in low permit prices, in the market regime.

Our third main finding is that the market reduced variable abatement costs by 11% at a constant level of emissions. This estimate is derived from plant-specific marginal abatement cost (MAC)

[4]This range reflects variation in point estimates based on how emissions are imputed during periods when CEMS did not transmit pollution data from plants (see Section 4).

curves that we estimate from the universe of 8,433 treatment plant bids to buy and sell permits.[5] We exploit within-compliance-period variation in plant permit bids to estimate MAC functions specific to each plant and compliance period. Though the MAC functions are estimated with data only from treatment plants' bids, the experiment ensures that the distribution of MAC functions is the same in the control group. We therefore compare costs across regimes, at several emissions levels, by evaluating the treatment plants' abatement cost functions at either the efficient, market distribution of emissions or the distribution of emissions actually observed in the control group.

Finally, we combine our pollution and cost estimates, including also the fixed costs of setting up the market, to conduct a benefit-cost analysis of a potential market expansion. This analysis finds that, under a range of assumptions on the mortality damages from pollution, the benefits of the market exceed costs by at least twenty-five times, reflecting the high mortality costs of air pollution and the low costs of abatement in the market.

This paper contributes to the literatures in development and environmental economics. A main theme of the research on development and the environment has been the Herculean challenge 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).[6] Our findings add to this literature by showing that, at least in Gujarat, high *private* abatement costs are not the primary culprit for high emissions (see also Greenstone et al., 2022, on this theme) and that markets can be an effective way to reduce pollution.

The conventional wisdom in environmental economics on market-based regulation is that emis-

[5] We believe these data are a novel contribution in the literature where price data on executed trades has been the best case to date. For example, Ellerman et al. (2000) discuss permit prices and how prices compare to ex ante expected abatement costs in the Acid Rain program. Klier, Mattoon and Prager (1997) and Shapiro and Walker (2020) present summary statistics on transaction prices in RECLAIM and US air pollution offset markets.

[6] The one prior example of a market targeting particulates with which we are familiar was in Chile and in fact regulated boiler capacity, not emissions, because of limitations in pollution monitoring at the time (Montero, Sanchez and Katz, 2002). This coarse proxy removed the prospect of low-cost abatement because most particulate abatement happens after combustion, and plants could receive no credit for actions such as running end-of-pipe abatement equipment or changing fuels. Eventually, the market disbanded after many plants switched to natural gas in response to a fall in prices further weakening the value of this proxy.

sions markets abate pollution at lower cost than extant command-and-control regulations (Ellerman et al., 2000; Fowlie, Holland and Mansur, 2012). A general challenge for the evidence on this point is that plants regulated with markets are vastly different from those left out, making it hard to identify a valid counterfactual.[7] As a result, existing evaluations of emissions markets require strong assumptions to estimate what emissions and plant costs would have been without a market (Fowlie, Holland and Mansur, 2012; Martin, de Preux and Wagner, 2014; Borenstein et al., 2019). This paper experimentally assigns plants to the market and command-and-control regimes to provide causal estimates of the effects of an emissions market on both pollution and costs.

The paper proceeds as follows. Section 2 introduces India's status quo environmental regulation and the new emissions market and then describes the data, sample and experimental design incorporating this market. Section 3 provides evidence that the market functioned well. Section 4 presents estimates of experimental treatment effects on emissions. Section 5 describes the model and estimates marginal abatement costs from bids. Section 6 compares the costs of abatement under the two regulatory regimes and provides a benefit-cost analysis. Section 7 concludes.

2 Context and Experimental Design

We begin by describing the status quo regulatory regime (2.1). We then lay out the history of our engagement with both central and state governments (2.2), the development of the Surat market and the market rules that applied to treatment plants (2.3). In the second half of this section we describe the experimental design (2.4), our data sources (2.5) and the characteristics of sample

[7] Prior research on emissions markets makes this point explicitly. 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." Early work on the EU Emissions Trading System (ETS) for carbon dioxide highlighted the difficulty of estimating whether the ETS reduced emissions *at all*, in an environment with uncertainty about emissions and aggregate shocks (Ellerman and Buchner, 2008). This problem has remained prominent in the literature on the EU ETS. Martin, 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."

plants (2.6).

2.1 The Command and Control Status Quo

Under the status quo command and control regulations, the "command" mandates plants to install pollution-control equipment. The "control" is an intensity standard that limits pollution concentrations (for particulate matter, the limit is typically 150 mg/Nm^3). State Pollution Control Boards (SPCBs) enforce these regulations by visiting plants, measuring pollution and imposing sanctions.[8]

We collaborated with GPCB on prior research that showed incomplete compliance in the status quo. The regulator is limited in part by poor information about pollution emissions (Duflo et al., 2013). While sanctions can be large, they are typically reserved for severe offenses, leaving the many plants with routine violations with weak incentives to abate (Duflo et al., 2018). Pollution concentration limits are uniform across plants without regard to their age, size, fuel or abatement capital. Because of this uniformity, despite large differences across plants, there may be significant heterogeneity in marginal abatement costs across different plants.

2.2 Laying the Groundwork for India's First Emissions Market

The market grew out of a deep collaboration between our research team and environmental regulators to build the monitoring, regulatory and trading infrastructure for an emissions market. We proposed, at a 2010 conference of Indian state environmental regulators, that states could use emissions markets as a regulatory tool. India's Ministry of Environment and Forests then solicited a white paper on emissions markets that we co-authored (Duflo et al., 2010). The conference sparked scoping discussions on emissions markets with the State Pollution Control Boards in India's three leading industrial states, Gujarat, Maharashtra and Tamil Nadu.

The monitoring infrastructure needed to establish an emissions market had to be built before a market could be launched, in any state. Emissions markets rely on Continuous Emissions Mon-

[8]The Air Act (1981) established a command and control framework for regulating industrial pollution. The Water Act (1974) established SPCBs as environmental regulators for water pollution standards. The Air Act expanded these powers to include air pollution. SPCBs can introduce additional regulation in highly polluted regions.

itoring Systems (CEMS), for which no Indian standards existed circa 2010. A Central Pollution Control Board (CPCB) panel, including co-author Sudarshan as a member, in 2013 drafted technical standards for CEMS usage in India (Central Pollution Control Board, 2013). These standards were developed with a pilot market in mind but also enabled a nationwide movement towards the adoption of CEMS.[9] Gujarat was the first of the three interested states to mandate CEMS devices when the standards were published, which prompted our research-policy collaboration with GPCB on the development of an emissions market. GPCB pushed the roll-out and testing of CEMS at scale, a laborious process that required a new, private ecosystem to install and maintain the devices (Sudarshan, 2023). Tamil Nadu and Maharashtra have since adopted CEMS for monitoring larger plants. As of 2024, we are also working with the Maharashtra Pollution Control Board on the design of a market for sulfur dioxide.

To start an emissions market GPCB needed to issue regulations, establish a trading platform and build market participants' capacity. Together with our research team at JPAL South Asia, we collaborated with GPCB on each of these steps. GPCB selected NCDEX e-Markets Limited (NeML), a subsidiary of a leading Indian commodity exchange, to host the market. GPCB, NeML, the co-authors of this paper and our research team jointly developed market rules. The Forest and Environment Department, Government of Gujarat formally notified the market on June 4th, 2019 (via notification GVN-2019-17-GPCB-SFS-1-2019-ETS-T). The market is supervised on an ongoing basis by a Market Oversight Committee (MOC) chaired by the Chairman, GPCB and with additional members including the GPCB Regional Officer, GPCB Environmental Engineers, NeML, the head of the South Gujarat Textile Processor's Association (SGTPA), the concerned industry association, and this paper's co-authors. The GPCB and the SGTPA hosted a series of stakeholder capacity-building workshops to train GPCB's own officials and regulated plants on the rules of the market, penalties for non-compliance, how to participate in auctions and how to trade.

[9] In 2014, the CPCB mandated installation of Continuous Emissions Monitoring Systems (CEMS) for large plants in 17 manufacturing sectors across India. This mandate is restricted to large plants and the CEMS data from this mandate is largely not used in the enforcement of regulation; for example, CEMS readings can direct the regulator's priorities but cannot be used as the legal basis for sanctioning plants in the status quo.

2.3 Surat Emissions Market Design

The Surat emissions trading scheme is the world's first particulate emissions market as well as India's first market for any pollutant. GPCB chose to locate the emissions market in Surat, an industrial hub with a population of 7 million, because the city is critically polluted and has a high contribution of point-source industrial emissions to ambient pollution levels.[10] The market is a standard cap-and-trade design in which plants are allocated, or can buy, permits granting the right to emit and face fines for non-compliance if emissions exceed their permit holdings. Here we detail the market design.

Cap.— The cap in the market limits the *load*, or mass, of pollution emitted. GPCB initially capped particulate emissions at 280 tons per month. This cap was an approximation that assumed plants would run at the maximum available capacity and produce emissions at the maximum concentration allowed under the status quo the whole time. Once the market began, GPCB got better estimates of load from CEMS devices and judged the initial cap had been set too high. GPCB then adjusted the cap downwards, in steps, to 170 tons per month (see Appendix Table A2).

Permit allocation.— Each permit allows a plant to emit one kilogram of particulate matter (PM). PM includes suspended particles of all sizes, since neither CEMS devices nor manual monitoring in the status quo differentiate by particle size, as is common in ambient monitoring of PM. Permits are valid only for one compliance period of four to six weeks' duration (see Appendix Table A1). Permits expire at the end of each compliance period so plants cannot bank or borrow across periods. At the start of each compliance period, 80% of issued permits were given to plants for free, proportional to a scale measure, emissions capacity, calculated prior to the market.[11] The balance of 20% of permits were sold to plants at auction.

Permit trade.— The main means of trade is a uniform price, multi-unit two-sided auction held weekly on Tuesday. Plants can bid to purchase or offer to sell permits in the auction by

[10]Roughly one-third of ambient fine particle pollution in and around Surat city comes from industrial plants; twice as large as that of the next most significant source, transportation, at 16% (Guttikunda, Nishadh and Jawahar, 2019). A current assessment for Surat is at: `https://urbanemissions.info/india-apna/surat-india`.

[11]Emissions capacity, measured in tons of steam equivalent per hour, is the sum of the output capacities of the boiler and thermic fluid heater, the two main fuel-burning pieces of equipment in sample plants.

submitting bid price-quantity pairs to the market operator. The market-clearing price is the least price at which permit supply weakly exceeds permit demand. 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 during the first auction, they continue to be offered in subsequent weekly auctions. Plants can also buy or sell permits during the week via over the counter (OTC) trades at the most recent auction's permit price. This OTC price restriction was adopted to encourage auction participation and limit price volatility. Market rules impose tight holding limits to prevent plants from gaining market power over permits.[12]

Price collar.— Prices were limited to between ₹5 and ₹100 per kg.[13] GPCB supported the floor price with a commitment to buy back permits at the floor at the end of each compliance period. GPCB supported the ceiling price with a commitment to sell permits at the ceiling at the end of each compliance period in unlimited quantity.

Compliance and penalties for non-compliance.— Compliance is enforced with financial penalties. At the start of the market, plants had to post an environmental bond known as an Environmental Damage Compensation Deposit (EDCD).[14] The size of EDCD varied with plant scale, and for most plants was ₹200,000, well in excess of permit expenditures at market prices. After a compliance period ended, at which point emissions were known with certainty, plants had a one-week true-up period to make further trades. At the end of the true-up plant permit holdings were compared against their total emissions in the period. Plants with insufficient permits were fined twice the ceiling price for every unit of emissions above their permit holdings with the fine deducted from the EDCD (which the plant was then required to top-up).

[12]The limit for each plant was the greater of 1.5 times that plant's initial allocation or 3% of the aggregate market cap for the compliance period.

[13]This range was informed by engineering estimates that particulate matter abatement by equipment commonly used in the sample, could occur at an average cost of ₹10 to 40, depending on type of equipment installed and plant scale. The ceiling price was sufficiently high that plants would rather abate than pay the ceiling price for permits.

[14]While these EDCD bonds were new, GPCB commonly used other environmental bonds for violating plants to guarantee against future non-compliance. The authority to impose fines derives 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).

Differences between the market and the status quo.— To summarize, the treatment emissions market differs from control regulation in three main respects. First, the compliance obligation for each treatment plant is tradeable, and ultimately determined by the market-level cap, rather than a fixed plant standard. Second, treatment plants are regulated for total pollution load (i.e., mass), rather than concentration at one point in time. Third, penalties for noncompliance in the market are financial and set based on ex ante rules.

2.4 Experimental Design

The initial sample included the 342 plants under the GPCB regional office in Surat with the highest air pollution potential, as captured by the following criteria: (i) the plant must consume solid fuel (coal or lignite, mostly), (ii) have a boiler capacity of at least one ton of steam per hour, and (iii) a stack diameter of at least 24 cm (required for CEMS). All sample plants were mandated to install CEMS. Half of sample plants were randomly assigned to the market treatment. After treatment assignment but before the market commenced, the GPCB gathered more information on plants' operating status, and deemed ineligible any plants that were closed or operated only seasonally. This restriction left 162 and 156 plants in the treatment and control groups, respectively. Appendix Table A3 summarizes plant attrition by treatment arm.

Figure 1 displays sample plants and ambient fine particulate (PM2.5) concentrations. The average pairwise distance between sample plants is 11 km, far below 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 PM2.5 concentrations in Surat are 10 to 20 times the WHO standard of 5 $\mu g/m^3$.

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. They were aware that the real market would start after the mock periods. Six real compliance periods followed from September

16th, 2019 to March 22nd, 2020.[15] Market operations were then suspended as part of a nationwide Covid-19 lockdown that closed all sample plants. The market restarted in December, 2020 and we have data from four additional compliance periods up to when the second (Delta) wave of Covid hit in April, 2021.[16] The treatment thus spans roughly one year of market operations spread over one-and-a-half calendar years.

2.5 Data sources

The paper relies on three main sources of data.

Pollution measurement.—CEMS devices were installed in sample plants to provide high-frequency data on particulate matter (see Appendix C.1 for more information on CEMS). The CEMS devices themselves have no direct effect on emissions.[17] This finding may be surprising, because a lack of information is a major constraint on status quo regulation (Duflo et al., 2013, 2018). However, CEMS installation, in the control group, was not accompanied by any change in regulation and existing regulations specify that only in-person samples are legally admissible as a basis for penalizing plants with high emissions. CEMS could therefore be used to direct or coordinate in-person visits but not to impose penalties. This regime is common to both the CEMS installations in our experiment and the larger Central Pollution Control Board mandate of CEMS for select industries across India (see footnote 9).

To incentivize CEMS uptime, plants that failed to report CEMS data for any period of time had missing emissions data replaced according to a rule that became increasingly punitive as the share of data missing rose (see Appendix Table C2).[18] Emissions with this replacement for missing data were called "validated emissions" and used to calculate compliance. The replacement rule created

[15]The initial schedule envisaged periods starting on the first of the month but the GPCB pushed it back by two weeks to allow plants more time to calibrate their CEMS devices.

[16]The market restarted after the more severe Delta wave shutdowns. However several plants did not reopen immediately or reduced operating hours. Permit prices were correspondingly often at the floor. The GPCB began plans to expand coverage to the control and then implemented this expansion. For these reasons we prefer not to use data after the Delta wave in our main analysis. We report results including these periods as a robustness check in Appendix F.5.

[17]The CEMS deployment was evaluated in a randomized control trial prior to the market launch. This trial shows that CEMS alone had no effect on plant pollution emissions (Appendix C).

[18]Missing CEMS data might occur for different reasons including reasons outside plants control such as extended internet and electricity outages or device malfunctions. Replacement rules treated all missing data alike.

stronger incentives for treatment plants to comply because non-compliance meant plants would have higher emissions and need to buy more permits. For this reason, treatment plants reported data at higher rates throughout the experiment. CEMS data availability for control plants largely caught up in later compliance periods and average weekly data reporting rose to over 85% by the end of the sample (Appendix Figure C2).

We use different emissions imputation rules in our data analysis to try to form an unbiased, rather than a punitive, estimate of non-reported emissions.[19] Appendix C describes our treatment of missing data. There are two steps to imputing data at the plant-month level. First, we impute missing daily observations within a week with the emissions rate from other days in that week or weeks in that month for the same plant. Second, we use several alternative rules, including no further imputation, to fill in plant-week observations when a plant does not report at any time during the week. These rules include using the emissions rate from the same plant at other times in the experiment (rule A) or the emissions of plants in the same treatment arm in the same month (rule B). Plant-weeks are summed within a month and within a plant to get plant-months (for plants with multiple stacks, all of the above steps first happen at the stack level before aggregating across stacks to the plant level). We discuss the robustness of our estimated treatment effects on emissions to different imputation rules with the empirical results in Section 4.

Trading data.—Plants could trade on the platform of the market operator (NeML) via either auction bids or over-the-counter trades, as described in Section 2. We observe initial allocations of permits to plants and data on the universe of plant bids, consisting of a price and quantity pair, irrespective of whether they resulted in a trade. For example, a plant willing to sell permits only at a very high price might submit a series of unfulfilled bids. With initial allocations and subsequent transactions we construct the full history of plant permit holdings and can observe, in particular, whether plants hold enough permits at the end of each compliance period to cover their emissions.

[19]The emissions replacement rule for the market was designed to punish non-reporting plants by over-estimating their emissions. Since control plants complied with CEMS requirements more slowly, using the market replacement rules for missing data in control plants in our analysis would increase control emissions and bias upward in magnitude the treatment effect on emissions.

Plant surveys.—An in-person baseline plant survey was conducted from December 2018 to January 2019 and a phone-based endline survey wave in November 2020. The surveys had two parts, general and technical. The general part was administered to the plant owner or manager, and covered economic variables like inputs, outputs, sales and energy use. For the technical part, our team observed abatement equipment installed on every point source of emissions in the plant and recorded the characteristics of all emissions sources and abatement equipment. They also interviewed plant staff about costs of equipment operation. We use the baseline data to characterize status quo regulation and the endline data to investigate investments in abatement equipment and their cost.

2.6 Summary Statistics

Table 1 summarizes plant covariates at baseline by treatment arm. Sample plants are large factories with high energy and related input costs, though many are formally classified as "small scale," based on their capital stock at the plant's establishment. The average control plant spends \$350,000 a year on electricity (Panel A). The boiler, the plant's main pollution source, costs \$112,000 annually to run, excluding fuel expenditures. The sample is balanced across a wide range of input, output, equipment, and pollution metrics.

The "command" portion of regulation works well: all plants have installed some air pollution control device (APCD). Table 1, panel B shows that 97% of control plants (98% of treatment) have a cyclone installed, 88% (80%) have a bag filter installed, 61% (64%) have a scrubber installed and 8% (12%) have an electrostatic precipitator. Installation rates are inversely proportional to the cost and efficacy of abatement equipment. Cyclones are inexpensive but have a low efficacy, reducing PM emissions by $60-90\%$ and PM2.5 by only $0-40\%$. Larger plants with multiple emissions sources are mandated to install more expensive APCDs like scrubbers, which remove more than 95% of PM.

The "control" portion works less well, as many plants violate pollution standards. At baseline, pollution concentrations and mass rates are balanced across treatment arms (Table 1, panel C). PM concentration in stack gas average 169 mg/Nm^3 in the control group and 179 mg/Nm^3 in the treatment group. Both *average* emissions levels exceed the SPM *maximum* standard of 150 mg/Nm^3. Roughly 30% of plants in both arms have pollution concentrations above the standard. The distribution of concentrations prior to the market's launch by treatment arm is shown in Appendix Figure B1.

Table 2 shows the evolution of bid counts, prices, and quantities. Table rows show statistics by compliance period and columns show the number, quantity and price of bids submitted. 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 ₹12.70 per kg (standard deviation ₹16.65 per kg) (column 6). Both the level and 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 bidding activity volume is large, especially since permits equalling 80% of the overall market cap were allocated to plants before the first auction.

3 Trade and Compliance in the Emissions Market

Emission market analysis often assumes compliance with market rules, particularly the key rule that plants must hold permits for each unit of emissions. This assumption may fail in low-state-capacity environments such as India where non-compliance is widespread in the status quo regime. It is therefore important to demonstrate that the emissions market functioned, both for the policy reason of validating that an emissions market can work in this setting and to understand the first stage for our analysis of how the market affects emissions and costs. Below, we provide a descriptive analysis of the treatment market, demonstrating that (i) trade was vigorous; (ii) final permit holdings differed from initial allocations, consistent with unobservable cost heterogeneity; and (iii) plants complied with the permit holding rule almost perfectly.

3.1 Permit Trade

Figure 2 depicts the weekly time series of permit prices (panel A) and quantities traded (panel B). The scattered data points show the weekly mean permit bid. Panel A's solid line reflects weekly clearing prices, which vary between blue and black to indicate the change in 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 red dashed line represents the price floor of ₹5 per kg.

Market-clearing prices range from the price floor of ₹5 per kg to ₹16 per kg depending on compliance period and the particular week. Prices were generally lower in the pre-Covid-interruption compliance periods (1-6), when the cap was looser, and higher once 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.[20] Prices generally rise within a compliance period after the regulator sells many permits at the floor price in the first week.[21] 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. A similar pattern of declining trading prices was observed at the start of the US Acid Rain program market for sulfur dioxide (Schmalensee et al., 1998).

Panel B plots daily permit quantities traded as a fraction of the compliance period cap. The double-sided auction on the first Tuesday of a compliance period typically causes a spike in quan-

[20]This 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. After the market closes and this uncertainty is resolved, prices should converge to the ceiling, if the market is short and firms face non-compliance penalties, or to the floor, if the market is long and excess permits will be sold back to the regulator. The market-clearing mechanism of a single auction after the close of the compliance period may mute this end of period price volatility.

[21]This pattern may be surprising because, in principle, plants could arbitrage price differences within a period by buying permits earlier and selling any excess at the floor price at the end of the compliance period without risk. The simplest explanation for why plants do not exploit this arbitrage is that the potential arbitrage profits are constrained by tight holding limits (see footnote 12). We note that the average bid prices in the first week of the compliance period are well above the floor price, though the market-clearing price typically remains at the floor. In addition, in interviews about their trading strategies, a couple plants mentioned that they are endowed with a free allocation of permits and they do not look to buy until they consume part of this endowment.

tity. Overall trade volume is significant, reaching up to 20% of the monthly cap, or more, on some 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.

3.2 Permit allocations

In the Surat market, permits totaling 80% of the cap were allocated *pro rata* based on a plant's total heat capacity, the regulator's best ex ante measure of emissions capacity. If there is unobserved heterogeneity in costs across plants, plants should trade away from their initial allocations depending on underlying abatement costs. To test this idea, Figure 3 plots the distribution of plant emissions as a percentage of their initial permit allocation in each compliance period. Plants that emit exactly what they were allocated appear at 100%, while plants that emit twice what they were allocated appear at 200%. Because only 80% of the total cap is freely allocated per period, with the rest auctioned, aggregate emissions as a percentage of the initial allocation will equal 125%, unless the price hits the floor or ceiling.

In every compliance period, emissions are widely dispersed with respect to initial permit allocations. Most plants fall between 50 – 200% of their initial allocation, with the relatively modest share between 100 – 125% revealing that a significant number of plants became substantial net buyers or sellers of permits. This dispersion indicates that the market meets two criteria: (i) low transaction costs, as lants are unconstrained by initial permit allocations and trade to adjust permit holdings; and, (ii) unobserved heterogeneity among plants, since the heat capacity on which initial allocations are based turns out to be a noisy proxy for ultimate emissions. Differences across plants in capacity utilization, emissions rates, and marginal abatement costs may all contribute to dispersion relative to the capacity-based measure used for permit allocation.

3.3 Compliance with Market Rules

Figure 4 plots the distribution of emissions across plants as a fraction of permit holdings at the end of each of the ten true-up periods that followed the ten compliance periods (rather than emissions as a fraction of *initial* allocations, as shown in Figure 3). Any plant that emits more than its final permit holdings (i.e., more than 100% in the figure) is non-compliant. Plants that emit less than their permit holdings (i.e., less than 100%), on the other hand, "leave money on the table" by not selling their excess permits to other plants or back to the regulator at the floor price.

Compliance, defined as emissions during the compliance period being equal to or less than permit holdings at the end of the true-up period, is nearly perfect. Across all panels, plants hold enough permits to cover their emissions in 99% of plant-periods. We observe only a few scattered non-compliant plants (see periods 1, 3 and 8).[22] By contrast, only 66% of treatment plants and 72% of control plants were compliance with concentration standards at baseline (see Figure B1). We believe compliance was high in the market because the regulator established that violators would be penalized immediately and a non-discretionary rule for fines proportional to permit shortfalls made strict enforcement credible.[23] Plant compliance in the market is an endogenous outcome reflecting a new regulatory regime with rule-based fines. In the status quo, regulators could impose severe penalties, up to closing a plant down. However, the most severe penalties were costly to impose, seldom and unpredictably applied, and did not scale smoothly with the magnitude of a violation (Duflo et al., 2018).

Lastly, plants did not leave much money on the table. The mass in the histograms is stacked at 100%; the vast majority of plants hold permits that exactly equal or only slightly exceed their total emissions at the end of each period. Looking down the first column of distributions, and then down the second, we can see that more plants left money on the table in early compliance periods,

[22] The permit holdings in Figure 4 and our calculations of compliance include permits that GPCB gave to plants in period 7 during the compliance period and above their typical allocations. Period 7 was the initial post-Covid-lockdown period and many plants were not operating but had high imputed emissions. These plants petitioned GPCB that because they were closed they should not be imputed at high rates, and GPCB accepted this argument by adjusting their permit holdings to cover the imputed emissions. We include the adjusted permits in our baseline calculation of compliance because GPCB authorized them. Without these adjustments, 37 plants would have been non-compliant (on the basis of imputed emissions) in period 7. Appendix Figure F3 repeats Figure 4 deducting these permit adjustments from plants' holdings. We find that plants comply in 97% of plant-periods, instead of 99%, without these adjustments.

[23] Emissions for two plants exceeded permit holdings during the initial compliance period. Plant A emitted 3928 kg against permit holdings of 3456 kg and Plant B emitted 4716 kg against permit holdings of 1456 kg. These plants were levied Environmental Damage Compensation (EDCD) in accordance with market rules. Plant A paid the EDCD and then topped up their environmental bond. As plant B had failed to post the required bond, the regulator ordered plant closure. Plant B posted a bond and paid a penalty of ₹652,000, more than ten times the cost to buy permits on the market to cover emissions in the period. The regulator allowed the plant to reopen after two weeks.

when market participants had limited experience and the clearing price was relatively low. In later compliance periods almost all plants hold only the permits they need to cover their emissions.[24] The precision of permit holdings suggests that plants understood the incentives for permit trade and that transaction costs in the market were low.

4 Experimental Results on Pollution Emissions

We now exploit the randomized assignment of plants to treatment to evaluate if the emissions market reduced pollution emissions. Figure 5 displays weekly mean per plant emissions in kilograms per month, from April 2019 to April 2021, by treatment arm. The solid (blue) and dashed (grey) lines represent treatment and control plants, respectively. Vertical lines separate market compliance periods. The Covid-19 lockdown, denoted interregnum on the figure, is shaded in light blue and divides the sample into early (1 to 6) and late (7 to 10) compliance periods.

We include pre-experiment data in Figure 5, despite poor data reporting then. The time pattern of reporting rates for treatment and control groups, depicted in Appendix Figure C2, explains the apparent drop in monthly mean emissions in both treatment and control groups prior to the experiment. Initial reporting is low and concentrated among larger plants, while treatment plants report more than control plants. As data reporting increased, smaller plants began reporting, lowering average emissions for reporting plants. Control plant reporting also rose such that the treatment-control gap in reporting narrowed to a few percentage points by the end of the experimental period. As described in Section 2.5, the main pollution series in Figure 5 imputes missing plant emissions using observations from the same week or month for the same plant.

Figure 5 yields two findings. First, treatment plants largely followed the market rules and emissions therefore met, or nearly met, the cap in all compliance periods, in contrast to the poor compliance with the concentration standard in the status quo. We plot the mean emissions per plant

[24] On average across all compliance periods plants consumed 95% of their permits and 78% of plants held permits, at the end of the period, exactly equal to their emissions (down the the last kg). The share of plants holding exactly the permits they needed rose from 51% in the first period to 91% and 84% in the last two, respectively, and the fraction of permits consumed rose from 87% in the first period to 97% and 95% in the last two.

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, near or below this level, sometimes substantially below (around the Diwali holiday, in November, many plants briefly cease operations and emissions fall). Aggregate emissions exceeded the cap by 3% in period 8 and plants were penalized for their excess emissions in accord with the market rules. The seeming over-compliance in early compliance periods reflects the market replacement rule for missing data (Appendix Table C2).[25]

Second, pollution emissions are consistently lower in the treatment than in the control group. By the start of compliance period 1, in September 2019, treatment plants emit roughly 300 kilograms per month less particulate matter than control plants. The treatment-induced gap in average emissions remains throughout despite marked increases in control plant reporting. We use a regression analysis at the plant-month level to estimate the treatment effect size:

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

$Pollution_{it}$ 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 effects. We restrict the data to the period when the experiment was running, when reporting was highest, and report the robustness of our estimates to alternative imputation rules for missing emissions (see also Appendix C). Standard errors are clustered at the plant level.

Table 3 reports the results. Columns 1 to 4 use pollution series that do not impute across plant-months and therefore drop plant-months with no data. Columns 1 and 2 report unweighted regressions. To capture treatment effects on the full plant sample, columns 3 and 4 regressions are reweighted by the inverse probability of a plant reporting emissions (DiNardo, Fortin and Lemieux, 1996), where we use baseline observable characteristics to predict a plant's reporting probability. Columns 5 to 8 report specifications that impute missing pollution observations using two different

[25] We 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 4).

imputation rules (see Appendix C). Briefly, Rule A, in columns 5 and 6, imputes a stack missing emissions in a given month at its mean emissions from other months 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. Even-numbered columns include year-month fixed effects and odd-numbered columns do not.

The market treatment significantly reduced PM emissions. Without imputation and reweighting, column 2 reports a treatment effect on log emissions of -0.193 log points (standard error 0.076 log points). Re-weighting gives very similar estimates (column 4). The treatment effect on pollution is larger with either imputation rule than with the raw data (columns $5-8$ compared to columns $1-4$). 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. Imputing missing data increases the magnitude of treatment effects because imputations tend to replace missing control observations for log particulate emissions with values higher than the mean among reporting control plants. Thus using imputed control plant emissions raises control emissions and the estimated difference between treatment and control emissions.[26] Appendix F.5 shows that the results are also robust to subsetting the sample to only include data either before or after the Covid lockdown.[27]

A natural next question is the impact of the market treatment on plant abatement costs. There are two challenges to comparing costs across treatment arms. First, the lower emissions in the treatment group mean that a comparison of costs across arms will not isolate the effect of the emissions market on costs holding constant the level of emissions. Second, permit market bids are the best measure of marginal costs, but they are only available in the treatment group, because only the

[26] In Appendix Table C4, we perform a bounding analysis of the treatment effect on pollution that allows for *differential* imputation rules by treatment arm (though we have no evidence of differential pollution during periods of missing data). We find that, given relatively low rates of missing data in the treatment, emissions at treatment plants have to be imputed at a much higher rate than emissions at control plants in order to meaningfully reduce the estimated treatment effect on emissions.

[27] In specifications without imputation or with imputation Rule A there is no statistically significant difference in the treatment effect before and after the lockdown. In specifications with imputation rule B, the treatment effect is statistically smaller in magnitude (less negative) after the lockdown but remains large, negative and statistically significantly different from zero.

treatment group participated in the market. For these reasons, we now introduce a revealed preference approach based on permit bidding data to estimate marginal abatement cost functions. We then describe how these functions can be used to evaluate costs for *any* distribution of emissions, and therefore to compare costs across treatment arms holding emissions constant.

5 Estimating Marginal Abatement Cost Curves

Marginal abatement cost (MAC) curves are the key to measuring the costs of virtually any environmental regulation. With MAC curves, one can evaluate marginal costs for any policy that implies an allocation of emissions to different plants, or integrate these functions to recover the variable costs for any proposed emissions reduction. While MAC curves are the theoretical foundation for the study of regulation, they are not observable, which makes it difficult to apply these ideas.

We address this problem by using our extraordinarily rich bid data to estimate plants' MAC curves. Section 5.1 outlines the two assumptions needed to interpret plants' emissions market bids as measures of their marginal abatement costs. Section 5.2 describes how we estimate MAC curves using bidding data and Section 5.3 presents the results, emphasizing the heterogeneity in estimated MACs across treatment plants. We argue that the distribution of MAC curves in the treatment group represents the distribution in the control group also and so can be used to compare abatement costs across these regimes.

5.1 Emission Market Bids as Measures of Marginal Abatement Costs

We start with the two assumptions required to interpret plant permit bids as observations of marginal abatement cost. The first assumption is that each plant i seeks to minimize its compliance costs in a given period by solving

$$\min_{E_i} Z_{i0} + Z_i(E_i) + P(E_i - A_i), \tag{2}$$

where E_i is the plant choice of emissions, Z_{i0} is the fixed cost of abatement, $Z_i(\cdot)$ gives variable abatement expenditures as a function of emissions, P is the equilibrium permit price, known to the plant, and A_i is the regulator's free allocation of permits to the plant. The substance of this cost minimization assumption is that plants trade-off their own, in-house emissions abatement with the purchase of permits to lower compliance costs.

Variable abatement costs $Z_i(\cdot)$ include costs associated with running pollution abatement equipment more frequently, changing inputs like filters or chemicals more often, or devoting more labor to operation and maintenance of machines. We expect that $Z_i' < 0$ and $Z_i'' > 0$; abatement expenditures decrease as emissions increase but at a rate that decreases in magnitude as emissions grow. (Equivalently, marginal costs of abatement are increasing in abatement.) Plants are already mandated to install pollution abatement equipment in the status quo, incurring a fixed cost Z_{i0} (see Table 1, Panel B). We will document in Section 5.2 that these fixed expenditures do not change in response to the emissions market treatment.

The second assumption is that plants are price-takers in the market for permits. This assumption is appropriate in our setting because no plant holds a large share of the permit market. For example, the 90th, 95th and 99th percentiles of plant permit allocations as a share of the emissions market cap are 0.6%, 0.9% and 4.5%, respectively. Market rules also imposed holding limits to prevent speculative price manipulation (see footnote 12). Under this assumption, the solution to (2) satisfies the first-order condition

$$-\frac{\partial Z_i(E_i)}{\partial E_i} \quad \equiv \quad MAC(E_i) = P. \tag{3}$$

This is the familiar condition that the plant's marginal abatement costs at a given level of emissions equal the permit price. This condition holds regardless of the plant's fixed abatement costs Z_{i0} and initial permit allocation A_i. Therefore, under the two assumptions of cost minimization and no market power for permits, plant marginal costs of abatement, which are not observable, equal their permit bids, which we record in our data. Equation (3) therefore provides a basis for estimating

MAC curves.

5.2 Marginal Abatement Cost Curve Estimation

We use permit bid data from treatment plants, participating in the emissions market, to estimate plants' MAC functions. Bids vary both across the ten compliance periods and within each period, because most plants submit multiple bids per period (see Table 2, column 2 and Appendix Figure B2). Our main specification is an iso-elastic marginal abatement cost curve that allows for higher or lower cost *functions* for each plant-period:

$$\log b_{itk} \quad = \quad \beta_1 \log E_{itk} + \xi_{it} + \varepsilon_{itk}. \tag{4}$$

The dependent variable is the log of plant i's bid number k in period t, our measure of plant marginal abatement cost at a given emissions level. The main explanatory variable $\log E_{itk}$ is the log of the implied plant emissions if bid k in period t were to be executed.[28] The error consists of a plant-period effect ξ_{it} and an idiosyncratic term ε_{itk} that we discuss below. We vary specification (4) by both (i) using plant observable characteristics or plant fixed effects in place of plant-period effects and (ii) allowing heterogeneity in β_1 for plants with different observable characteristics. We discuss these variants with the estimates and focus on this main specification here.

The main parameter of interest is β_1, the elasticity of marginal abatement costs with respect to emissions. We expect that marginal abatement costs are decreasing in emissions (increasing in abatement) such that $\beta_1 < 0$. The main challenge in estimating the abatement cost elasticity, in general, is that a plant's chosen emissions level E_{it} will be endogenous to abatement cost shocks. If plants with high abatement costs at a given time choose higher emissions levels, then the estimated abatement elasticity would be biased upwards ($\hat{\beta}_1 > \beta_1$).

[28]Practically, we calculate the emissions associated with a bid as the plant's permit holdings if that bid were executed, because, given near-perfect compliance, emissions are equivalent to permit holdings in the market. 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, $E_{it,k=1} = 1,500 + 500 = 2,000$ is the sum of the initial allocation and the amount the bid seeks to buy. Generically, let $\mathscr{K}(k) = \{k' : k' < k, k' \text{executed}\}$ be the set of bids already executed at the time k is offered. Then $E_{itk} = A_{it} + \sum_{k' \in \mathscr{K}(k)} B_{itk'} + B_{itk}$ where buy bids are represented as positive quantities B and sell bids with negative quantities.

Our specification addresses this challenge and allows for unbiased estimation of β_1 via ordinary least squares using variation in bids within a plant-period. This assumes that variation in $\log E_{itk}$ *within* a period reflects plant bids at different points along their MAC curve, up to an idiosyncratic error. Formally, this assumption is the familiar one of mean conditional exogeneity $\mathbb{E}[\varepsilon_{itk}|E_{itk},\xi_{it}] = 0$. While plants almost certainly emit more when they face a high abatement cost shock (e.g., due to a positive demand shock), we find it plausible that they do not adjust emission levels to cost shocks at high frequency. For example, a plant cannot change its set production schedule in a week simply because of higher than expected ash content of its fuel supply.[29]

In order to align the data with the time horizon of plant abatement choices we estimate (4) in a restricted sample of bids from only the first half of each compliance period, while also exploring estimates in other samples.[30] In our model, plants choose whether to comply by reducing emissions or buying permits, which is a good characterization of decisions in the early portion of a compliance period. As the end of a period approaches, however, plants no longer face the same trade-off, as within-period plant emissions are largely sunk. At the extreme, after the compliance period has ended and the true-up period begins, emissions are fixed, so we expect plant demand for permits to be inelastic and bids no longer to reflect marginal abatement costs.

Alternative step function marginal abatement cost functions.—Equation (4) specifies a smooth, iso-elastic marginal cost function. This functional form may not represent plants' underlying technology and therefore marginal costs at emissions levels far from those at which plants bid in the data. A simple engineering model of abatement would specify that each piece of capital abatement equipment can abate some fraction of emissions at a fixed marginal cost, beyond which

[29]One formal economic justification for this identifying assumption is that plants form unbiased expectations of their emissions, and therefore marginal costs, at the time of bidding, but are uncertain about later shocks and therefore their exact emissions level and marginal cost. For example, assume plants anticipate emissions $\widetilde{E}_{itk} = E_{it}\nu_{itk}$ with $\nu_{itk} \perp E_{itk},\xi_{it}$ and $\mathbb{E}[\log \nu_{itk}] = 0$. plants bidding their expected marginal costs yields specification (4) with a residual $\varepsilon_{itk} = \beta_1 \log \nu_{itk}$ based on the forecast error.

[30]This strategy of restricting the sample to the first half of compliance periods is feasible because 3,120 out of the 8,433 total bids were submitted in the first half of a period and 2,775 were offered by plants that submitted multiple bids in that time. Variation in bids to estimate β_1 within a plant-period comes from plants making multiple bids in the first half of a compliance period. There are a total of 1,560 (=156 plants $\times$ 10 compliance periods) plant-by-period cells and in 1,140 of these cells plants made more than 1 bid.

a more expensive piece of equipment must be used. Our data include the type and specifications of abatement equipment in each plant, allowing us to consider this alternative functional form for MAC curves. Appendix E provides a step function abatement cost model based on technologies available to the plants and describes how we use bid data to estimate the height of the steps. We compare the fit of this alternative model to our main specification in Section 6.1.

5.3 Marginal Abatement Cost Function Estimates

Elasticity of abatement cost with respect to emissions.—Table 4 reports β_1, the estimated MAC elasticity for plant emissions, using alternative specifications of equation (4). Column 1 reports a specification that controls for only plant heat capacity, the relevant measure of plant scale for emissions, analogous to the horsepower of a car engine. Columns 2 and 3 report specifications with period fixed effects and plant and period fixed effects, respectively, which absorb all time-invariant plant characteristics such as heat capacity. Column 4 reports our preferred specification (4) with plant-by-period fixed effects that non-parametrically control for all plant-period abatement cost shocks. In column 5, we additionally allow the coefficient β_1 to vary with the type of air pollution control devices (APCDs) the plant has installed at baseline.

Our preferred column 4 estimate reports an elasticity of bid prices (MAC) with respect to emissions of -0.609 (standard error 0.087). This estimate validates the standard intuition that marginal costs of abatement fall as emissions rise (e.g., marginal costs are increasing in abatement). A comparison of the column 4 estimate to those in columns 1 to 3 demonstrates 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), towards zero. Even with plant and period fixed effects (column 3), the elasticity is less than half the magnitude of our preferred estimate. The upward bias in these estimates suggests that plant emissions are endogenous to cost shocks: plants choose to emit more when they have a high plant-period shock to marginal abatement costs. In column 5 we relax the assumption of a common elasticity of abatement across plants. We find that the MAC curve is slightly more elastic (in absolute terms) for plants which installed the less expensive APCDs (cyclones and bag filters versus scrubbers and ESPs) installed, but that this difference is quantitatively small and not statistically significant.

Three additional pieces of evidence support our preference for the Table 4, column 4 specification. First, the iso-elastic MAC curve fits the bid data well. Figure 6 plots log bid prices, after residualizing on plant-period fixed effects, against log plant emissions. Bid prices (marginal costs) rise at lower levels of emissions. The line of best fit corresponds to the constant elasticity of marginal abatement costs with respect to emissions $\hat{\beta}_1 = -0.609$ estimated in Table 4, column 4. The figure shows that there is wide variation in log emissions and that the fit of the iso-elastic MAC curve is quite good across this wide range.

Second, we test the joint significance of plant-period fixed effects, relative to a model with plant-period random effects. The data strongly reject the random effects model in favor of the plant-period fixed effects model (p-value < 0.001, column 4). This clear rejection supports the idea that heterogeneity across plant-periods is an important determinant of marginal costs, suggesting the potential for gains from trade in the emissions market (see Section 6).

Third and finally, Appendix Figure F1 supports our sample restriction to bids from the first half of each period. The grounds for this restriction are that plants can no longer trade-off own abatement against permit purchases once their emissions are sunk. Consistent with this idea, Appendix Figure F1 shows that permit bids are inelastic with respect to emissions when two weeks or fewer remain in the compliance period. This inelastic demand makes sense near the end of a period, because plants should then have a fixed, high willingness-to-pay to avoid the high penalties they will incur if they do not buy enough permits.

Heterogeneity in marginal abatement cost curves.—The estimates in Table 4 contain different MAC curves for each plant-period. To explore this heterogeneity, we calculate the fitted value of each plant's MAC curve over the range of emissions between a value close to zero and the high level of uncontrolled emissions, which the plant would emit if it did not run any air pollution control devices (see Appendix E.1 for how uncontrolled emissions are calculated). These fitted values exponentiate the estimates of equation (4) for each plant to relate the marginal abatement cost (rupees per kg) to the quantity of plant emissions (kg). Therefore, they depend on both the plant-period fixed effects, $\hat{\xi}_{it}$ and the elasticity of marginal abatement costs with respect to emissions $\hat{\beta}_1$. They allow for level differences in costs across plants but assume a common elasticity of marginal abatement costs with respect to emissions.

Figure 7 plots these curves for the plants in the treatment group in compliance period eight. The figure shows wide heterogeneity in the level of MACs at any constant emissions level. As an illustrative example, the figure includes a vertical line to indicate a policy requiring all plants to have the average plant emissions load. At this hypothetical standard, plants have a wide range of marginal abatement costs; for example, the ratio of marginal abatement costs for plants at the 75th percentile of costs to ones at the 25th percentile is 2.1 to 1. In practice, loads are not uniform in the control group due to imperfect compliance and the fact that the actual standard limits concentration, not load. The triangle markers on the plot mark the level of marginal abatement costs for each firm at a representative draw of emissions load from the control group distribution of emissions. As in the case of the strictly uniform standard, there is wide variation in marginal abatement costs across plants. The implication of this heterogeneity is that compliance costs will not be minimized at this allocation of emissions. To minimize costs, the regulator would need to know the MAC curves for all plants. Without such knowledge, command-and-control approaches to regulation will be relatively costly, since they will not equate marginal costs across plants.

Estimating changes in fixed abatement costs.—The goal of MAC estimation is to be able to compare abatement costs across regulatory regimes. We have estimated these curves using bid data only from the treatment plants. To use the treatment MAC curves to compare costs, we then need to assume that control plants have the same distribution of MACs. Ex ante, this will be true by design, as treatment plants are randomly assigned from the experimental sample of plants. However, it is possible that the treatment caused plants to change their MAC curve, for example, by investing in new pollution abatement capital. Such changes to fixed costs would not be reflected in bids but may cause treatment and control plants to have different ex post distributions of MAC curves.

We directly test the assumption that the treatment does not cause plants to change their pollution abatement capital using data collected via a phone survey in November 2020. In the survey, plants reported on abatement equipment costs and other costs for operating the boiler and related equipment (the boiler house includes the equipment that consumes fuel and therefore produces emissions). Table 5 reports plant-level treatment effect regressions of abatement capital costs on a treatment status indicator. Treatment has a small, negative and statistically insignificant effect on total abatement capital (column 1). This lack of an impact on abatement capital is consistent with the high rates of APCD installation in sample plants at baseline (Table 1). In the Surat market, as is the norm for other emissions markets, the market was imposed on top of a pre-existing equipment mandate.[31] Columns 6 to 10 consider other boiler house input costs including capital, labor, electricity, and fuel. Again, we find insignificant effects overall and no consistent pattern by component. We conclude that plant abatement capital did not change in the treatment group and that it is therefore valid ex post to treat the marginal abatement cost curves in the treatment and the control groups as being drawn from the same distribution.

6 The Gains from Trade and a Benefit-Cost Analysis

This section compares abatement costs under the market to those under the status quo command-and-control regime using the estimated MAC curves. Section 6.1 discusses the model's fit to realized market outcomes in the treatment group. Section 6.2 calculates the gains from trade in the market regime. Section 6.3 compares the market's costs and benefits using our cost and emissions estimates and external estimates of pollution damages. Appendix D details these steps.

[31] 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 SO_2 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.

6.1 Model Fit

This subsection calculates the market price that would prevail in a particular period, given the estimated plant-specific MAC curves (4). The purpose is to compare modeled market prices, which depend on our MAC estimates and plant conduct assumptions, with observed market prices.

Given the MAC curve estimates for all the plants in the market, we solve for the market-clearing permit price at any given cap Q_t. Define the function $E_{it}(P_t) = \widehat{MAC}_{it}^{-1}(P_t)$ as the inverse of the estimated MAC curve for plant i; that is, the level of emissions plant i would choose at a given market price. The equilibrium permit price is the price at which aggregate emissions equal the cap:

$$E_t(P_t^*) = \sum_i E_{it}(P_t^*) = Q_t. \tag{5}$$

The equilibrium price is unique because the MAC curves imply that emissions for each plant monotonically decrease in price. The resulting allocation of emissions E_{it} across plants is efficient because all plants set their marginal cost of abatement at E_{it} equal to the market price.

Figure 8 shows the modeled market prices P_t^*, calculated in this way, alongside the actual market prices. The dashed (black) line represents mean bids, the dotted (black) line represents mean clearing prices, and the solid (blue) line indicates market-clearing prices simulated by the model. The predicted market prices fit mean bid price and their fluctuations across periods well. Bids and simulated prices are relatively high in period 1, fall to ₹8 to ₹10 for periods $2-6$, and rise in the final four compliance periods to ₹10 to ₹12. Predicted prices are likewise similar to the market-*clearing* price, but the fit is less good. They are consistently above actual clearing prices in period $1-6$, with the gap narrowing in periods $7-10$. The initial gap in model fit and later convergence may reflect plants learning about compliance costs over time. Overall, we conclude that this simple model describes market prices well.

Model fit with alternative step function marginal cost.—Because the main tools for particulate abatement are distinct air pollution control devices, we may think that marginal abatement cost has a step function form, with the cost on each step corresponding to the operation of a different device. In Appendix E, we simulate market outcomes using this alternative, step-function model of marginal abatement costs. We find this model, fitted to the same bidding data as used above, fits observed market prices and emissions poorly. Market-clearing prices exceed actual prices (see Appendix Figure E2). Moreover, the step function technology implies a distribution of emissions that is lumpy and dispersed, relative to both the data and modeled outcomes with iso-elastic marginal costs (Appendix Figure E3). One interpretation of this poor fit is that the step function model takes an inflexible, all-or-nothing view about how plants run their abatement devices and rules out, in particular, that plants may have ways to adjust emissions even conditional on their abatement capital and technology mix. The remainder of the paper reports results using our preferred iso-elastic model of marginal abatement cost functions based on (4).

6.2 Gains from trade in the market regime

We start by illustrating the gains from trade using the estimated MAC curves for two plants. This illustration could be drawn from a textbook on the benefits of trading, except that it uses data from real plants participating in a live emissions market. Next, we describe the steps necessary to calculate the gains from trade in the market as a whole and report estimates of these gains.

Illustration of the potential for gains from trade.—To illustrate the mechanism for gains from trade in the emissions market, Figure 9 plots the fitted MAC curves for two plants in the market, using the Table 4, column (4) estimates. We give the plants the pseudonyms "Surat Polyfilm" (panel A) and "Mahadev Textiles" (panel B) for confidentiality. Marginal abatement costs (prices) are on the vertical axis and emissions (quantities) are on the horizontal axis.

In the example, we use MAC curve estimates drawn from compliance period 8 and we represent the command-and-control regime as a fixed load standard for all plants. The load standard, indicated by the dashed vertical line, is set to the average emissions load of 1,090 kg per plant. The dashed horizontal line at ₹11.25 per kg shows the market clearing price that sets total emissions in the hypothetical two-plant market equal to total emissions in the command-and-control regime. Below, we will broaden our estimates of the gains from trade to cover all plants and periods and to use more realistic distributions of emissions in the command-and-control regime.

The differing marginal abatement cost functions for the two plants create an opportunity for profitable trade. At the command-and-control load standard of 1,090 kg, Surat Polyfilm's marginal abatement cost is roughly ₹6 per kg (panel A) and Mahadev Textiles' about ₹14 per kg (panel B). This difference in marginal abatement costs means that the two plants could gain from trade: both plants would reduce their costs if Surat Polyfilms abates an additional kg of particulate matter and sells a permit to Mahadev Textiles at a price between ₹6 per kg and ₹14 per kg.

The movement from a fixed load standard to a market would cause large, non-marginal changes in plant emissions. Under the market, each plant emits until its MAC curve intersects the market-clearing price (equation 3). The market price here is ₹11.25 per kg, at which price Surat Polyfilms would cut its emissions by approximately 66%, to 366 kg, while Mahadev Textiles would increase its emissions by approximately 66%, to 1,814 kg. The reallocation of emissions reduces costs for both plants. Each plant's total variable abatement cost, depicted in pale red, is the area beneath the MAC curve from the chosen emissions level up to its maximum emissions level. When Surat Polyfilms reduces emissions from the command-and-control limit to 366 kg it incurs additional abatement costs, indicated by the cross-hatched area, but more than covers these costs with additional permit revenues, because its marginal abatement costs lie below the market price in this range. The shaded blue area above the MAC curve and below the price is its gain from trade in the market (₹2,533). On the other side of the market, Mahadev Textiles (panel B) increases its emissions, lowering abatement costs by the cross-hatched area beneath its MAC and purchasing 724 permits (=1,814 kg - 1,090 kg) at ₹11.25 per kg. Mahadev's costs fall because its marginal abatement costs exceed the price of a permit at all emissions levels between 1,090 kg and 1,814 kg; the shaded blue region depicts its gain from trade (₹1,279). Thus the transition to an emissions market reallocates emissions and reduces plant costs on both sides of the market.

Estimating total abatement costs in the market regime.— To quantify the market-level gains from trade, we calculate total abatement costs for all plants at the emissions levels that they would choose in both the market and command-and-control regime. The discussion of model fit (6.1) and the example above show how we do this calculation in the market regime. First, we invert each plant's MAC function to find its emissions at the market clearing price, as in (9). Second, we calculate plant variable abatement costs as the area underneath its MAC curve between its unconstrained emissions level and its chosen emissions, and add these costs across plants. We repeat this calculation for a full range of potential market caps.

The blue line in Figure 10 depicts total variable abatement costs across all plants in the market regime. At each market cap, this curve is calculated as the sum of all plants' total variable abatement costs for compliance period 8, as shown in Figure 7, over market caps ranging from near zero to 300 metric tons. The vertical dashed line at left depicts emissions at 170 tons, 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 aggregate abatement cost curve.

The most striking finding about the market-level abatement cost curve is that abatement costs rise only slowly in response to reductions in pollution. For example, 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 4, because total abatement costs include inframarginal abatement costs and are therefore less elastic to emissions than marginal costs. The cost estimates indicate that substantial improvements to Gujarat's air quality are available for relatively small increases in plant abatement costs.

Estimating total abatement costs in the command-and-control regime.— To compare costs across regimes, we also need to estimate the total abatement costs that treatment plants would have incurred under the status quo, command-and-control regime. We estimate these costs using two observations. First, because the market was introduced in an experiment, the distribution of MAC curves we have estimated for the treatment plants will be the same as that for the control plants. Second, we can use the control distribution of emissions levels to represent the stringency of the command-and-control regime. We therefore evaluate costs in the command-and-control regime by evaluating treatment cost *functions* at control emissions *levels*.

This approach to the command-and-control regime treats the world as it is, rather than applying a single uniform standard, as in our simple example of the gains from trade. Roughly 30% of plants in the control group exceed the *de jure* emissions limits in the command-and-control regime (Figure B1). We therefore model the command-and-control regime not as a single standard, as on paper, but with a distribution of possible emissions levels. We simulate many sets of draws of emissions levels for all plants in the market and evaluate the aggregate costs of abatement for each set by calculating the area under each plants' MAC curve, from the emissions draw up to its unconstrained emissions level, and summing these costs across plants. We consider several alternative ways to draw emissions; our preferred method draws residual emissions levels conditional on plant heat capacity. Appendix D.2 explains how we conduct these simulations.

Returning to Figure 7, the triangle markers on each marginal abatement cost curve illustrate one draw from the distribution of control emissions levels for all plants. The figure shows that, at this draw, plant emissions load is not uniform across plants and that the implied marginal abatement costs vary widely across plants.

To isolate the gains from trade we seek to compare abatement costs at the same aggregate emissions level, though treatment plants emitted less than control plants on average (Table 3). For each candidate aggregate emissions cap in the market, therefore, we scale the distribution of control emissions loads up or down by a common factor to match the cap. For example, consider a plant that had estimated emissions of 1,090 kg when aggregate emissions across all plants is 240 tons. We assume that if aggregate emissions were cut to 120 tons this plant would emit 545 kg. This approach assumes that changes in stringency in the command-and-control regime change emissions at all plants proportionally.

Now returning to Figure 10, the black line depicts total variable abatement cost as a function of aggregate emissions under the command-and-control (higher, black line) regime (again using compliance period 8 estimates). At each point in the curve, the total variable abatement cost is the sum across all plants, and costs for each plant are estimated with $S = 100$ simulation draws of emissions (residualized on plant heat capacity) from the control group. One can visualize the total cost, in a single simulation draw, as the sum of areas under the MAC curves in Figure 7, down to the emissions draw for each plant.

The comparison in Figure 10 for period 8 shows that the market regime lowers total variable abatement costs, relative to the command-and-control regime, at any level of emissions. At the control emissions of 240 tons, the market reduces total variable abatement costs by 10% (moving down the vertical dashed line). Because total abatement costs are not very elastic with respect to emissions, the emissions market would cut total emissions by 43% (moving left along the horizontal dashed line) at the same variable abatement costs as in the status quo. 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.

We estimate, across all plants and compliance periods, that the actual market outcome falls in this range of reducing *both* costs and emissions at the same time. Appendix Table D1, panel A summarizes the results on abatement costs. 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 2, row A), so that the market cuts costs by 11%. 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. We report similar results using different methods to draw counterfactual emissions levels from the control group, with or without conditioning on observables, for example. We find that the cost differences among the alternative representations of the command and control regime are small both in absolute terms and relative to the difference in cost between the market and command-and-control regimes.

6.3 Benefit-cost Analysis of Emissions Market Expansion

The preceding analysis has quantified the benefits of the market in terms of emissions reductions and abatement costs. Table 6 uses these estimates as inputs for a benefit-cost analysis of expanding the emission market to all 906 industrial plants that burn solid fuel in Surat. We compare the benefits of lower ambient concentrations of particulate matter increasing life expectancy, valued in monetary terms, against the emission market's operational and abatement costs. All comparisons take as given the existing stock of abatement capital in Surat (Appendix G discusses the inputs in depth).

Table 6 reports on the benefit-cost analysis for emissions reductions of 10%, 30% and 50% in columns 1 to 3, respectively. The first two reductions are within the range of our experimental data, with the 30% reduction roughly equal to the experimental treatment effect documented above. The third extrapolates outside that range, using the emissions market's total abatement cost function, founded on our estimates of individual MAC functions.

Table 6, Panel A shows the costs of expanding the market, which are comprised of the fixed costs of CEMS monitoring and changes in total variable abatement costs. Improved monitoring is a necessary condition to start an emissions market. The administrative data record that it costs roughly \$5,000 per plant-year to set up a monitoring system. This cost is fixed with respect to the targeted emissions reduction. For 10% and 30% emissions reductions, these additional costs are counterbalanced by estimated *savings* in abatement costs of \$1,242 and \$648 per plant-year, respectively, associated with switching to the market. Even at a 50% emissions reduction, the savings from moving to the market nearly cancel out the additional abatement costs from further pollution abatement; on average, plant abatement costs increase by just \$77 annually. Scaling these costs to cover all eligible plants in Surat yields aggregate annual costs of \$3.4 million, \$3.9 million, and \$4.6 million for 10%, 30%, and 50% reductions in emissions, respectively.

On the benefits side, we consider only the mortality benefit of lower pollution extending people's lives, ignoring other possible sources of benefits (better health, productivity improvements, etc.). We convert emissions reductions into a monetary value using three factors: how emissions change ambient pollution (panel B), the value of increased lifespans (panel C), and how lower ambient pollution lengthens lifespans (panel D). The first factor scales the 10%, 30%, and 50% emissions reductions from industry by the industrial share of ambient PM2.5 concentrations, drawn from the best available source apportionment study (Guttikunda, Nishadh and Jawahar, 2019). This scaling yields estimated reductions in ambient PM2.5 of 2.8 $\mu g/m^3$, 8.5 $\mu g/m^3$, and 14.2 $\mu g/m^3$, respectively (panel B). The next factor is the gain in life years in Surat for each cut in ambient pollution. We estimate this relationship with the elasticity of life expectancy with respect to ambient PM2.5 from Ebenstein et al. (2017) and report robustness to other estimates (panel E). The third factor is the value of a statistical life (VSL), the willingness-to-pay for a reduction in mortality risk, which we set at $665,000 ($9,500 per statistical life-year), following the value chosen by Nair et al. (2021) (panel C). The estimated benefits of the three emissions reductions scenarios are then the monetized values of life-years saved for each year that the ETS is in force, $282 million, $847 million, and $1,412 million, across the three respective emissions reductions.

Panel E reveals that the estimated benefits of scaling the market to all eligible plants greatly exceed the estimated costs. Row 1 uses the mortality dose-response estimated in Ebenstein et al. (2017) and rows 2 through 4 use alternative estimates. For a 30% reduction in industrial emissions (column 2), the benefit to cost ratio ranges from roughly 25:1 to more than 200:1, depending on the estimate of the elasticity of life expectancy with respect to ambient PM2.5 (across rows within panel E). Because a large share of the costs of the market are fixed costs of monitoring, and abatement costs are only moderately convex, the benefit-cost ratio rises with greater reductions in emissions over this range.

7 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 new regulation 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%, holding emissions constant. More broadly, we estimate that emissions can be reduced without capital investment and 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.

Our experimental setting allowed for low-cost abatement in part because the existing regime mandated the installation of air pollution abatement equipment but could not adequately enforce its use. This situation, where a market is layered on top of a command-and-control mandate, is a frequent starting point for introducing markets (see footnote 31). We expect that longer-run changes may augment the efficiency and environmental benefits estimated in this experiment, for two reasons. First, the efficiency gains from other emissions markets have been attributed largely to reallocation towards the most cost-effective kinds of abatement capital (Fowlie, 2010; Chan et al., 2018; Colmer et al., 2023), whereas the gains in efficiency we estimate are due only to reductions in variable abatement costs. Second, regulators often use the lever of a market cap to tighten standards over time. Indeed, the GPCB started such a process during our experiment by tightening the cap when the market revealed that aggregate abatement costs were relatively low.

We believe that this proof-of-concept for an emissions market has broad policy relevance. Based on a review of the performance of the treatment market in Surat, GPCB decided to expand its scope. Surat's control plants were included in the market in September, 2022. In September 2023, the GPCB launched a second particulate emissions market for plants in Ahmedabad, Gujarat's largest city and major industrial hub. Currently, the GPCB is exploring expanding the market to additional industrial clusters and pollutants. The Maharashtra Pollution Control Board (MPCB) has started the development of a statewide market for sulfur dioxide emissions from thermal power plants and other industrial sources. The coauthors of this paper are advising MPCB and are in discussions with several other Indian states on how to use environmental markets.

The larger question remains why environmental quality is so poor in many developing countries. Our results suggest that industrial air pollution remains high not because of high abatement costs, at the level of individual plants, but due to high fixed costs of monitoring and enforcement. State capacity may limit the use of sophisticated regulatory instruments, but 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 pollution while *lowering* abatement costs.

Pollution markets, as a policy tool, are not exclusively reserved for high-income countries. The Gujarat evidence shows a market shifting abatement to firms with low abatement costs, just as markets have done in rich countries. In developing countries, markets can also have an additional benefit, of helping to solve the first-order problem of low compliance with environmental regulation, in any form. Given pollution's high costs to human well-being in these countries, emissions markets have great potential to raise both environmental quality and economic growth.

References

Apte, Joshua S., Michael Brauer, Aaron J. Cohen, Majid Ezzati, and C. Arden Pope III. 2018. "Ambient PM2.5 Reduces Global and Regional Life Expectancy." *Environmental Science and Technology Letters*, 5(9): 546–551.

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*.

Chan, H Ron, B Andrew Chupp, Maureen L Cropper, and Nicholas Z Muller. 2018. "The impact of trading on the costs and benefits of the Acid Rain Program." *Journal of Environmental Economics and Management*, 88: 180–209.

Chu, Cindy. 2009. "Thirty years later: the global growth of ITQs and their influence on stock status in marine fisheries." *Fish and Fisheries*, 10(2): 217–230.

Coase, Ronald H. 1960. "The Problem of Social Cost." *The Journal of Law & Economics*, 3: 1–44.

Colmer, Jonathan, Ralf Martin, Mirabelle Muûls, and Ulrich J Wagner. 2023. "Does pricing carbon mitigate climate change? Firm-level evidence from the European Union emissions trading scheme." *CRC TR 224, Discussion Paper*.

Correia, Andrew W, C Arden Pope III, Douglas W Dockery, Yun Wang, Majid Ezzati, and Francesca Dominici. 2013. "The effect of air pollution control on life expectancy in the United States: an analysis of 545 US counties for the period 2000 to 2007." *Epidemiology (Cambridge, Mass.)*, 24(1): 23.

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.

Dechezleprêtre, Antoine, Daniel Nachtigall, and Frank Venmans. 2023. "The joint impact of the European Union emissions trading system on carbon emissions and economic performance." *Journal of Environmental Economics and Management*, 118: 102758.

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.

Duflo, 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).

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan. 2013. "Truth-telling by Third-party Auditors and the Response of Polluting Firms: Experimental Evidence from India." *The Quarterly Journal of Economics*, 128(4): 1499–1545.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan. 2018. "The Value of Regulatory Discretion: Estimates From Environmental Inspections in India." *Econometrica*, 86(6): 2123–2160.

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, and Barbara K Buchner. 2008. "Over-allocation or abatement? A preliminary analysis of the EU ETS based on the 2005–06 emissions data." *Environmental and Resource Economics*, 41: 267–287.

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. 2010. "Emissions trading, electricity restructuring, and investment in pollution abatement." *American Economic Review*, 100(3): 837–69.

Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur. 2012. "What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NOx Trading Program." *American Economic Review*, 102(2): 965–93.

Greenstone, Michael, and Christina Hasenkopf. 2023. "Air Quality Life Index Annual Update." Energy Policy Institute at Chicago.

Greenstone, Michael, and Rema Hanna. 2014. "Environmental Regulations, Air and Water Pollution, and Infant Mortality in India." *American Economic Review*, 104(10): 3038–72.

Greenstone, Michael, and Ted Gayer. 2009. "Quasi-experimental and experimental approaches to environmental economics." *Journal of Environmental Economics and Management*, 57(1): 21–44.

Greenstone, Michael, Guojun He, Ruixue Jia, and Tong Liu. 2022. "Can Technology Solve the Principal-Agent Problem? Evidence from China's War on Air Pollution." *American Economic Review: Insights*, 4(1): 54–70.

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.

Jayachandran, Seema. 2022. "How Economic Development Influences the Environment." *Annual Review of Economics*, 14(1): 229–252.

Klier, Thomas H, Richard H Mattoon, and Michael A Prager. 1997. "A Mixed Bag: Assessmentof Market Performance and Firm Trading Behaviour in the NOx Reclaim Programme." *Journal of Environmental Planning and Management*, 40(6): 751–774.

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.

Ryan, Nicholas, and Anant Sudarshan. 2022. "Rationing the commons." *Journal of Political Economy*, 130(1): 210–257.

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.

Shapiro, Joseph S, and Reed Walker. 2020. "Is Air Pollution Regulation Too Stringent?" National Bureau of Economic Research Working Paper 28199.

Stavins, Robert N. 2003. "Experience with Market-Based Environmental Policy Instruments." Elsevier Handbook of Environmental Economics.

Sudarshan, Anant. 2023. "Introduction to Development Engineering." In . , ed. Temina Madon, Ashok J. Gadgil, Richard Anderson, Lorenzo Casaburi, Kenneth Lee and Arman Rezaee, 161–182. Springer International Publishing.

Zhou, Xuehua, Zhaoyu Cao, Yujie Ma, Linpeng Wang, Ruidong Wu, and Wenxing Wang. 2016. "Concentrations, correlations and chemical species of PM2. 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.

8 Figures

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

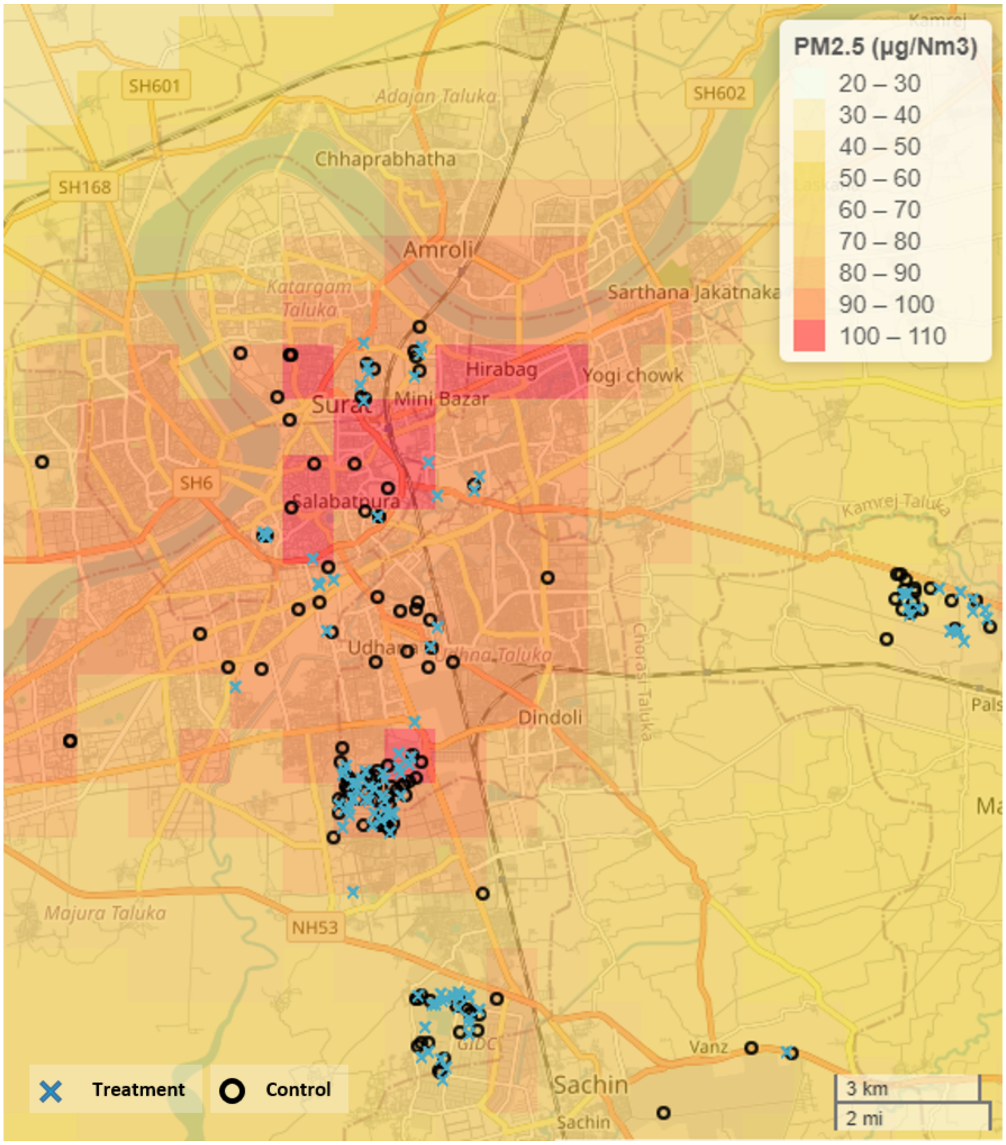

The figure shows ambient PM2.5 $\mu g/m^3$ concentrations in Surat, Gujarat averaged over the year 2018, overlaid with the locations of sample plants. The ambient pollution data is from Guttikunda, Nishadh and Jawahar (2019). As a basis for comparison, India's National Ambient Air Quality Standard for PM2.5 is 40 $\mu g/m^3$ and the WHO standard is 5 $\mu g/m^3$. 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

A. Permit prices

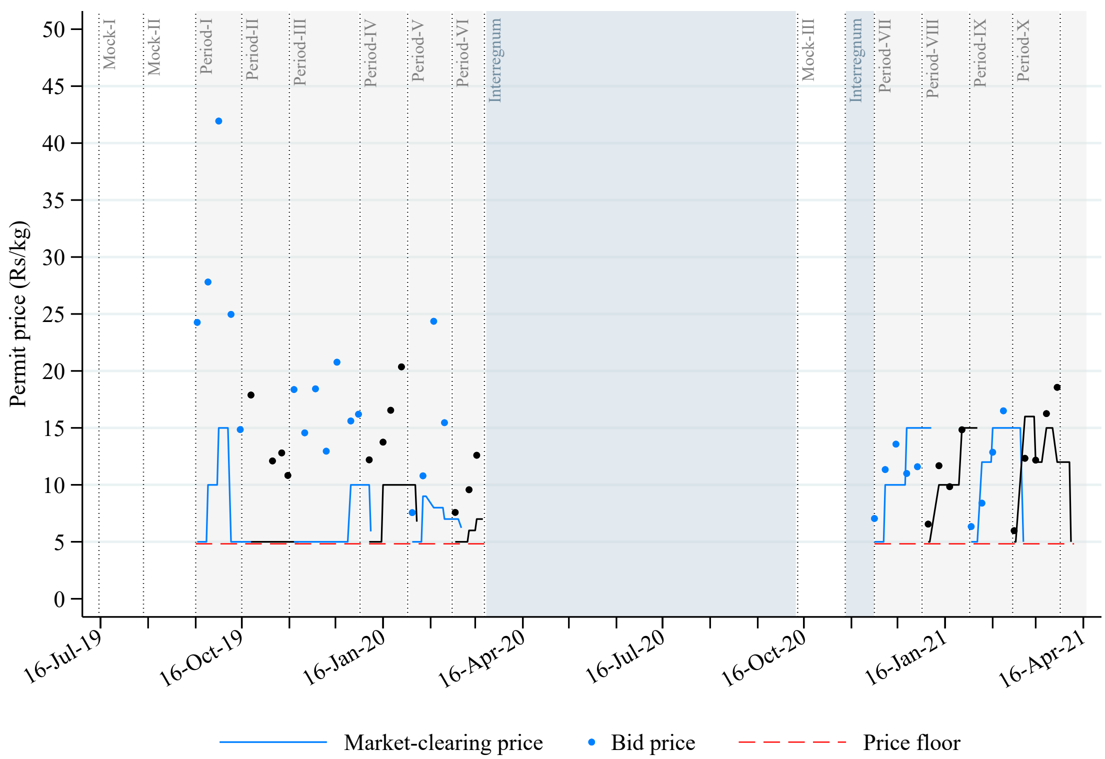

B. Permit quantites

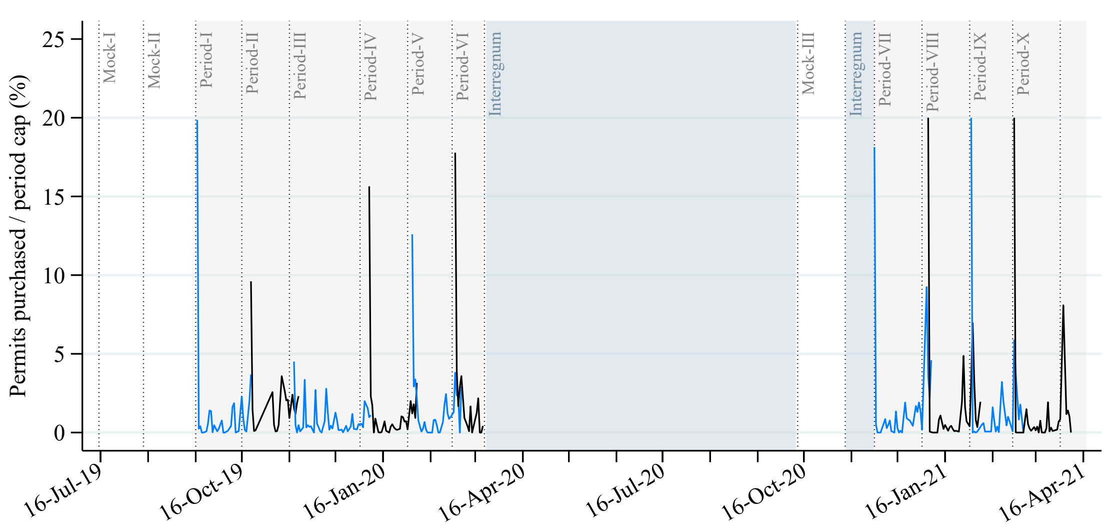

This figure shows weekly permit prices (panel A) and quantities (panel B) from September 2019 to April 2021. In panel A, the scattered points are the mean bid prices (both sale and purchase) and the solid line the market-clearing price. Since permits of different vintages, from two consecutive compliance 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 ₹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 Initial Permit Allocation by Compliance Period

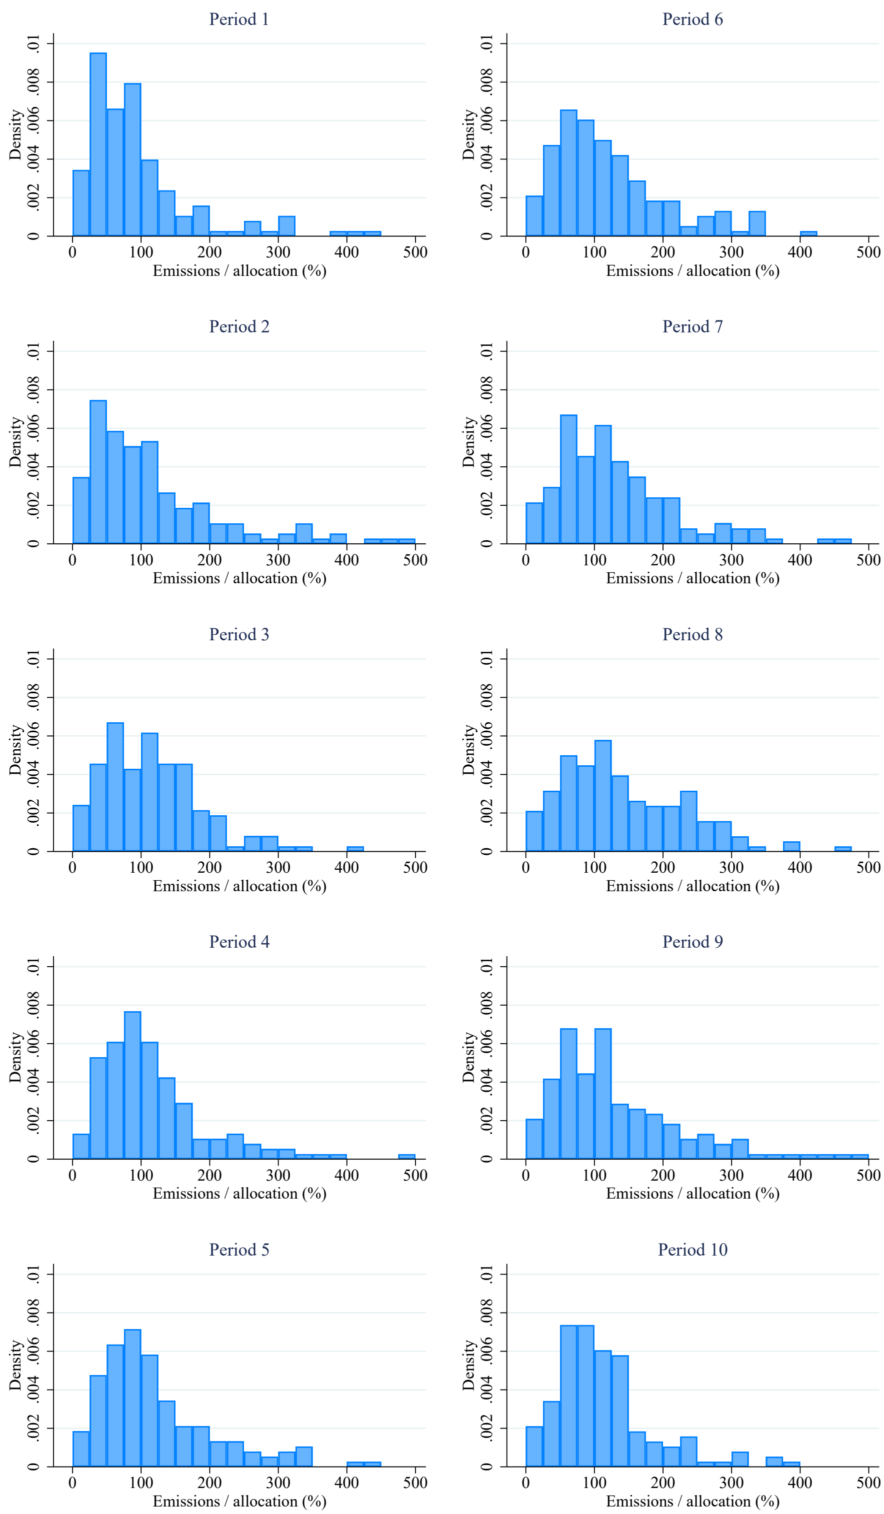

This figure plots the distributions of (emissions / initial permit allocation $\times$ 100%) across treated plants ($N = 156$) by compliance period, truncated at the 97.5th percentile. Emissions data and permit holdings are from the administrative records of the market operator. Emissions are the validated 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 4: Distribution of Emissions over Final Permit Holdings by Compliance Period

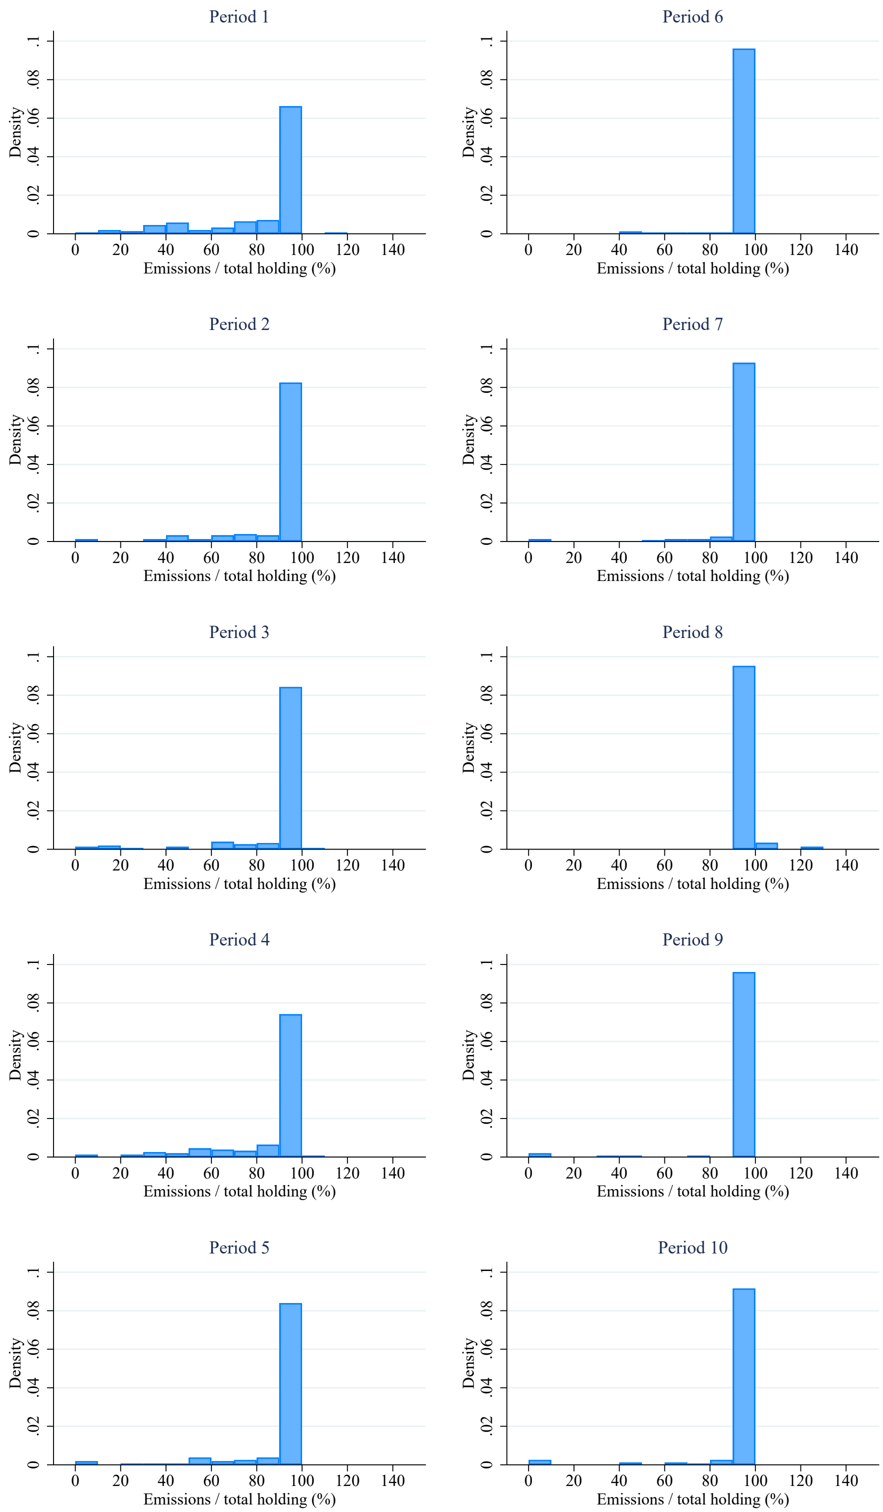

This figure plots the distributions of (emissions / final permit holdings $\times$ 100%) across treated plants ($N = 156$) by compliance period, truncated at the 99.5th percentile. Final permit holdings are the total number of permits a plant held at the end of the true-up period after each compliance 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 5: PM Emissions by Treatment Status

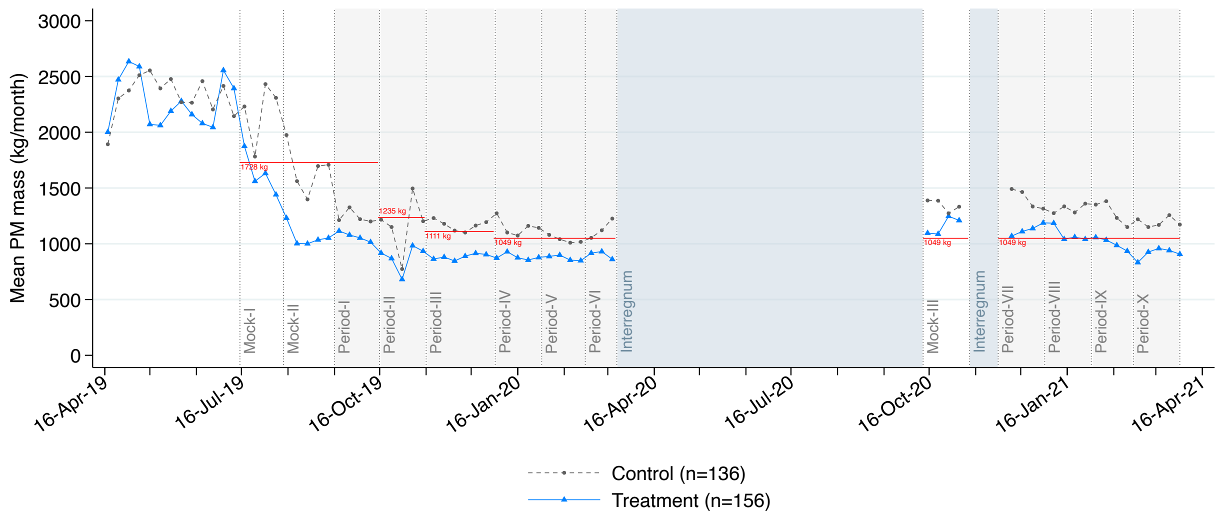

The figure shows the weekly mean plant PM emissions in kilograms (calculated at a monthly rate equivalent) from April 2019 to March 2021 by treatment status. The treatment group is represented by the solid (blue) line and the control group by the dashed (grey) line. The grey regions mark the ten compliance 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 C2). 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: Estimation of Marginal Abatement Cost Elasticity with respect to Emissions

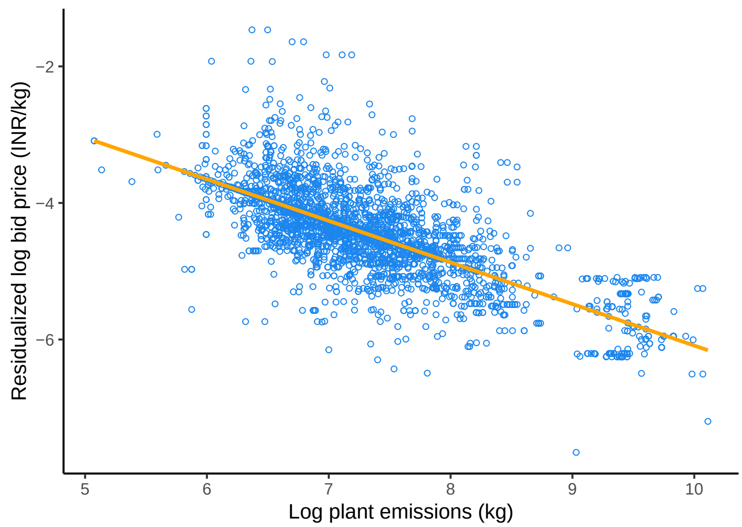

This figure visualizes the estimation of the marginal abatement cost elasticity with respect to emissions, as specified in equation 4. The data are restricted to 3,120 bids offered by all plants in the first halves of all compliance periods. The vertical axis is the log bid price residualized on plant-period fixed effects. The horizontal axis are the log plant emissions that would result if a bid was executed. The linear fit shows the iso-elastic curve of best fit for the marginal abatement cost curve.

Figure 7: Marginal Abatement Cost Curves for Treatment Plants

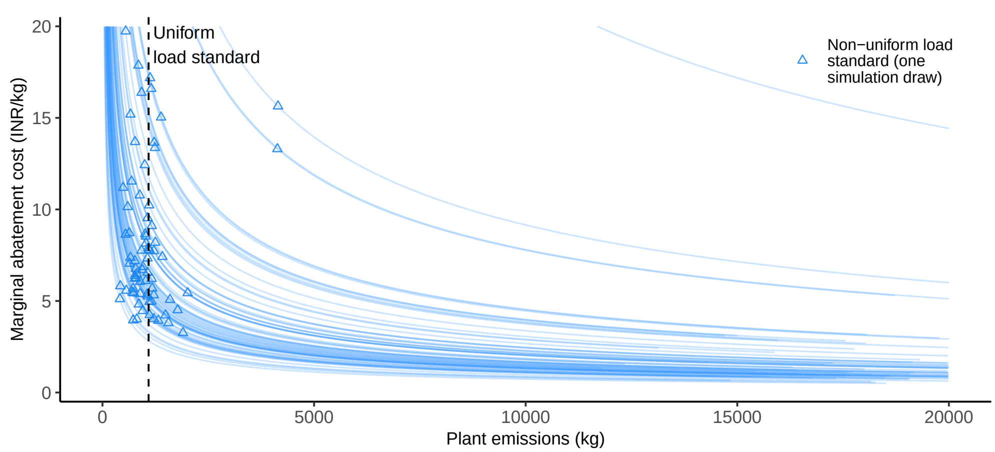

The figure illustrates estimated marginal abatement cost curves for all plants that bid in period 8. The domain of each curve extends upward to the uncontrolled emissions level for each plant. The triangles correspond to plant emissions under one simulation of the counterfactual command-and-control regime in which emissions rate are allowed to vary with plant capacity along with an idiosyncratic error term. This corresponds to the regulatory regime simulated in Table D1 Panel A Row B4.

Figure 8: Model Fit to Market-Clearing Prices

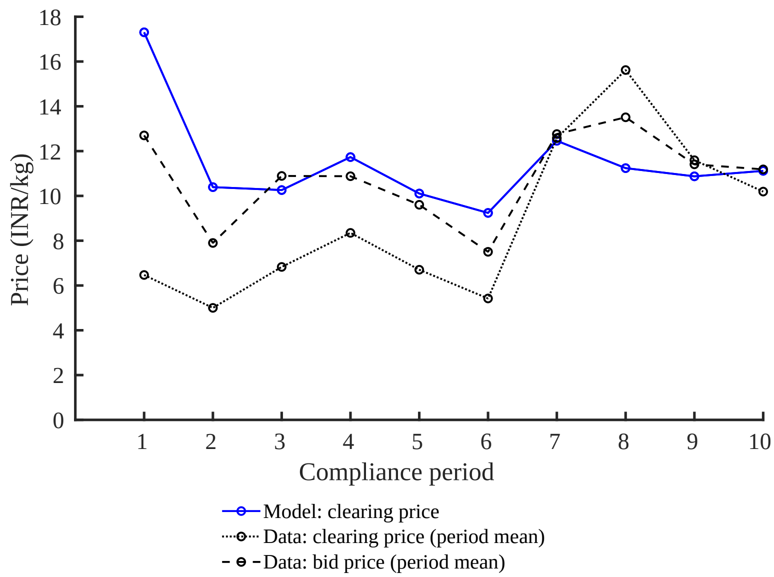

The figure shows the fit of the model to the time series of market and bid prices by compliance period. The solid (blue) line is the time series of market-clearing prices in the fitted model. The model is fit based on bids in the first half of each compliance period. The dashed (black) line is the time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

Figure 9: An Example of the Gains from Trade in the Market

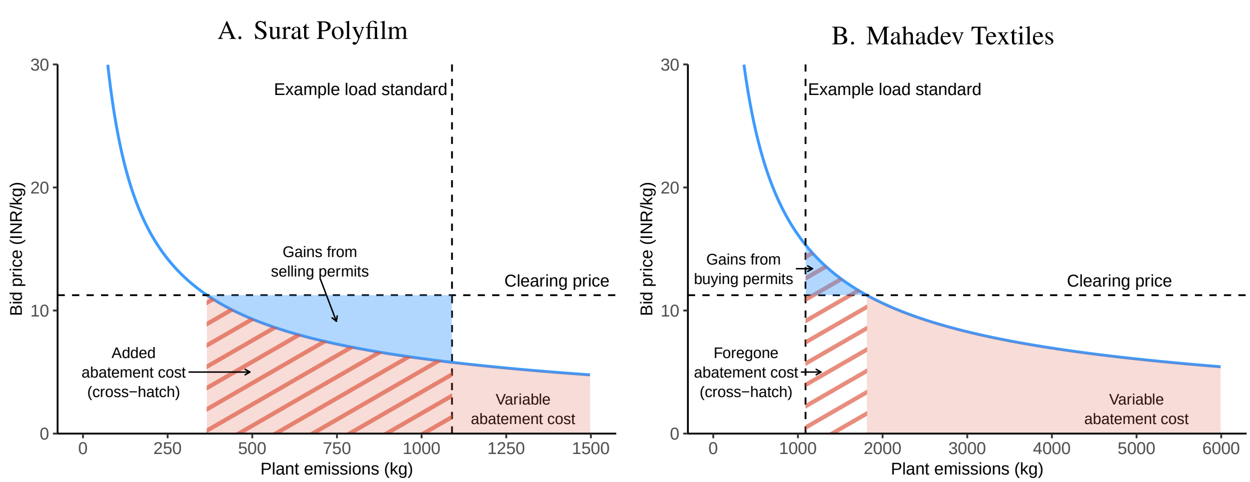

The figure visualizes gains from trade on the estimated MAC curves for "Surat Polyfilm" (panel A) and "Mahadev Textiles" (panel B), both pseudonyms. The MAC curves are fit as seen in Figure 7. The vertical dashed line gives a hypothetical load standard set at the average emissions per plant-month, and the shading shows how trading permits allows for price savings for both permit buyers and sellers compared to a command-and-control regime.

Figure 10: Variable abatement costs by regime

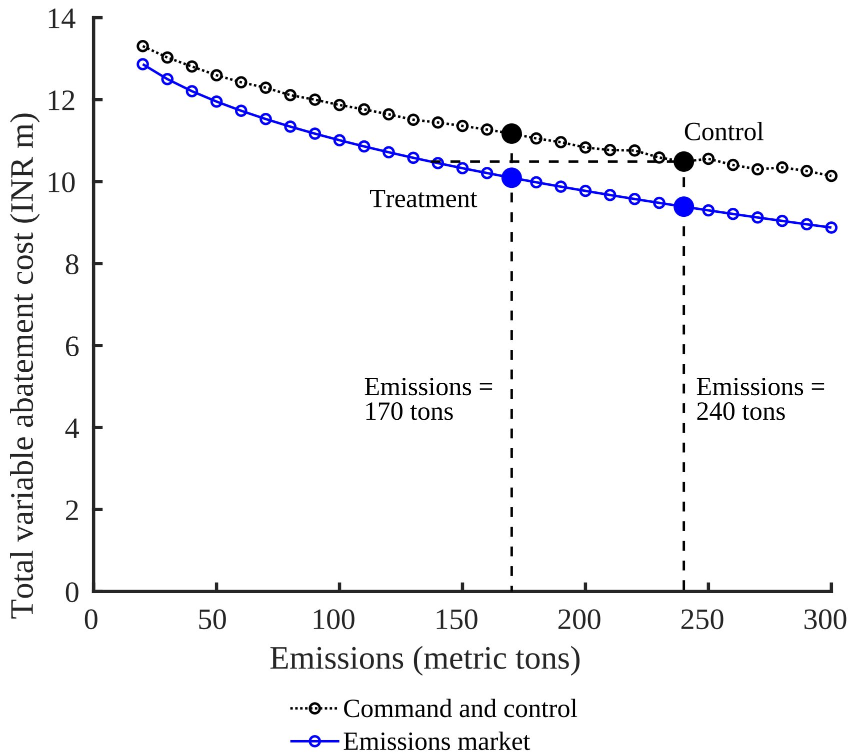

The figure shows the total (not marginal) variable abatement costs by regulatory regime as estimated for compliance period 8. The dotted (black) curve shows the total variable abatement cost curve under command-and-control and the solid (blue) curve under the emissions market. The command-and-control regime uses a capacity-based emissions rate with error to set emissions targets for each plant, as described in Section 5. 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.

9 Tables

Table 1: Balance of plant characteristics by treatment status

	Treatment (1)	Control (2)	Difference (3)
Panel A: Plant Measures			
Total electricity cost (1,000 USD)	467.6	345.8	121.9
	[869.0]	[327.0]	(78.5)
Log(plant total heat output)	15.6	15.6	0.012
	[0.62]	[0.50]	(0.065)
Size as recorded on environment consent (1 to 3)	1.37	1.37	0.0052
	[0.64]	[0.62]	(0.075)
Small-scale (size=1)	0.72	0.71	0.0063
	[0.45]	[0.46]	(0.054)
Large-scale (size=3)	0.086	0.075	0.011
	[0.28]	[0.26]	(0.032)
Number of stacks	1.08	1.04	0.035
	[0.41]	[0.21]	(0.038)
Textiles sector (=1)	0.85	0.87	-0.025
	[0.36]	[0.33]	(0.041)
Gross Sales Revenue in 2017 (1,000 USD)	13010.5	9755.0	3255.5
	[43534.7]	[39788.2]	(5043.2)
Panel B: Plant Abatement and Investment Cost			
Boiler house employment	36.9	32.3	4.62
	[32.9]	[29.4]	(3.69)
Boiler house capital expenditure (1,000 USD)	199.9	171.4	28.5
	[405.0]	[196.6]	(38.3)
Boiler house operating cost (1,000 USD)	140.4	112.4	28.0
	[206.3]	[84.2]	(18.3)
APCD: Cyclone present	0.98	0.97	0.0100
	[0.14]	[0.17]	(0.019)
APCD: Bag filter present	0.80	0.88	-0.079*
	[0.40]	[0.33]	(0.043)
APCD: Scrubber present	0.64	0.61	0.030
	[0.48]	[0.49]	(0.058)
APCD: ESP present	0.12	0.075	0.045
	[0.33]	[0.26]	(0.035)
Panel C: Plant Pollution Measures			
Plant total PM mass rate (kg/hr)	3.62	3.60	0.027
	[4.94]	[3.82]	(0.52)

Plant mean PM concentration (mg/Nm^3)	179.0	168.8	10.2
	[156.1]	[150.2]	(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

			Bid quantities (kg)			Bid prices (INR/kg)		
Period	Count	Bids per plant	All bids	Buy bids	Sell bids	All bids	Buy bids	Sell bids
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1	1,525	9.8	557	428	620	12.70	10.03	14.00
		(10.4)	(1,172)	(604)	(1,361)	(16.65)	(14.47)	(17.47)
2	600	3.9	474	475	474	7.90	6.39	8.43
		(4.1)	(747)	(535)	(810)	(10.72)	(9.16)	(11.18)
3	1,084	7.0	329	344	323	10.89	9.02	11.67
		(6.7)	(533)	(600)	(502)	(11.36)	(12.24)	(10.89)
4	806	5.2	323	332	319	10.88	7.90	12.03
		(6.1)	(551)	(483)	(575)	(9.37)	(6.77)	(9.96)
5	767	4.9	376	449	350	9.60	6.49	10.72
		(7.0)	(515)	(559)	(496)	(10.80)	(1.74)	(12.36)
6	296	1.9	463	533	426	7.50	5.84	8.38
		(3.2)	(558)	(559)	(556)	(6.33)	(3.11)	(7.34)
7	646	4.1	400	468	325	12.76	10.29	15.46
		(4.5)	(533)	(580)	(466)	(6.55)	(4.81)	(7.11)
8	783	5.0	418	501	249	13.51	12.56	15.47
		(6.3)	(588)	(671)	(298)	(13.92)	(16.53)	(4.96)
9	962	6.2	353	397	257	11.41	9.89	14.78
		(9.7)	(423)	(458)	(314)	(6.65)	(7.00)	(4.15)
10	964	6.2	383	428	336	11.18	8.40	14.04
		(8.8)	(533)	(532)	(531)	(9.10)	(5.65)	(10.92)
Total	8,433	54.1	412	430	399	11.25	9.47	12.52
		(51.0)	(708)	(565)	(795)	(11.56)	(10.50)	(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 Imputed Months				Imputed Months			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ETS Treatment (=1)	-0.178**	-0.193**	-0.177**	-0.194**	-0.282***	-0.282***	-0.316***	-0.316***
	(0.076)	(0.076)	(0.075)	(0.075)	(0.074)	(0.075)	(0.057)	(0.057)
Year-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 year-month 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: Elasticity of marginal cost with respect to emissions

	log(Bid price)				
	(1)	(2)	(3)	(4)	(5)
log(Emissions as bid)	-0.100	-0.143**	-0.269***	-0.609***	
	(0.061)	(0.063)	(0.084)	(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**	0.138***			
	(0.041)	(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	146	146	142
Observations	3120	3120	3120	3120	3082

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. To avoid biasing SEs singleton FEs are dropped during estimation, and as a result the elasticity estimates in columns (4) and (5) are based on 2,775 and 2,753 non-singleton bids respectively, though in the table we list the full set of bids and plants for which we have data. 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 5: Treatment effects on abatement costs in survey data

	Abatement capital costs ($1000s)					Boiler house input costs ($1000s)					
	All	Cyclone	Bag	Scrubber	ESP	Total	Capital	Labor	Electricity	Fuel	Materials
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
ETS Treatment (=1)	-3.467	0.602**	0.530*	-0.222	-4.281	11.26	-7.178	1.561	25.21*	26.87*	-0.142
	(3.089)	(0.266)	(0.318)	(0.407)	(3.344)	(26.31)	(19.05)	(3.332)	(13.53)	(15.35)	(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 6: Benefit-Cost Analysis of Scaled-up ETS in Surat

	Emissions reduction				
	10%	30%	50%	Units	Source
	(1)	(2)	(3)	(4)	(5)
Panel A. Annual costs from scale-up of emissions trading					
1. Monitoring costs per plant	5,000	5,000	5,000	\$/year	Author estimates
2. Abatement cost Δ per plant	−1,242	−648	77	\$/year	Author estimates
3. Total costs	3.4	3.9	4.6	\$m/year	= (A1 + A2)×906
Panel B. Reduction in pollution					
1. Ambient $PM_{2.5}$ conc.	88.5	88.5	88.5	$\mu g/m^3$	Guttikunda et al. (2019)
2. Industry share	0.3	0.3	0.3		Guttikunda et al. (2019)
3. Reduction in industry $PM_{2.5}$	10	30	50	%	
4. Reduction in ambient $PM_{2.5}$	2.8	8.5	14.2	$\mu g/m^3$	= B1×B2×B3
Panel C. Gain in life-years					
1. Mortality impact	0.1	0.1	0.1	years/$(\mu g/m^3)$	Ebenstein et al. (2017)
2. Gain in life expectancy	0.3	0.8	1.4	years	= C1×B4
3. Life expectancy	70	70	70	years	World Bank
4. Per year of ETS	0.004	0.012	0.020	years	= C2 / C3
5. Population	7.5	7.5	7.5	m	World Pop. Review
6. Total gain in life-years	29,736	89,208	148,680	years	= C4×C5
Panel D. Value of gain in life-years					
1. Value of statistical life	665,000	665,000	665,000	\$	Nair et al. (2021)
2. Value of 1 year of life	9,500	9,500	9,500	\$/year	= D1 / C3
3. Value of gain in life-years	282	847	1,412	\$m	= C6×D2
Panel E. Benefit-cost ratio					
1. Ebenstein et al. (2017)	83:1	215:1	307:1		= D5 / A3
2. Pope, Ezzati and Dockery (2009)	52:1	134:1	191:1		
3. Correia et al. (2013)	30:1	77:1	110:1		
4. Apte et al. (2018)	10.0:1	26:1	38:1		

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 D1 Panel I. 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 PM_{10} concentrations. We convert their findings using a 0.65 $PM_{2.5}$-to-PM_{10} ratio (Zhou et al., 2016). Columns 1 to 3 show the benefit-cost analysis for reductions in pollution of 10%

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 F gives additional empirical results to support those in the main text. Appendix G 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 incompetely 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 (1)	Control (2)	Difference (3)
Panel A: Plant Measures			
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)
Gross Sales Revenue in 2017 (1,000 USD)	12614.5	13628.3	-1013.8
	[42698.0]	[53258.9]	(5680.6)
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. The survey consists of three main sections: a general section, a technical section, and a pollution 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. Of 304 plants covered in the technical survey 289 have only a single stack. 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 Gujurat 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 mg/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. A total of 66% of plants are in compliance at baseline in the treatment group and 72% in the control group.

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

Figure B1: Distribution of Pollution Before the Experiment

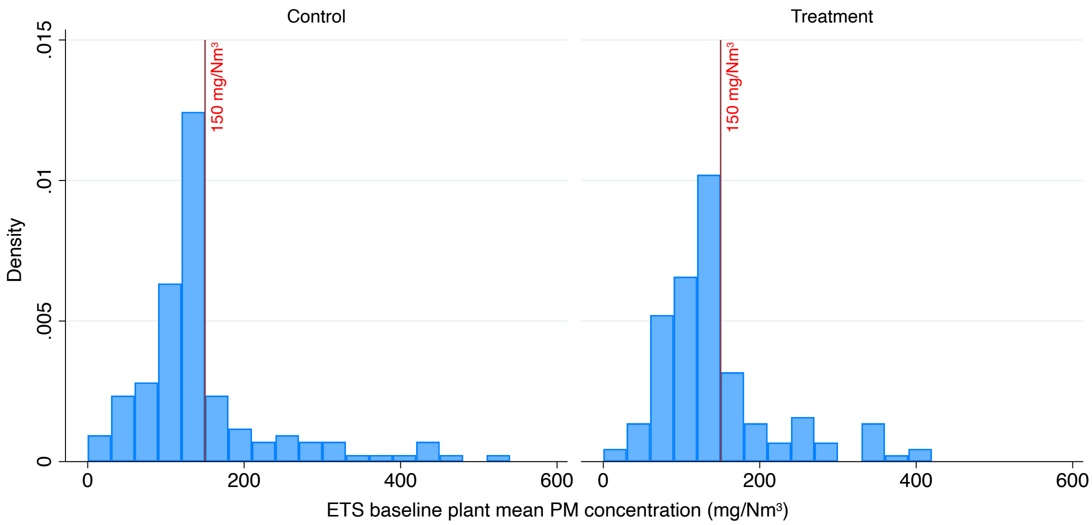

This figure shows the distributions of the plant PM concentration by treatment status as measured by manual iso-kinetic stack sampling at the ETS baseline (December 2018 to January 2019). One PM sample was collected from each industrial stack by a third-party laboratory. The histograms are truncated at the 95th percentile (520 mg/Nm^3). The red, vertical lines indicate the regulatory concentration standard of 150 mg/Nm^3. 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.

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 B1 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 B1: Trading data summary statistics

	All	Purchase	Sale
Panel A: Order			
Order quantity (kg)	411.61	429.50	398.78
	(707.98)	(565.09)	(794.52)
Order price (Rs/kg)	11.25	9.47	12.52
	(11.56)	(10.50)	(12.10)
Order price (Rs/kg), weighted by quantity	9.23	8.42	9.86
	(8.49)	(8.71)	(8.27)
Observations	8433	3520	4913
Panel B: Trade			
Trade quantity (kg)	360.42	389.36	327.18
	(564.45)	(544.09)	(585.37)
Trade price (Rs/kg)	9.29	9.19	9.40
	(7.39)	(9.31)	(4.20)
Trade price (Rs/kg), weighted by quantity	8.40	8.16	8.73
	(6.17)	(7.26)	(4.21)
Observations	3775	2018	1757

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

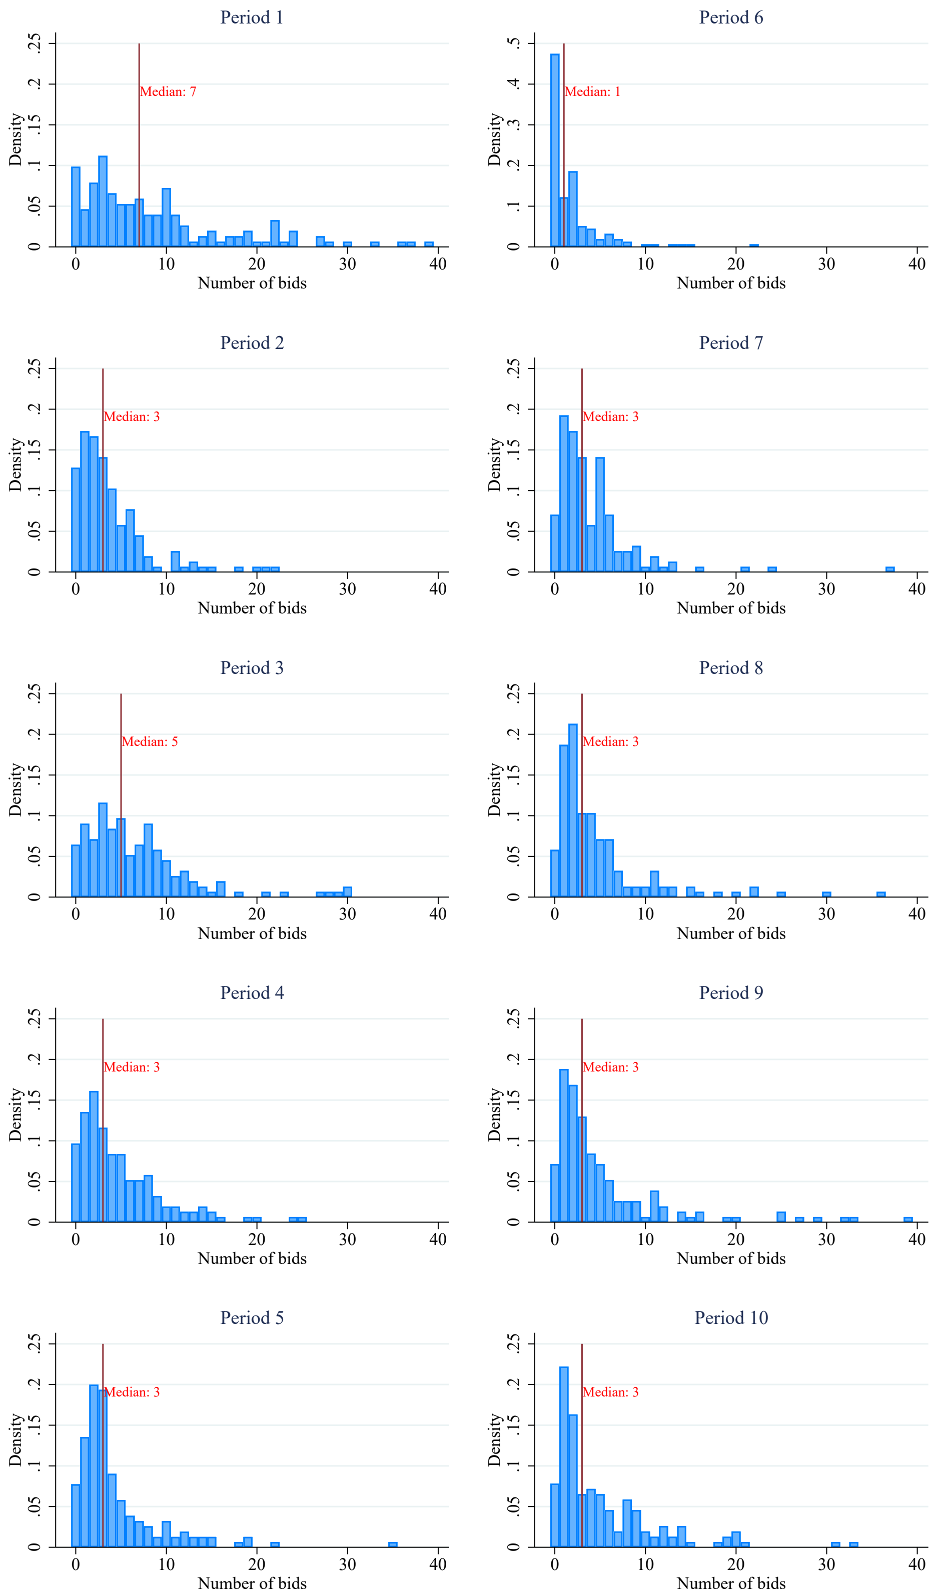

This figure presents the distributions of number of bids placed per plant by compliance period, truncated at 40 (about 99th percentile). The bin width is 1. The red line indicates the median number of bids placed.

C Appendix: Pollution Monitoring

C.1 Measuring plant emissions

We describe how we construct plant-level monthly average PM mass (in kilograms). CEMS provides stack-level daily reporting hours and uncalibrated daily average PM mass rate (kg/hr) or PM concentration (mg/Nm^3). 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.[32] 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. Plants can choose what kind of CEMS device to install (Type 1 or Type 2), so long as the device meets technical standards and passes the calibration test. Generally the devices calibrated to mass rates (Type 1) are simpler and less expensive than having a combination of a concentration device and a flowmeter to measure volume. Most plants therefore install Type 1 devices, without any mandate to do so from the regulator. 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/Nm^3).

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

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

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

where m and c are stack's calibration factors. Any negative calibrated value is set to missing. We

[32] This definition was an artifact of the market's initial timing. The first compliance period began on the 16th of September, following two mock periods running from the 16th of July. GPCB intended to start on the 1st of July but pushed the market back by two weeks to allow a grace period and increase CEMS reporting at the start of the market.

convert the mass rate to concentration, or vice versa, using

$$\text{PM\_Conc}^{cal}_{i,d} = \frac{1000^2 \text{ PM\_Rate}^{cal}_{i,d}}{(3600 \text{ max\_velocity}_i) \times \text{stack\_area}_i},$$

where max_velocity is the maximum flue velocity (m/s) of calibration samples, and stack_area is the stack cross-sectional area (m^2).

C.2 Imputation

Before we do any imputation, we truncate outliers in the non-missing data. Within each stack-day we take the average of the reported PM kg/hr at the minute-level. This gives us a set of stack-days across the entire sample. We then set to missing all those stack-days which fall above the 99th percentile.

Our final unit of observation is stack-month emissions (which are then summed within plants to get plant-months). To calculate emissions at the stack-month level we sum the emissions at the minute-level across all the minutes in the stack-month. Thus, our imputation aims to fill in all missing stack-minutes. We impute in two stages: first within the stack-week observation unit, and second within or across stack-months.

The first step is within stack-week imputation. Emissions are imputed for all missing stack-minutes in stack-weeks which have any data reported. To do this, we fill in all minutes in each week with the stack's average reported emissions during that week.

The second step is to either perform within stack-month imputation or across stack-month imputation. How we do this depends on which imputation rule we use. The three options are as follows:

- **The "no imputation" rule** only performs within stack-month imputation. Any remaining missing minutes from the first step (i.e. those in months with some data reported but in weeks with no data reported) are imputed at the average reported value across the entire month. This procedure does not perform any across stack-month imputation and leaves stack-months with no reported data as missing.

- **Rule A ("Stack Experiment")** imputes stack-weeks which do not report any data at the average level of the reported values for that stack over the entire ETS experiment (excluding the interregnum periods).

- **Rule B ("Treatment")** imputes stack-weeks which do not report any data at the average level of the reported values for that appropriate Treatment/Control group for that month (i.e. if that stack is part of the control group we average over just control plants in that month).

In Figure C1 we show the distribution of reporting frequency among those stack-months which require some intra-month imputation.

C.3 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 C2 shows the data availability of CEMS data on pollution by treatment arm. Figure C3 replicates Figure 5, from the main text, but with alternate imputation rules for missing data. Table C1 summarizes the level of log PM emissions by imputation rule and Figure C4 compares the distribution of pollution under different rules.

Figure C1: Share of data available within month for months with partial data

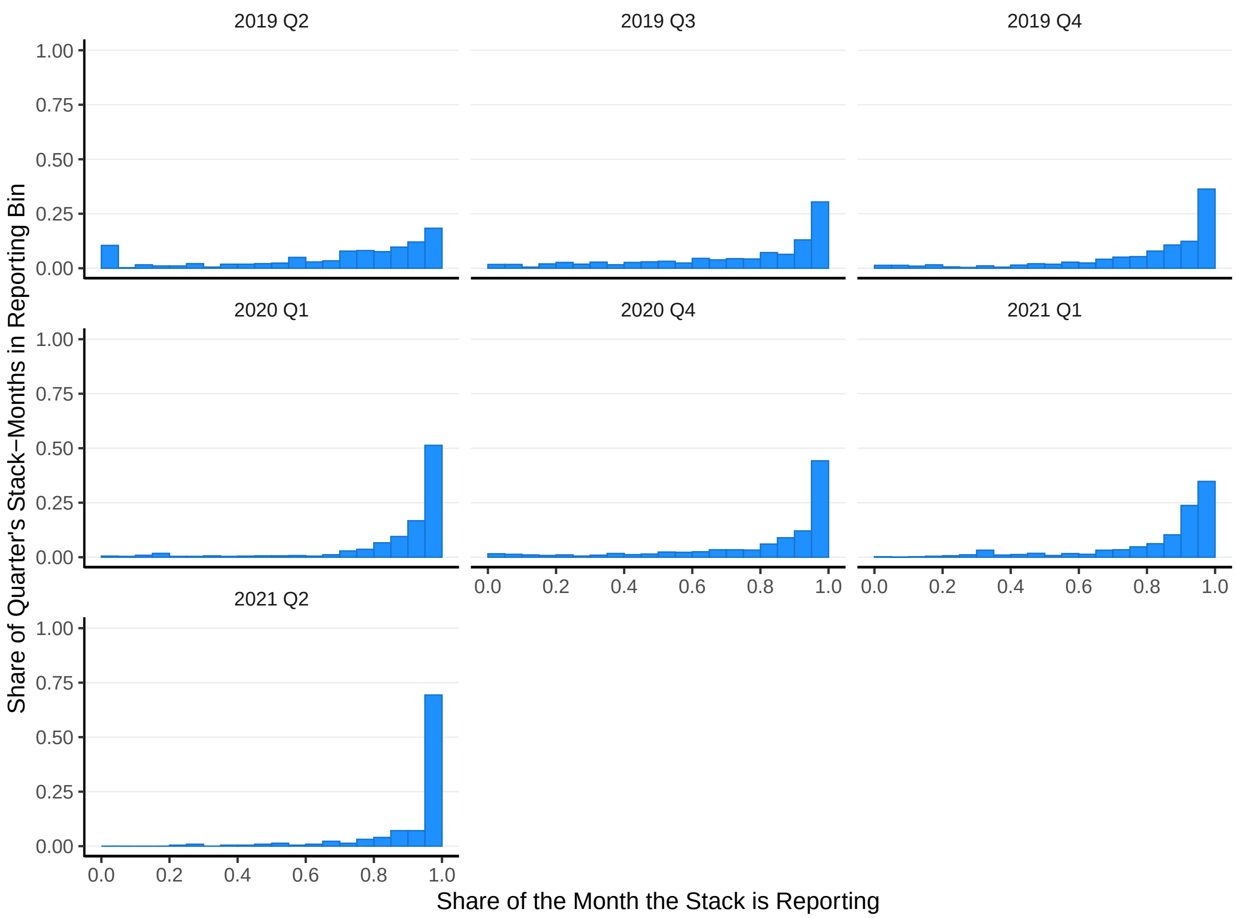

This figure plots a histogram of data availability at the stack-month level. The only stack-months included are those which require intra-month imputation, which are those with some, but not complete, minute-level reporting throughout the month. This represents 73% of the stack-months in the sample. The y-axis represents the portion of all plant-months in that panel which fall into the corresponding bin. Each panel represents a different quarter-year of the sample, excluding interregnum periods.

Figure C2: Data availability from CEMS by treatment status

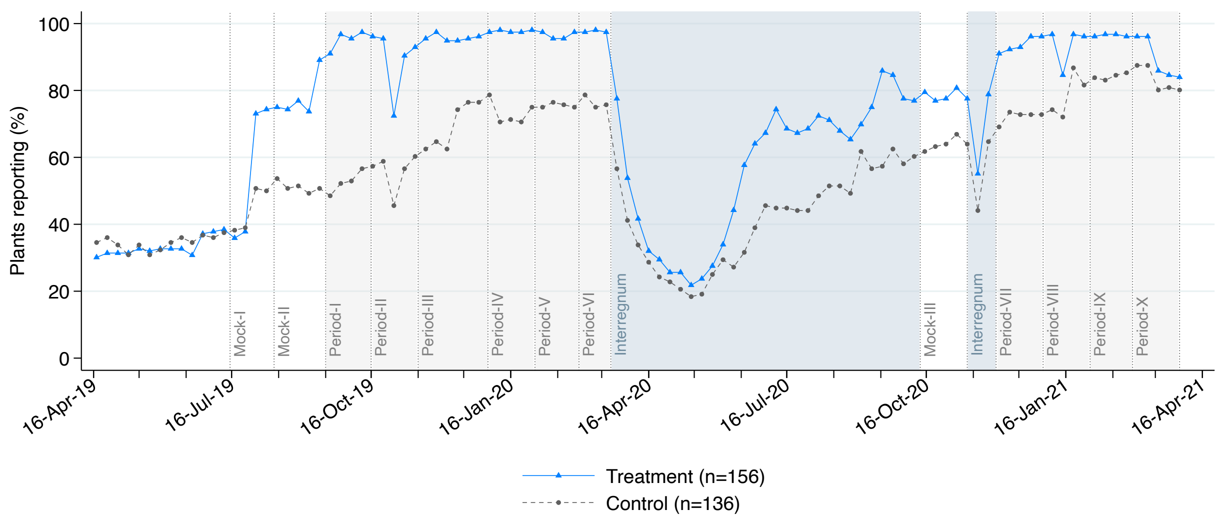

The figure shows the percentage of plants reporting, at weekly frequency, from April 2019 to March 2021. The missing pollution readings are imputed within a stack-week, but not across stacks or weeks. This sample consists of 292 plants that had at least one day of PM data from CEMS devices during the 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 C3: PM emissions by treatment status

A. Rule A: Stack-Experiment

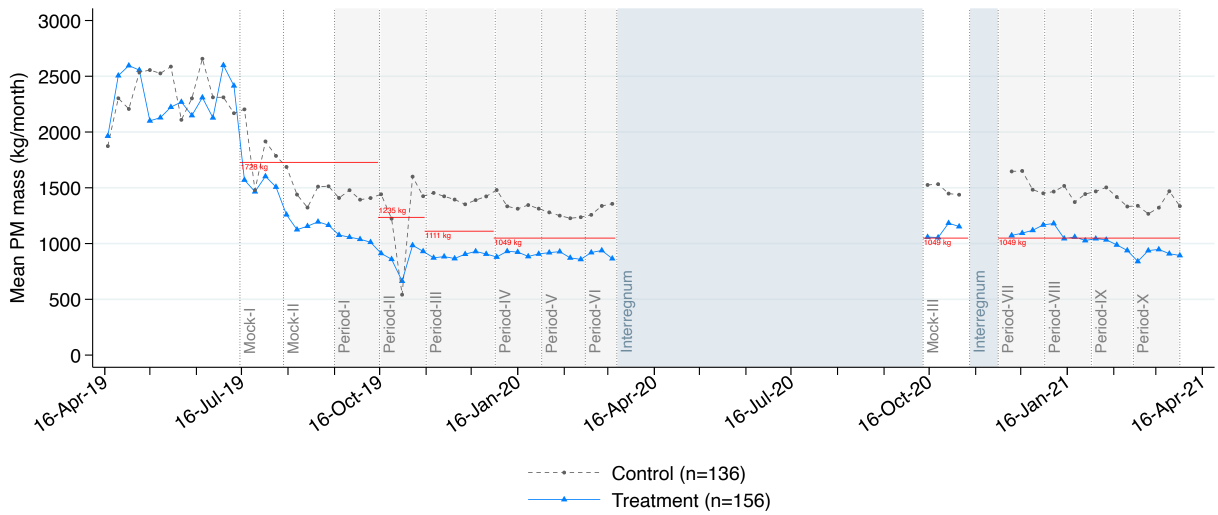

B. Rule B: Treatment-Month

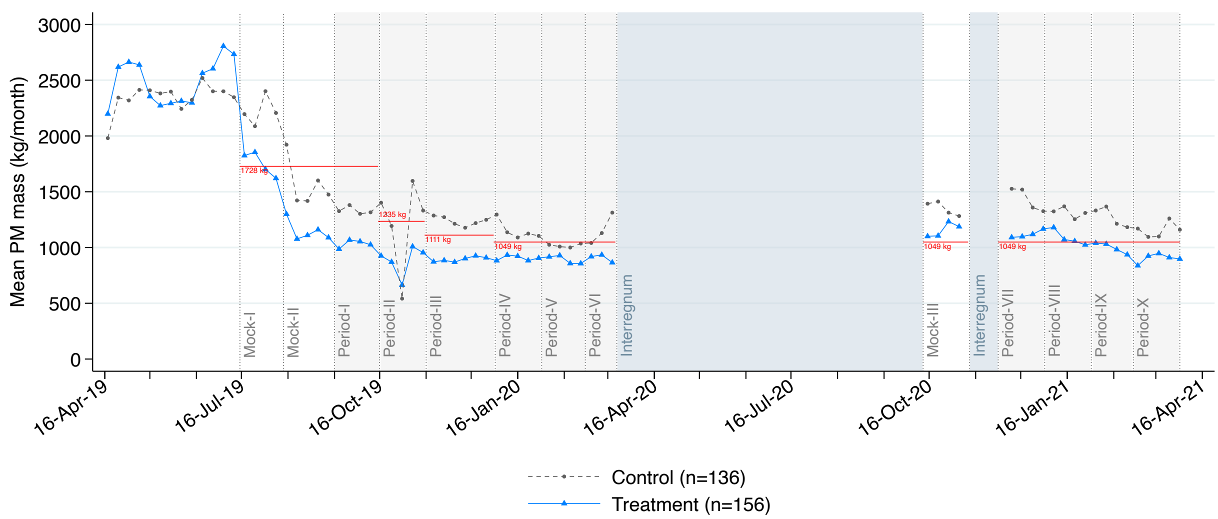

The figure shows the weekly mean per-plant PM emissions in kilograms calculated at a monthly rate equivalent, from April 2019 to March 2021. In the top panel, the missing pollution readings are imputed within stack-week, and then within stack-experiment. In the bottom panel, they are imputed within stack-week, 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 C1: Mean of the log(PM emissions) by imputation rules

	Control	Treatment	All
No Imputation	6.67	6.52	6.58
	[1336]	[1899]	[3235]
Rule A: Stack-Experiment	6.80	6.54	6.66
	[1768]	[2028]	[3796]
Rule B: Treatment-Month	6.88	6.59	6.72
	[1768]	[2028]	[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. Observational unit is stack-month, excluding interregnum and mock trading periods, across 292 potential plants in the whole sample.

C.4 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: filling gaps in the emissions record, *and* penalizing plants for non-reporting.

Table C2 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.

C.5 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

Figure C4: Kernel density of PM emissions by treatment status

A. PM emissions

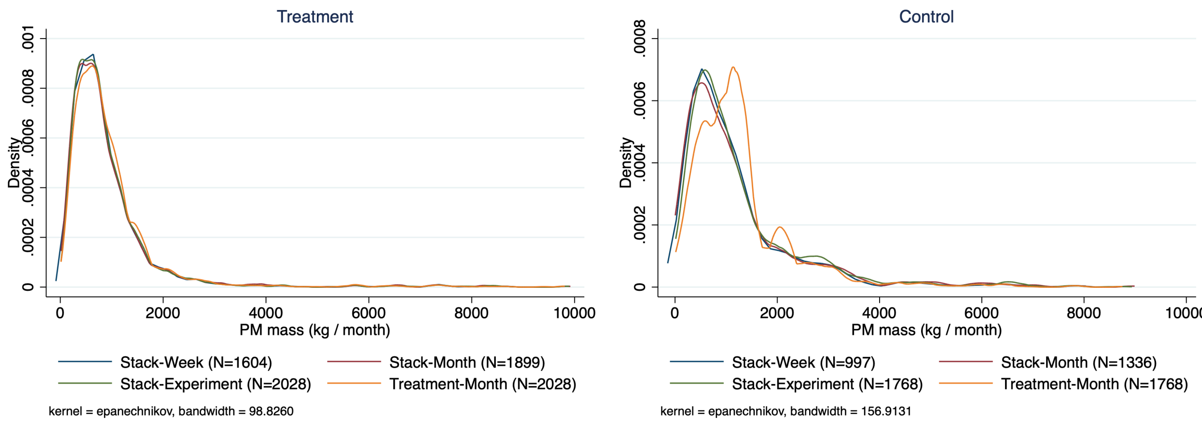

B. log(PM emissions)

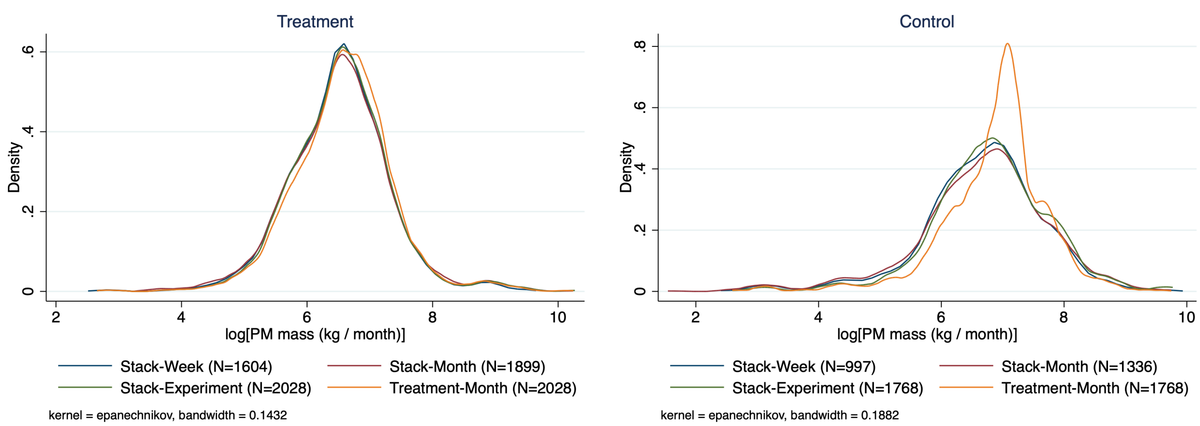

This figure plots the kernel density of PM emissions (kg/month) in Panel A and log(PM emissions) in Panel B, both by treatment status, in different stages of imputation described in Section C.2. Stack-Week corresponds to the emissions variable after step 2. Stack-Month, Stack-Experiment, and Treatment-Month 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. 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. 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.

regulatory action be based on pollution samples collected manually.

Here we ask whether CEMS monitoring, even without accompanying regulatory changes, might itself change plant behavior. We worked with GPCB to rollout CEMS as a randomized

Table C2: Imputation Rules for Missing CEMS Data

% Data Available During Week	Imputation for Missing Stack Data Values (kg/hr)
$> 95\%$	Stack's own mean operating emissions load during the week
80-95%	Stack's own 75^{th} percentile emissions load during the week
50-80%	Stack's own 90^{th} percentile emissions load during the week
1-50%	Stack's own 90^{th} percentile emissions load during the three months prior 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.

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 emission samples carried out by the environmental regulator as part of their inspection schedule. We run the following regression:

$$y_{it} = \beta_0 + \beta_1 Treat_i \times Post_t + \alpha_i + \gamma_t + \varepsilon_{it}$$

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 C3 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 C3: Effects of CEMS Installation on Plant Emissions

	PM, mg/Nm3 (1)	Log(PM) (2)
Treatment Effect	0.432	-0.0601
	(23.84)	(0.0912)
Observations	796	796
Year-Month FE	Yes	Yes
Plant fixed effects	Yes	Yes
Plants	197	197
R^2	0.3384	0.4276
Mean dependent variable	142.8	4.757

Dependent variable emission measures are from GPCB's regularly scheduled manual samplings. Unit of observation is plant. Sample is restricted to plants in Phases 1 and 3 of CEMS rollout. The treatment indicator in the regression is set to the interaction of the plant being in Phase 1 and not being in the control group. The treatment indicator is therefore set to 0 for the union of all experiment control plants and all Phase 3 plants. Standard errors are clustered at the plant-level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C.6 Treatment effect on PM emissions with different imputation rules for the control and treatment groups

This section examines how treatment effects vary based on the stringency of the imputation rules applied to control and treatment plants. Table C4 presents the results. On the main diagonal, where the imputation rules are assumed the same, the treatment effect is as large or larger in magnitude as the preferred estimates. The estimated treatment effect grows larger with more punitive (higher quantile) imputation rules because control plants are missing more data than treatment plants. Thus increasing emissions for missing data increases control emissions more than treatment emissions and increases the magnitude of the estimated treatment effect.

Off the main diagonal below, the treatment imputation rule is assumed to be relatively higher than for the control. For example, in the first column, Rule A is maintained for imputation in the control group, roughly imputing emissions at the mean for the same plant during times when it is reporting. The rows give the estimated treatment effect if the imputation rule for the control group remains at Rule A but the imputation rule for the treatment group increases, corresponding to higher emissions assumptions in the treatment group when treatment CEMS are not reported. The treatment effect is similar to the main estimate, although somewhat reduced, even when missing emissions in the treatment group are imputed at the 80th percentile of treatment group emissions when reporting. The treatment effect is negative but not statistically significant if treatment group emissions are imputed at the 90th percentile of treatment group emissions when reporting, and close to zero if treatment group emissions are imputed at the market imputation rule, which fills in extremely high levels of emissions as punishment when a plant does not report.

Table C4: Treatment effect on PM emissions (log(PM mass/month)) with different imputation rules for the control and treatment groups

Imputation rule – treatment		Imputation rule – control				
		Rule A	p70	p80	p90	Market
		(1)	(2)	(3)	(4)	(5)
	(1) Rule A	-0.282***	-0.368***	-0.461***	-0.591***	-0.904***
		(0.074)	(0.076)	(0.076)	(0.077)	(0.072)
	(2) p70	-0.243***	-0.329***	-0.422***	-0.552***	-0.865***
		(0.074)	(0.075)	(0.076)	(0.077)	(0.072)
	(3) p80	-0.192**	-0.278***	-0.371***	-0.501***	-0.814***
		(0.074)	(0.075)	(0.076)	(0.077)	(0.072)
	(4) p90	-0.109	-0.196**	-0.288***	-0.418***	-0.731***
		(0.075)	(0.076)	(0.077)	(0.078)	(0.073)
	(5) Market	0.008	-0.078	-0.171**	-0.301***	-0.614***
		(0.076)	(0.077)	(0.078)	(0.079)	(0.074)

This table reports estimated treatment effects on PM emissions, as in Table 3, column (5) of the main text, using different imputation rules for the treatment and control groups. 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. For each cell, the row describes the imputation rule used for treated plants and the column the imputation rule used for control plants. Rule A is stack-experiment imputation. 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). p70 imputes missing values of a stack's daily PM mass rate using the stack's 70$^{\text{th}}$ percentile of PM mass rate across the experiment (July 2019 to March 2021, excluding interregnum). p80 and p90 are identical to the p70 imputation rule except that they use the 80$^{\text{th}}$ and 90$^{\text{th}}$ percentiles of PM mass rate respectively. Market is the market imputation rule described in Table C2. All regressions control for plant characteristics including capital expenditure, operating cost, log(total heat output), mean boiler installation year, and their corresponding indicators for missing values. Robust standard errors in parentheses are clustered at the plant level with statistical significance indicated by $^{*}p<0.10$; $^{**}p<0.05$; $^{***}p<0.01$.

D Appendix: Model specification and abatement costs

D.1 Model specification

Abatement technology.—Plant i chooses the level of variable abatement expenditures Z_{it} in each compliance period $t = 1, 2, \ldots, 10$. 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 capacity H_i. Heat capacity 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 $\overline{E}_i$ such that $Z_{it}(\overline{E}_i) = 0$. 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.—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 capacity, $A_{it} \propto H_i$. Let P_t be the equilibrium price of permits, known to the plant.

Under the two assumptions of cost minimization and no market power, the plant chooses emissions to minimize the total cost of compliance:

$$\min_{E_{it}} Z_{i0} + Z_{it}(E_{it}) + P_t(E_{it} - A_{it}). \tag{6}$$

The first-order condition for the plant's problem under these assumptions is

$$-\frac{\partial Z_{it}(E_{it})}{\partial E_{it}} \;=\; MAC(E_{it}) = P_t. \tag{7}$$

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.

To calculate total costs we integrate marginal abatement costs to obtain each plant's total variable cost function. The marginal abatement cost function (4) we assume is consistent with a simple representation of total variable abatement costs:

$$Z_{it}(E_{it}) = e^{\beta_0+\tilde{\xi}_{it}} H^{\beta_2} \left(\frac{1}{\beta_1+1}\right) \left(\overline{E}_i^{\beta_1+1} - E_{it}^{\beta_1+1}\right), \qquad \beta_1 \in (-1,0), \tag{8}$$

where the parameters are common with (4). We estimate the parameters $\{\beta_1, \tilde{\xi}_{it}\}$ of the abatement cost function (8) using the marginal abatement cost specification (4). Moving from marginal to total variable abatement costs introduces a constant of integration. In (8), the constant $\overline{E}_i$ 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.

D.2 Calculating abatement costs by regime

We wish to compare abatement costs across regulatory regimes, but permit bids are only available in the treatment group. In this subsection we describe how we use the marginal abatement cost functions estimated in the treatment group to calculate abatement costs by regime.

Abatement costs in the emissions market.—In the market 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

$$E_t(P_t^*) = \sum_i E_{it}(P_t^*) = Q_t. \tag{9}$$

The equilibrium price is unique because 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 $E_{it}^* = E_{it}(P_t^*)$.

With our empirical specification of abatement costs we can solve for total variable abatement costs $Z_t^{ETS}(Q_t)$ at any proposed cap. The first step for a given cap Q_t is to solve (9) to find the equilibrium price P_t^*. With the estimated marginal abatement cost functions (4), the empirical inverse MAC function for each plant is:

$$E_{it}(P_t) \;\; = \;\; P_t^{1/\hat{\beta}_1} e^{-\hat{\xi}_{it}/\hat{\beta}_1}. \tag{10}$$

This function gives a plant's emissions as a function of the permit price. Substituting into (9), we can then find aggregate emissions at any price and solve for the equilibrium price $P_t^*(Q_t)$ for a given cap. We then calculate plant emissions with (10), evaluate plants' variable abatement costs (8) and sum across plants to find aggregate costs Z_t^{ETS}. The result of these steps is that we can write aggregate costs as a function of the aggregate emissions cap, $Z_t^{ETS}(Q_t)$.

Abatement costs in the command-and-control regime .—We estimate the stringency of regulation in the command-and-control regime in the control group. 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 *de facto* non-compliance with the intensity standard and fairly wide dispersion in emissions rates, rather than a point mass at the standard $\overline{R}$ (Appendix Figure B1).

We therefore estimate costs in the command-and-control regime by evaluating MAC functions at the observed emission rates in the control group. We represent emissions with plant-specific emissions rates $\overline{R}_{it} = E_{it}/H_i$ per unit of capacity. Since we observe emissions rates in the control group, it is straightforward to develop expressions for total emissions and total variable abatement costs. Total status quo emissions $E_t^{CC} = \sum_i H_i \overline{R}_{it}$ 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, $Z_{it}(H_i \overline{R}_{it})$. Summing across plants, total variable abatement costs under command-and-control are $Z_t^{CC} = \sum_i Z_{it}(H_i \overline{R}_{it})$.

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. The de jure standard is a uniform concentration standard. There is widespread noncompliance even with this standard. We do not think this non-compliance is likely to equalize marginal abatement costs across plants. 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. We therefore assume in our baseline case for the command-and-control regime that plant emissions rates are independent of plant marginal abatement costs.

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* $R_{it} = \overline{R}$; (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 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:

$$\log R_{it} = \beta_{0t} + \beta_{1t} \log H_i + \varepsilon_{it}. \tag{11}$$

The remaining three regimes we consider follow this approach: (iii) *capacity-based emissions rate* $R_{it} = \exp(\widehat{logR_{it}})$; (iv) *capacity-based emissions rate with error* $R_{it} = \exp(\widehat{logR_{it}}\varepsilon^s_{it})$ for draws $\varepsilon^s_{it} \perp \hat{\xi}_{it}$ from the residuals of (11); (v) *capacity-based emissions rate with correlated error*, similar to (iv), but with draws ε^s_{it} that are slightly negatively correlated ($\rho = -0.1$) with marginal abatement cost shocks $\hat{\xi}_{it}$.[33] 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.

In counterfactuals we evaluate costs not only at the distribution of emissions in the control group in the data, but also at higher or lower levels of emissions. We assume that a differently stringent command-and-control regime would scale up or down all emissions rates by a common factor δ. In the control group, we estimate fitted emissions rates across plants $\{\widehat{R}_{it}\}$ using one of the five regimes described above. We then calculate a scaling factor $\delta(Q_t) = Q_t/E^{CC}_t$ to meet emissions level Q_t. We evaluate plant-specific costs at alternate stringencies to calculate aggregate costs $Z^{CC}_t(\delta(Q_t)) = \sum_i Z_{it}(\delta(Q_t)H_i\widehat{R}_{it})$.

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

[33] This 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).

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.

Comparison of abatement costs across regimes.—Appendix Table D1, panel A explores the robustness of the finding that the the market regime reduces total variable abatement costs to different approaches to assigning each plant's command-and-control emissions. We use the same two reference levels of aggregate emissions, 170 tons (columns 1 to 3) and 240 tons (columns 4 to 6). Row A reports equilibrium market price and total variable abatement costs under the market. Rows B1 to B5 present 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.

There are two main findings. 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 2, row A). 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, row A versus column 5, row B4).

Second, the cost differences among the alternative representations of the command and control regime 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.[34] 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.

[34]On 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 column 2, 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).

We prefer the simulation draws that condition on plant heat capacity and draw over residualized emissions levels, since emissions rates do vary systematically with plant size (heat capacity). 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.

Table D1: Variable abatement costs under alternative regulatory regimes

	Emissions = 170 tons			Emissions = 240 tons		
	Price (INR/kg)	Cost (INR m)	ΔCost (%)	Price (INR/kg)	Cost (INR m)	ΔCost (%)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Iso-Elastic MAC Curve						
A. Emissions market	12.2	10.1	–	9.9	9.3	–
B. Command and Control						
1. Constant emissions rate		10.9	8.0%		10.2	10.0%
2. Constant emissions rate, with error		11.2	11.4%		10.6	14.1%
3. Capacity-based rate		10.9	8.2%		10.3	10.3%
4. Capacity-based rate, with error		11.3	11.8%		10.7	14.6%
5. Capacity-based rate, correlated error		11.4	12.9%		10.8	15.9%
Panel B: Step-Function MAC Curve						
A. Emissions market	15.3	6.0	–	11.6	5.0	–
B. Command and Control						
1. Constant emissions rate		7.9	32.4%		7.2	44.1%
2. Constant emissions rate, with error		7.9	32.6%		7.2	43.9%
3. Capacity-based rate		7.9	31.9%		7.1	42.9%
4. Capacity-based rate, with error		7.9	32.1%		7.1	42.9%
5. Capacity-based rate, correlated error		8.0	33.8%		7.3	45.5%

The table shows the results of counterfactual simulations under different regulatory regimes. Each row represents a different regime. Each panel corresponds to a different functional form assumption on the plant level marginal abatement cost curve. Within each panle the first row is the emissions market. The second through final rows in each panel 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: Marginal Cost Curve Specification

E.1 Uncontrolled Emissions

Without any abatement particulate emissions can be very high. We let $\overline{E}_i$ represent the uncontrolled level of emissions for plant i. To calculate it, we first find the average flow rate across manual samplings of that plant. These manual samplings include the ETS baseline, the CEMS baseline, as well as possible calibration measurements in 2020 and 2022. We then assume a maximum possible outlet concentration of 2500 $\frac{mg}{Nm^3}$, which was determined by our field staff to be a conservative plausible maximum, and we additionally assume operation for 12 hours per day for the entire period. These together determine a maximum possible emitted mass of PM for the stack for any period:

$$\overline{E}_i \frac{kg}{period} = \text{Flow-Rate}_i \frac{Nm^3}{hr} \cdot 12 \frac{hr}{day} \cdot 30 \frac{day}{period} \cdot 2500 \frac{mg}{Nm^3} \cdot \frac{1}{1,000,000} \frac{kg}{mg}$$

E.2 Step Function

In the body of the paper we assume that plants marginal abatement cost curves (hereafter MAC curves) follow a convenient iso-elastic form. However, it is worth examining what happens if we were to assume a different form of marginal abatement cost curve: We here consider a step-function MAC curve alternative.

Functional Form.—The basic form of the step-function MAC curve, for a specific plant i in period t, is shown in Figure E1. The intuition for this functional form is that if plant i doesn't run their abatement technology they will emit at E_i^{max}, if they do run their abatement tech they will decrease their emissions down to E_i^{min}, and running their abatement tech costs ξ_{it}. In order to implement this MAC curve for each plant, we thus need, for each i,t, values of E_i^{min}, E_i^{max} and ξ_{it}.

We set $E_i^{\max} = \left(1 - \mathbf{1}\{i \text{ has cyclone}\} \cdot 0.8\right) \cdot \overline{E}_i$. See Appendix E.1 for how we set $\overline{E}_i$. This setup allows plants to run their cyclones "for free" when they have one, decreasing their max potential

Figure E1: Example Step-Function MAC Curve for Plant i in Period t

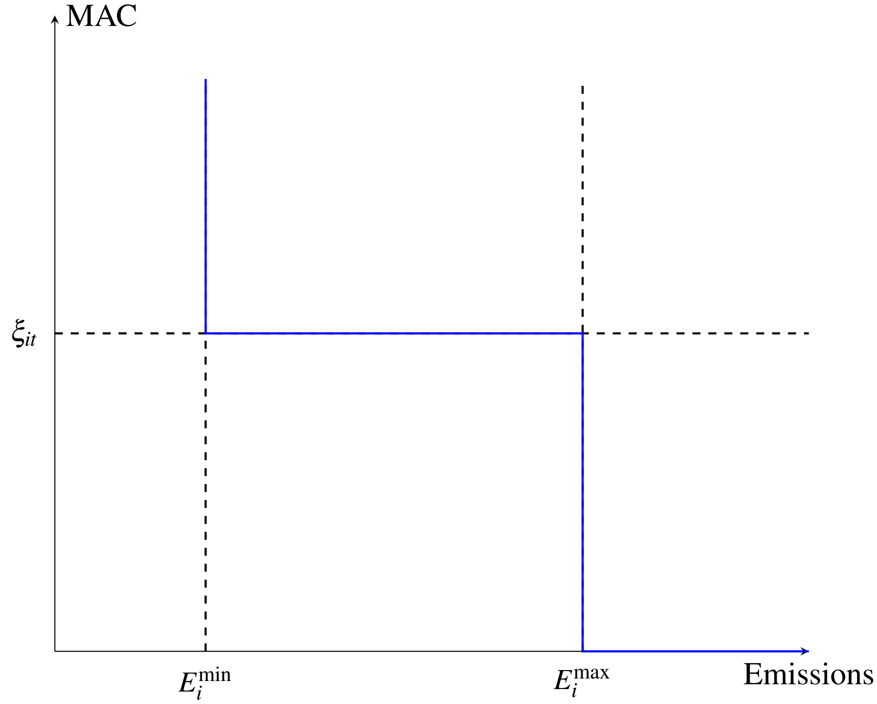

emissions by 80% (following the engineering estimates of cyclone efficiency reported in Table F1). Cyclones are mechanical abatement devices that run automatically based on the flow of the stack gas through the device. Thus, for those plants with a cyclone (which covers 98% of our sample) their abatement decision comes from whether or not they use further technology, with their cyclone being used at all times.

To calculate E_i^{min} we take the minimum of (a) the minimum observed level of emissions for plant i over any period, and (b) the abatement efficiency of plant i's most advanced technology (again following efficiencies in Table F1) times their E_i^{max}. The intuition for this is that the emissions of a plant over a period provide an upper bound on how far they are able to abate their emissions, while the engineering estimates of the efficiency of their abatement tech represents a more direct estimate of how far they can reduce their emissions beyond E_i^{max}.[35]

Lastly, to calculate ξ_{it} we again assume that plants bid their marginal costs of abatement, and

[35]For those plants which have a cyclone as their maximum abatement technology we set E_i^{min} to the minimum of (a) their minimum observed emissions over periods and (b) E_i^{max}, as the use of their cyclone has already been factored in to E_i^{max}.

thereby set ξ_{it} to the corresponding fixed-effect of a regression of first half of period bid prices on plant-period fixed effects:

$$b_{itk} = \xi_{it} + \varepsilon_{itk}.$$

Market Clearing Prices.— At a given market price P_t in period t, plant i will run their abatement equipment if and only if P_t is at least as high as their marginal abatement cost:

$$\text{Run}_{it}(P_t) = \mathbf{1}\{P_t \geq \xi_{it}\}$$

Therefore, E_{it} is a weakly decreasing function of P_t:

$$E_{it}(P_t) = E_i^{min} \cdot \text{Run}_{it}(P_t) + E_i^{max} \cdot (1 - \text{Run}_{it}(P_t))$$

We thus define the market clearing price in market period t, when there is market cap C_t as:

$$P_t^* = \min_{\mathbb{R}^+} P_t \text{ s.t. } \sum_i E_{it}(P_t) \leq C_t$$

The step functional form in which plants either don't abate at all or abate to their minimum means that it is possible for $\sum_i E_{it}(P_t^*) < C_t$. In this case, the discreteness of the abatement equipment leads to over-compliance.

At a given allocation of emissions across plants we use the model to calculate costs. The abatement cost function corresponds to the area under the MAC curve from E_{it} to ∞. However, the step functional form only allows E_{it} to be either E_i^{max} or E_i^{min} depending on Run_{it}. Abatement costs will then equal:

$$\begin{aligned} Z_{it}(E_{it}) &= \xi_{it} \cdot (E_i^{max} - E_{it}) \\ &= \text{Run}_{it} \cdot \xi_{it} \cdot (E_i^{max} - E_i^{min}) \\ &= \mathbf{1}\{E_{it} = E_i^{min}\} \cdot \xi_{it} \cdot (E_i^{max} - E_i^{min}) \end{aligned}$$

Command-and-Control As before we assume the command-and-control regime sets standards and associated emissions levels E_{it} for each plant-period. If this assigned $E_{it} < E_i^{min}$ then it will be impossible for plant i to abate to this level, so we re-assign $E_{it} = E_i^{min}$. We then assume that plants abate to achieve these levels of emissions, costing them:

$$Z_{it}(E_{it}) = \xi_{it} \cdot (E_i^{max} - E_{it})$$

Results.—We now duplicate several results from the paper using the alternative, step-function functional form for marginal abatement costs.

Figure E2 shows market prices calculated with the step-function form in red as a test of in-sample model fit. The step-function model tends to over-predict market prices, relative to the iso-elastic cost function model.

Figure E3 shows the distribution of emissions in the treatment market calculated with the step-function MAC (panels A and B), the iso-elastic MAC (panels C and D) and in the data. The step-function MAC leads to a multi-modal distribution of emissions. Because plants either run their equipment (if the cost of the step is low enough) or do not, emissions are dispersed and there are separate modes for plants based on these operating decisions. By contrast, the distributions of emissions for the iso-elastic MACs and in the data are smoother.

We carry the step-function results through to counterfactual analysis of market cost savings in Table D1, panel B. The rows of Panel B correspond to the different regimes discussed in , simply changing the assumed functional form of the marginal abatement cost curves. The main finding is that the alternative, step-function MAC model predicts much larger counterfactual cost savings from the emissions market. The reason for this result is that, in the step function model, if a plant is estimated to have low marginal abatement costs it will always have low costs, up to the maximum efficacy of a piece of equipment. The costs of misallocation of abatement are therefore large because the model extrapolates in-sample differences in cost over a large range of emissions. In the iso-elastic MAC model, by contrast, a low-cost plant cannot take over such a large share of counterfactual abatement in the market, because its own MAC would curve upwards.

Figure E2: Model Fit to Market-Clearing Prices with Step-Function Alternative

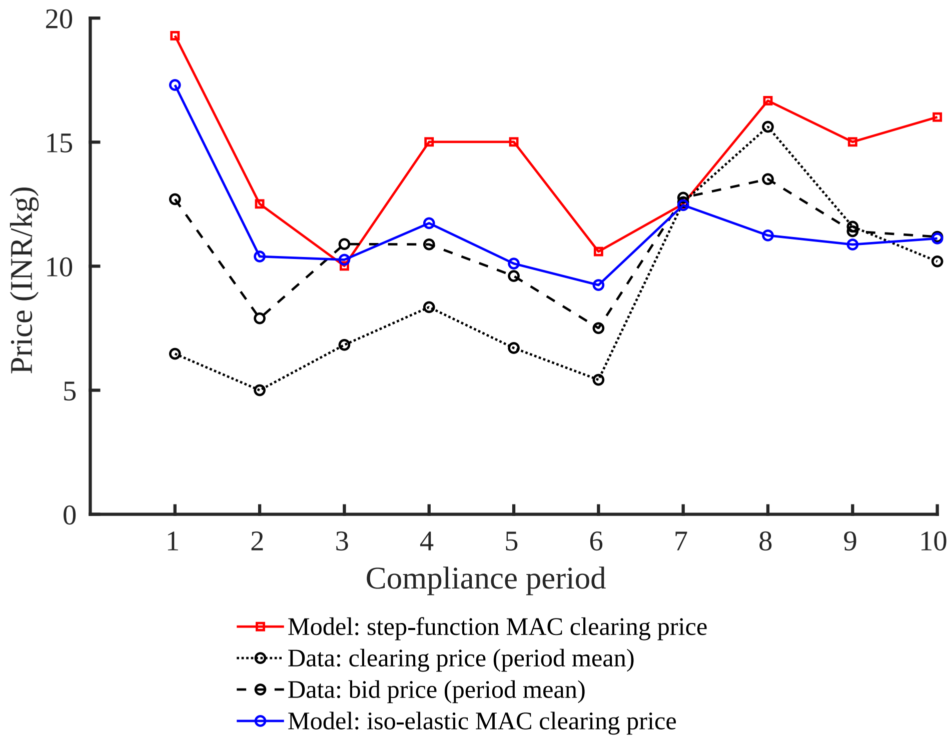

The figure shows the fit of the step-function and iso-elastic MAC models compared to the time series of market and bid prices by compliance period. The solid red line with square points is the time series of market-clearing prices in the fitted model with step-function MACs. The solid blue line with circular points is the time series of market-clearing prices in the fitted model with the original iso-elastic MACs. The models are 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.

E.3 Iso-elastic MAC with heterogeneous elasticities of abatement costs

Our main specification for marginal abatement cost curves (4) allows the mean log marginal abatement costs for each plant-period to differ but constrains all plants to have the same elasticity of MAC with respect to emissions. This part considers our counterfactual results if we allow heterogeneity in the MAC elasticity. In Table 4, column 5 allows the elasticity of MAC to differ by what abatement equipment a plant has installed. Table E1 replicates the counterfactual results of Table D1 Panel A, with the specification of column 5. We find that the magnitude and qualitative pattern of cost savings in the emissions trading regime relative to the control regime are similar to those reported in Table D1 Panel A.

Figure E3: Histograms of predicted versus observed emissions

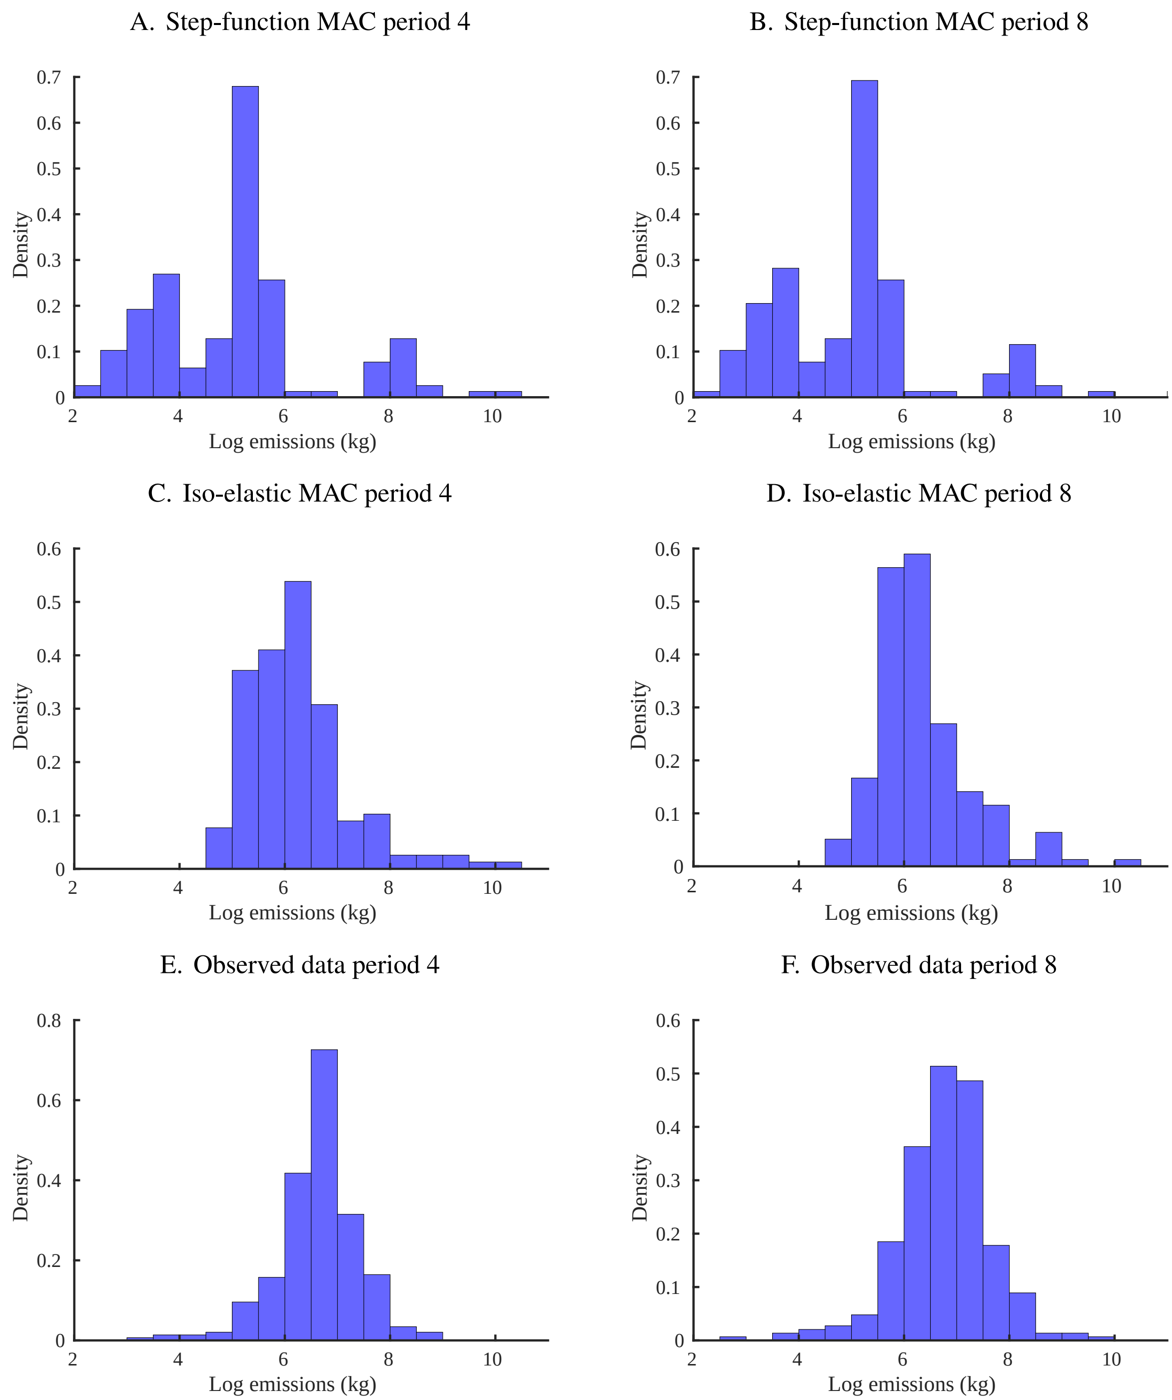

The figure shows predicted and observed emissions levels in 2 periods for 2 different MAC curve specifications. Panels A and B show predicted emissions when running our emissions market model under a step-function MAC curve in periods 4 and 8 respectively. Panels C and D show the same except using the iso-elastic MAC curve. Panels E and F show the observed distribution of emissions in those periods.

Table E1: Variable abatement costs under alternative regulatory regimes (with Heterogeneity by APCD)

	Emissions = 170 tons			Emissions = 240 tons		
	Price (INR/kg) (1)	Cost (INR m) (2)	ΔCost (%) (3)	Price (INR/kg) (4)	Cost (INR m) (5)	ΔCost (%) (6)
A. Emissions market	12.2	10.1	–	10.0	9.3	–
B. Command and Control						
1. Constant emissions rate		11.0	8.6%		10.3	10.7%
2. Constant emissions rate, with error		11.3	11.8%		10.7	14.7%
3. Capacity-based rate		11.0	8.9%		10.4	11.1%
4. Capacity-based rate, with error		11.3	12.3%		10.8	15.2%
5. Capacity-based rate, correlated error		11.4	13.0%		10.8	16.2%

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.

F Appendix: Additional Results

F.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 3, 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 F1 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 F1: Engineering estimates of abatement costs under ideal operating efficiency, if a cyclone is already operating

	Cyclone (1)	Bag Filter (2)	Scrubber (3)	ESP (4)
	Total Boiler Capacity = 3 TPH			
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
	Total Boiler Capacity = 6 TPH			
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
	Total Boiler Capacity = 8 TPH			
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
	Total Boiler Capacity = 15 TPH			
Capital costs (Rs/month, amort.)	13906.67	20860.00	26075.00	234675.00
Variable costs (Rs/month)	6000.00	9000.00	11250.00	101250.01
Emission reduction (%)	80.00	99.00	94.00	99.70
Assumed pollution (kg/month)	8781.49	8781.49	8781.49	8781.49
Emission abatement (kg/month)	7025.19	8693.67	8254.60	8755.14
Average abatement cost (Rs/kg)	2.83	3.43	4.52	38.37
Variable abatement cost (Rs/kg)	0.85	1.04	1.36	11.56

Note. Table displays engineering estimates of abatement cost for different APCDs and boiler capacities. We assume one cyclone is already operating when calculating the quantity of abatement, and we assume each APCD is purchased in isolation. Costs can be compared with those in other tables at a rate of INR 70 to USD 1. Capital costs are amortized to a monthly flow value. All plants are assumed to have a raw inlet concentration of 2,000 mg/Nm^3; in practice it can vary between 1,000 mg/Nm^3 and 10,000 mg/Nm^3. This is converted to a monthly mass rate via a volumetric flow rate collected at baseline, assuming continuous operation for 16 hours/day and 25 days/month. Of plants with boilers in our analysis sample, the boiler capacity (BC) distribution is: 11% have 2-3 TPH BC, 47% have 4-7 TPH BC, 36% have 8-14 TPH BC, 6% have 15+ TPH BC.

F.2 Treatment effects on capital installation

Table F2 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 F2: Treatment effects on the presence of abatement devices

	All APCDs	Components			
		Cyclone	Bag	Scrubber	ESP
	(1)	(2)	(3)	(4)	(5)
ETS Treatment (=1)	0	0.0233*	0.0650***	-0.0151	-0.0311
	(.)	(0.0134)	(0.0231)	(0.0310)	(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 $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

F.3 Model robustness checks

Heterogeneity in estimated elasticities by time of bid.—Section 5 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 F1 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 F1: Elasticity estimate by weeks remaining in the order period

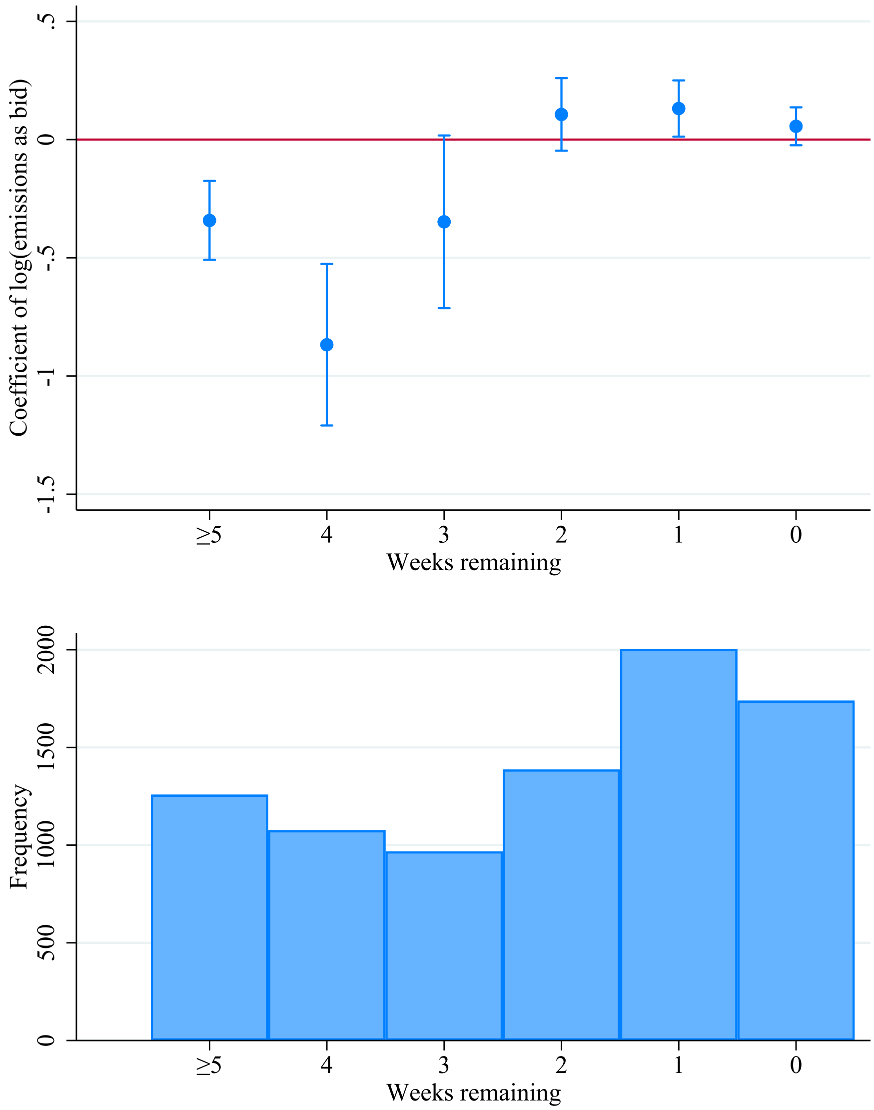

The top panel presents the coefficients of log(emissions as bid) from regressing ln(bid price) on log(emissions as bid) and plant × period fixed effects, estimated with different sample truncations defined by the number of weeks remaining in the order period. The bottom panel shows the number of bids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

F.4 Significance of Emissions Non-Reporting

In this section we investigate if differences between plants who report different amounts might be driving any of our results.

Table F3 shows the same balance table as Table 1 from the main paper, only now comparing those plants above versus below median levels of daily data reporting. There are some differences between high- and low-reporting plants, but they are not large. The main difference appears to be that high-reporting plants are somewhat larger, with higher sales revenue and boiler house capital expenditure (panel A). There are no differences in abatement equipment installation (panel B). On emissions, high-reporting plants have similar PM emissions mass rates and baseline PM concentrations to low-reporting plants, and high-reporting plants are somewhat more likely (10 pp with a standard error of 5.5 pp) to be above the emissions concentration standard at baseline.

We next examine whether abatement costs or being in the treatment group are related to plant's data reporting. To do this we calculate predicted abatement costs using baseline rates of equipment installation for different plants and the average costs from Table F1. We then regress reporting rates on predicted abatement costs and their interaction with treatment status to test whether plants with higher costs report less. The results of this regression are in Table F4. We do not find any significant effect of predicted abatement costs on reporting or differential reporting in the treatment.

Table F5 then estimates treatment effects on pollution controlling directly for plant reporting rates. We estimate a similar magnitude of average treatment effect conditional on reporting rates.

Table F3: Balance of plant characteristics by whether report more than median reporting

	Over Median (1)	Under Median (2)	Difference (3)
Panel A: Plant Measures			
Total electricity cost (1,000 USD)	466.5	344.1	122.4
	[833.7]	[401.9]	(78.1)
Log(plant total heat output)	15.6	15.6	0.042
	[0.62]	[0.48]	(0.065)
Size as recorded on environment consent (1 to 3)	1.41	1.31	0.10
	[0.66]	[0.59]	(0.074)

Small-scale (size=1)	0.68	0.75	-0.068
	[0.47]	[0.43]	(0.054)
Large-scale (size=3)	0.096	0.062	0.033
	[0.29]	[0.24]	(0.032)
Number of stacks	1.08	1.05	0.030
	[0.40]	[0.21]	(0.037)
Textiles sector (=1)	0.85	0.88	-0.028
	[0.36]	[0.33]	(0.041)
Gross Sales Revenue in 2017 (1,000 USD)	15125.8	6950.0	8175.8*
	[54715.6]	[13111.3]	(4609.6)
Panel B: Plant Abatement and Investment Cost			
Boiler house employment	36.4	32.7	3.78
	[32.4]	[30.0]	(3.70)
Boiler house capital expenditure (1,000 USD)	215.8	149.6	66.2*
	[403.6]	[174.7]	(36.8)
Boiler house operating cost (1,000 USD)	142.8	108.1	34.7*
	[202.5]	[83.9]	(17.8)
APCD: Cyclone present	0.98	0.97	0.012
	[0.14]	[0.17]	(0.019)
APCD: Bag filter present	0.82	0.86	-0.038
	[0.38]	[0.35]	(0.043)
APCD: Scrubber present	0.64	0.62	0.020
	[0.48]	[0.49]	(0.058)
APCD: ESP present	0.12	0.070	0.051
	[0.33]	[0.26]	(0.035)
Panel C: Plant Pollution Measures			
Plant total PM mass rate (kg/hr)	3.70	3.50	0.20
	[4.97]	[3.73]	(0.52)
Plant mean PM concentration (mg/Nm^3)	179.5	167.8	11.7
	[154.2]	[152.3]	(18.2)
Plant mean Ringelmann score (1 to 5)	1.32	1.41	-0.086*
	[0.39]	[0.41]	(0.047)
Above regulatory standard at ETS baseline (=1)	0.36	0.26	0.10*
	[0.48]	[0.44]	(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 plants who report above and below the median level across plants. Each plants level of reporting is calculated as the average minute-level CEMS data availability across the full sample period and across all stacks belonging to that plant. The only plants which are included in this table are those in the analysis set. This sample consists of 292 plants that had at least one day of PM data from CEMS devices during the ETS experiment. 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 F4: Treatment effect on reporting by predicted plant abatement costs

	Dependent Variable: Share of Day Not-Reporting
Treatment (=1)	-0.166***
	(0.035)
Predicted Abatement Cost	0.000
	(0.002)
Predicted Abatement Cost × Treatment	0.000
	(0.002)
R^2	0.13
Observations	304

Unit of observation is plant. Predicted Abatement Cost variable for industry set to the engineering estimate of the average abatement cost per kg from Table F1 assuming Boiler Capacity = 8TPH) for the most advanced abatement technology of the plant: Cyclone less advanced than bag-filter less advanced than scrubber less advanced than ESP. Share of day not reporting calculated at industry level as average over industry's stacks and over all days (excluding interregnum). Robust standard errors in parentheses. Statistical significance is indicated by $^{*}p<0.10$; $^{**}p<0.05$; $^{***}p<0.01$.

Table F5: Treatment effects on PM emissions (log(PM mass/month)) controlling for data availability

	No Imputed Months				Imputed Months			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ETS Treatment (=1)	-0.179**	-0.203***	-0.180**	-0.207***	-0.240***	-0.247***	-0.233***	-0.250***
	(0.077)	(0.077)	(0.076)	(0.076)	(0.073)	(0.073)	(0.061)	(0.061)
Share of Day Reporting	0.0000343	0.00127	0.000325	0.00148	-0.00222***	-0.00187**	-0.00445***	-0.00350***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Year-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	0.18	0.14	0.17	0.19	0.22	0.19	0.27
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 adding average availability. 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 year-month 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 $^{*}p<0.10$; $^{**}p<0.05$; $^{***}p<0.01$.

F.5 Impact of COVID-19 Pandemic

In this section we include several analyses trying to determine the impact of COVID-19 on the experiment.

First, we examine how net-demand (defined as a plant's total period emissions less their initial permit allocation) differed before and after the COVID interregnum. Figure F2 shows a scatter plot of the plant net permit demand before (compliance periods 1 to 6) and after (periods 7 to 10) the COVID-19 interruption. The scatter plot shows that plant net demands are highly correlated in the pre- and post-Covid periods. Plants that have higher emissions than permit allocations before the pandemic tend to also have higher emissions than permit allocations afterwards.

In Table F6 we then estimate the treatment effect on emissions separately for the pre- and post-COVID subsets of our sample. In specifications without imputation or with imputation Rule A there is no statistically significant difference in the treatment effect before and after the Covid-19 lockdown. In specifications with imputation rule B, the treatment effect is statistically smaller in magnitude (less negative) after the lockdown but remains large, negative and statistically significantly different from zero. The point estimates for the treatment effect on pollution are smaller post-lockdown, which would be consistent with a less tightly binding cap in a weaker economy.

In Table F7 we re-estimate the counterfactual market versus command-control results from Table D1 using both the iso-elastic and step-function MAC, only restricting the sample to the pre-COVID periods.

Lastly, Figure F3 displays emissions over final permit holdings without GPCB's period 7 permit adjustment. This figure is identical to Figure 4 other than that it removes additional permits granted to plants during compliance period 7, the initial post-Covid-lockdown period. Footnote 22 describes the adjustment in more detail.

Figure F2: Net demand before and after COVID

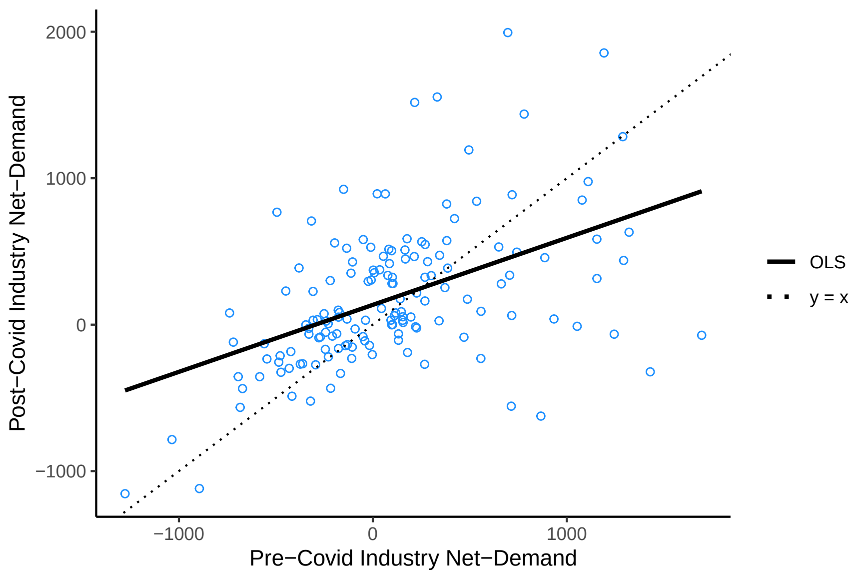

The figure shows net demand (emissions – initial allocation) for each industry averaged across pre vs post COVID compliance periods. Each point represents a single industry. The solid black line is the OLS fit for the data The dotted black line is the $y = x$ line. We omit the 9 industries with values of magnitude greater than 2000 for ease of visualization.

Table F6: Treatment effects on PM emissions (log(PM mass/month)) before and after COVID

	No Imputation				With Imputation			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ETS Treatment (=1)	-0.211**	-0.219**	-0.211**	-0.224***	-0.303***	-0.303***	-0.372***	-0.372***
	(0.087)	(0.088)	(0.085)	(0.085)	(0.077)	(0.077)	(0.060)	(0.060)
Post-Covid (=1)	-0.173**		-0.159**		-0.149***		-0.276***	
	(0.082)		(0.079)		(0.047)		(0.059)	
Treatment × Post-Covid	0.0652	0.0656	0.0668	0.0717	0.0538	0.0538	0.144**	0.144**
	(0.093)	(0.094)	(0.091)	(0.092)	(0.060)	(0.060)	(0.071)	(0.071)
Year-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	0.14	0.17	0.14	0.17	0.19	0.22	0.18	0.26
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. Post-Covid is defined as periods 7 to 10. 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 year-month 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 F7: Variable abatement costs under alternative regulatory regimes using only pre-COVID data

	Emissions = 170 tons			Emissions = 240 tons		
	Price (INR/kg) (1)	Cost (INR m) (2)	ΔCost (%) (3)	Price (INR/kg) (4)	Cost (INR m) (5)	ΔCost (%) (6)
Panel A: Iso-Elastic MAC Curve						
A. Emissions market	12.7	8.1	–	9.7	7.3	–
B. Command and Control						
1. Constant emissions rate		9.0	11.4%		8.3	13.6%
2. Constant emissions rate, with error		9.4	16.9%		8.8	20.4%
3. Capacity-based rate		8.9	10.9%		8.2	12.9%
4. Capacity-based rate, with error		9.4	16.7%		8.7	19.9%
5. Capacity-based rate, correlated error		9.5	18.2%		8.9	21.8%
Panel B: Step-Function MAC Curve						
A. Emissions market	16.3	5.8	–	12.0	4.8	–
B. Command and Control						
1. Constant emissions rate		9.0	54.6%		8.2	72.7%
2. Constant emissions rate, with error		9.0	54.7%		8.2	72.3%
3. Capacity-based rate		8.9	52.5%		8.0	68.7%
4. Capacity-based rate, with error		8.9	52.7%		8.0	68.5%
5. Capacity-based rate, correlated error		9.0	54.7%		8.2	71.8%

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. Data used for estimation is restricted to pre-COVID periods (periods 1 to 6) only.

Figure F3: Distribution of Emissions over Final Permit Holdings by Compliance Period without GPCB's Period 7 Adjustment

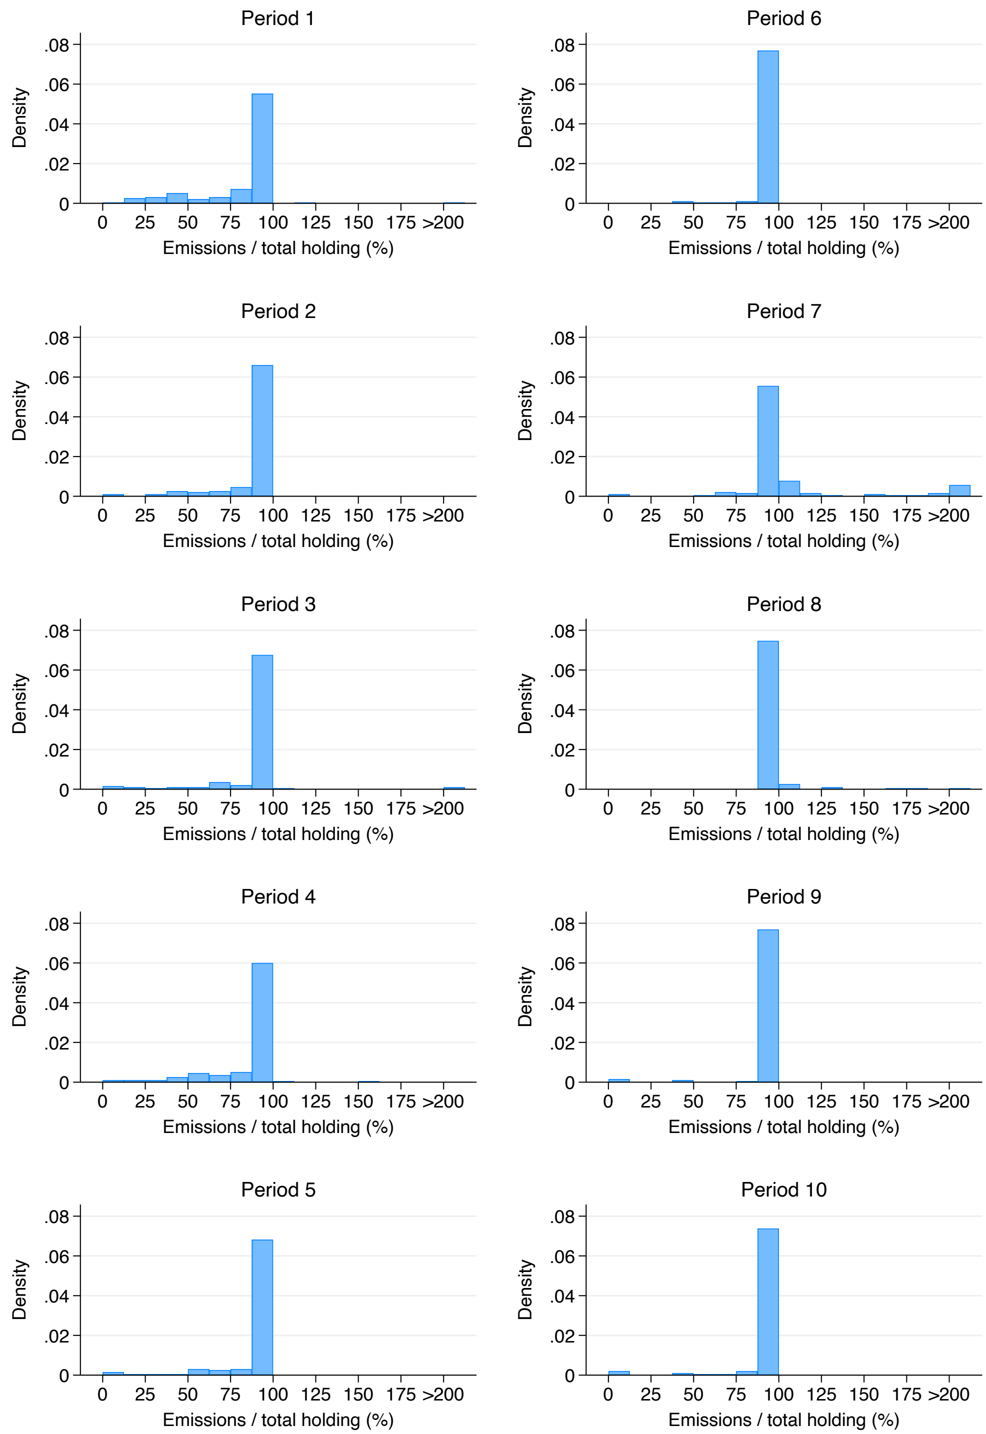

This figure plots the distributions of (emissions / final permit holdings $\times$ 100%) across treated plants ($N = 156$) by compliance period. Final permit holdings are the total number of permits a plant held at the end of the true-up period after each compliance period. Emissions data and permit holdings are from the administrative records of the market operator. Permit holdings are adjusted to remove those granted in GPCB's period 7 adjustment. 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.

F.6 Manual Sampling and CEMS Comparison

Figure F4 plots emissions measurements from manual samplings versus CEMS readings of the emissions from the same window of time during which the manual sampling was taking place. A regression line and the $y = x$ line are also shown. There is a high correlation between the manual samples and CEMS readings.

Figure F4: Simulatenous CEMS and sampling comparison

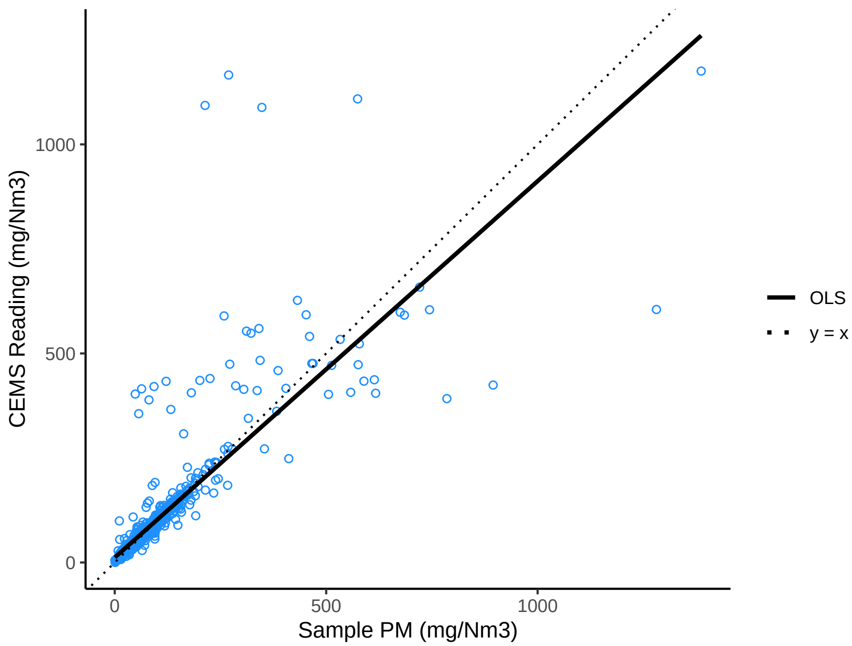

The figure plots CEMS readings against concurrent manual samplings. Unit of observation is an industry. The solid black line is the OLS fit for the data The dotted black line is the $y = x$ line. We restrict the graph to only those CEMS readings with at least 15% data availability during the appropriate window, and to those with concentrations less than 2000 mg/Nm3.

F.7 Impact of Device Types

In Table F8 we estimate the treatment effect on emissions levels separately for each potential CEMS device type (type 1, types 2, or both types). We do not find any significant effects of device type on treatment effect. The coefficient on Type 2 devices is positive, though statistically insignificant. A positive effect would be consistent with larger, more sophisticated plants selecting a more expensive piece of equipment.

Table F8: Treatment effects on PM emissions (log(PM mass/month)) depending on device type

	No Imputation				With Imputation			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ETS Treatment (=1)	-0.130	-0.147*	-0.135	-0.154*	-0.247***	-0.247***	-0.310***	-0.310***
	(0.087)	(0.087)	(0.085)	(0.085)	(0.084)	(0.084)	(0.065)	(0.065)
Device FE: Type 2	0.159	0.177	0.141	0.159	0.0955	0.0955	0.126	0.126
	(0.219)	(0.221)	(0.221)	(0.224)	(0.173)	(0.173)	(0.136)	(0.137)
Device FE: Both Types	0.362*	0.354*	0.362*	0.355*	0.289	0.289	0.179	0.179
	(0.212)	(0.209)	(0.204)	(0.201)	(0.210)	(0.211)	(0.156)	(0.156)
Treatment $\times$ Device FE: Type 2	-0.0910	-0.0934	-0.0639	-0.0652	-0.00556	-0.00556	-0.0495	-0.0495
	(0.264)	(0.266)	(0.267)	(0.271)	(0.228)	(0.228)	(0.190)	(0.190)
Treatment $\times$ Device FE: Both Types	-0.319	-0.302	-0.311	-0.299	-0.229	-0.229	-0.0798	-0.0798
	(0.269)	(0.268)	(0.266)	(0.265)	(0.261)	(0.262)	(0.214)	(0.214)
Year-Month FE		Yes		Yes		Yes		Yes
Imputation rule					Rule A	Rule A	Rule B	Rule B
Reweighted			Yes	Yes				
Mean dep. var (control)	6.65	6.65	6.64	6.64	6.76	6.76	6.87	6.87
R^2	0.14	0.18	0.14	0.18	0.18	0.22	0.17	0.26
Plants	279	279	279	279	279	279	279	279
Observations	3110	3110	3110	3110	3627	3627	3627	3627

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. We add fixed effects for the different types of abatement devices which an industry has across all of its stacks (Type 1, Type 2, or both). Type 1 Devices are the omitted level of device type fixed effect. Approximately 80%, 10%, and 10% of plants are set to Type 1, Type 2, and having both types, respectively. 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 year-month 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 $^{*}p<0.10$; $^{**}p<0.05$; $^{***}p<0.01$.

F.8 Prior Regulations in Other Markets

The market experiment in this paper is layered on top of existing regulations which mandate investment in abatement and monitoring technologies such as CEMS devices. This layered mandate is typical of the way in which markets have been implemented throughout the world and studied throughout the economics literature. In Table F9 we canvas 6 different emissions markets in the US and EU, give the prior regulations which plants were facing, whether those regulations were lifted once the market was put in place, and also give a citation to a paper studying those markets.

The table summarizes command-and-control regulations that were in place before or alongside major emissions markets. Each row considers one market. The "Program" column names the emissions market and the pollutant it covers. The "Prior Regulation" column lists relevant regulation which was introduced before or concurrent to the emissions market. CAAA refers to the Clean Air Act Amendments; NSPS refers to New Source Performance Standards (applies to all new sources of NOx and SO2, applying uniform national standard based on best adequately demonstrated technology); BACT refers to Best Available Control Technology (applies to all new sources of NOx and SO2 emitting significant amounts in attainment areas and is at least as strict as NSPS); RACT refers to Best Reasonably Available Control Technology. The "Prior Regulation Lifted" column indicates whether (or how) the prior regulations were adjusted at the point of introducing the emissions market. The "Paper on Market" column gives citations to papers studying the market.

Table F9: Prior regulations across emissions markets

Program (Pollutant)	Location	Year Instituted	Prior Regulation	Prior Regulation Lifted?	Paper on Market
Nitrogen Oxides Budget Program (NOx)	Eastern US	2003	• 1970 CAAA: NSPS, BACT • 1990 CAAA: Required subset of boilers in extreme areas to transition to low-pollution fuel • 1990 CAAA: Acid Rain program manted installing NOx monitoring and abatement normally met with control technology • In 1995 Ozone Transport Commission required existing sources to meet RACT limits and ran OTC NOx budget program	No	Deschênes, Greenstone and Shapiro (2017)
AB-32 cap-and-trade system (CO2)	California	2013	• Bill establishing market included complementary programs and modifications of existing programs, including LCFS, RPS, and efficiency mandates • LCFS adopted in 2009 and "require[d] the carbon intensity of transportation fuels to be reduced by at least ten percent in 2020" (ARB, 48). "In April 2011 California adopted a 33 percent RPS" (Appendix, 21). The efficiency mandates varied by different buildings and appliances (ARB, 37).	No, introduced concurrently	Borenstein et al. (2019)
RECLAIM: Regional Clean Air Incentives Market (NOx)	South Coast Air Basin in Southern California	1994	• 1970 CAAA: NSPS, BACT • In 1990 South Coast Air Basin was only nonattainment area for NOx emissions. Thus, the NOx emissions standards may have been more stringent here than other areas. • Prior to RECLAIM, command-and-control program emphasized advanced control technologies, which they adopted in late 1989.	Relaxed	Fowlie, Holland and Mansur (2012)
Acid Rain Program (SO2)	US	1990	• 1970 CAAA: NSPS, BACT • 1977 CAAA: New coal plants must operate with scrubbers and achieve a certain reduction in potential SO2 emissions (70-90%)	No	Joskow and Schmalensee (1998)
Regional Greenhouse Gas Initiative (CO2)	Northeast US	2009	• All states except Virginia had implemented an RPS • Complementary measures as part of RGGI separate from its emissions cap: RGGI auction revenues used for "energy efficiency purposes". • CAAA in force: favoring substitution to cleaner sources from coal powered generation	No	Murray and Maniloff (2015); Kim and Kim (2016)
European Union Emissions Trading System (CO2)	EU	2005	• Directive 2001/77/EC in 2001 set country specific targets for adoption of renewable energy production. • Directive 2003/30/EC in 2003 promoted biofuels for EU transport.	No	Deschênes, Greenstone and Shapiro (2017)

G 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 6 summarizes the analysis we describe below.

G.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.

G.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 PM2.5. A simple estimate of the impact of the ETS is that ambient pollution would fall by an amount equal to 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 $\mu g/m^3$. Guttikunda, Nishadh and Jawahar (2019) use pollution inventories combined with an atmospheric dispersion model to apportion ambient particulate concentrations in Indian cities to different sources.[36] The authors estimate annual average ambient PM 2.5 concentrations in the city at 88.5 $\mu g/m^3$, with 32% (or 28.32 $\mu g/m^3$) coming from local industry. Then the Surat ETS applied to all plants in the city would reduce fine particulate pollution by $0.30 \times 28.32 = 8.5 \mu g/m^3$ (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 $\mu g/m^3$ 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

[36] Their updated assessment for Surat is available at: `https://urbanemissions.info/india-apna/surat-india/`.

were to assume that an ETS were implemented in Surat for 70 years (roughly the current life expectancy in India), reducing pollution each year by 8.5 $\mu g/m^3$, 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.

References cited within Appendices

Apte, Joshua S., Michael Brauer, Aaron J. Cohen, Majid Ezzati, and C. Arden Pope III. 2018. "Ambient PM2.5 Reduces Global and Regional Life Expectancy." *Environmental Science and Technology Letters*, 5(9): 546–551.

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): 3953–3977.

Deschênes, Olivier, Michael Greenstone, and Joseph S. Shapiro. 2017. "Defensive Investments and the Demand for Air Quality: Evidence from the NOx Budget Program." *American Economic Review*, 107(10): 2958–2989.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan. 2018. "The Value of Regulatory Discretion: Estimates From Environmental Inspections in India." *Econometrica*, 86(6): 2123–2160.

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.

Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur. 2012. "What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NOx Trading Program." *American Economic Review*, 102(2): 965–993.

Guttikunda, Sarath K., K. A. Nishadh, and Puja Jawahar. 2019. "Air pollution knowledge assessments (APnA) for 20 Indian cities." *Urban Climate*, 27: 124–141.

Joskow, Paul, and Richard Schmalensee. 1998. "The Political Economy of Market-Based Environmental Policy: The U.S. Acid Rain Program." *Journal of Law and Economics*, 41: 37–83.

Kim, Man-Keun, and Taehoo Kim. 2016. "Estimating Impact of Regional Greenhouse Gas Initiative on Coal to Gas Switching Using Synthetic Control Methods." *Energy Economics*, 59: 328–335.

Murray, Brian C., and Peter T. Maniloff. 2015. "Why Have Greenhouse Emissions in RGGI States Declined? An Econometric Attribution to Economic, Energy Market, and Policy Factors." *Energy Economics*, 51: 581–589.

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.

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.

Sudarshan, Anant. 2023. "Introduction to Development Engineering." In . , ed. Temina Madon, Ashok J. Gadgil, Richard Anderson, Lorenzo Casaburi, Kenneth Lee and Arman Rezaee, 161–182. Springer International Publishing.