

NBER WORKING PAPER SERIES

TOO LITTLE TOO LATE? NATIONWIDE CARBON TRADING IN CHINA

Shuang Zhang
Douglas Almond

Working Paper 35924
<http://www.nber.org/papers/w35924>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2026

We thank the editor and three referees for helpful feedback in improving the paper. We thank Xiaohan Dong, Ruixue Li, Sophia Qin and Yuting Wang for excellent research assistance. For helpful comments, we thank Sam Asher, Patrick Bolton, Lorenzo Casaburi, Eugenie Dugoua, Richard Green, Ryan Kellogg, Mirabelle Muûls, Ralf Martin, Albert Park, Simon Quinn, Gregor Singer, Mar Reguant, Eric Zou, and seminar participants at Imperial College London, University of Zurich, CREST, LSE/Imperial Workshop in Environmental Economics, University of Warwick Data, Information and the Environment Workshop, and CEPR Summer Symposium in Climate Change and the Environment. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Shuang Zhang and Douglas Almond. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Too Little too Late? Nationwide Carbon Trading in China
Shuang Zhang and Douglas Almond
NBER Working Paper No. 35924
October 2026
JEL No. Q48, Q5, Q53, Q54

ABSTRACT

China launched a nationwide, rate-based carbon market in 2021. Covering the world's top CO₂ emitter, compliance depends on the intensity of carbon emissions rather than an absolute cap. We draw on the theoretical literature of rate-based systems to provide the first empirical evidence on an early adopter. Using plant-level data from satellite imagery and machine learning methods, we find that high-efficiency power plants increased electricity generation and CO₂ emissions by approximately 5% relative to low-efficiency plants. Because its carbon market regulates nearly all coal plants in China, we focus on the differential responses by plant efficiency, finding no evidence of a differential reduction in emission intensity (though these estimates are imprecise). Relative to coal plants in Indonesia and the Philippines, which had no carbon market, China's high-efficiency plants also increased emissions, while the estimate for its low-efficiency plants is small and insignificant. This pattern of responses suggests that emissions increases of high-efficiency plants exceeded substitution away from low-efficiency plants. Underscoring the production stimulus to more efficient Chinese plants, air quality worsened downwind of high-efficiency plants relative to low-efficiency plants, compromising air quality in relatively clean areas of China. If China's goal is to reduce its carbon emissions, then benchmarks will need to be tightened within the existing system or an alternative system that does not simultaneously incentivize electricity production will need to be adopted.

Shuang Zhang
Imperial College London
Department of Economics
and Public Policy
and CEPR
s.zhang@imperial.ac.uk

Douglas Almond
Columbia University
School of International
and Public Affairs
and Department of Economics
and NBER
da2152@columbia.edu

1 Introduction

Greenhouse gas emissions from emerging economies will drive the future course of climate change (Greenstone, 2024). China stands out for three reasons. First, it generates one-third of global emissions. Appendix Figure A.1 shows that its CO_2 emissions are more than double those of the US and more than three times those of the EU or India. GHG emissions continue to rise in China, while they are falling in developed countries. Second, China launched a nationwide carbon market in July 2021, the scale of which is unprecedented. The program began in the power sector, which alone accounts for 15% of global carbon emissions. Third, China’s market is rate-based rather than cap and trade, and no rate-based carbon market has been evaluated empirically. China’s design has been taken as a model for carbon trading elsewhere in Asia and the developing world.

Under cap-and-trade programs, compliance depends on keeping the absolute level of emissions below a preset cap. In contrast, China’s national ETS is rate-based, where compliance depends on emission intensity (the ratio of emissions to current output). Because allowances increase with output, the design implicitly subsidizes output. Theory predicts that plants below the intensity benchmark expand generation, while plants above it rely more on reducing emission intensity (Fischer, 2001; Fischer and Newell, 2008; Goulder et al., 2022, 2023). A rate-based system could therefore be less cost-effective, because it provides little incentive for reduction of intended output. It also has potentially better compatibility with economic growth (Goulder et al., 2022), making it appealing to other emerging economies. Indonesia launched a rate-based ETS in 2023, and India has announced that it will start a similar program in 2026. To our knowledge, these predictions have not been tested for a national carbon market. That China implemented relatively accommodating benchmarks in its ETS is suggested by recent carbon prices of around \$14 per ton.

Testing these predictions requires plant-level data on output and carbon emissions, which have not been available. While China took a major step in establishing a national carbon market, the challenges of imperfect monitoring and insufficient transparency are familiar (Karplus, Zhang, and Almond, 2018; Greenstone, He, Jia, and Liu, 2022). Indeed, China’s Ministry of Ecology and Environment has documented a number of cases of falsified carbon data submitted by power plants (Reuters, March 2022).¹

¹<https://www.reuters.com/world/china/china-slams-firms-falsifying-carbon-data-2022-03-14/>

Here, we analyze new satellite imagery data on carbon emissions for 413 power plants covering about half of China’s coal-fired generation. As these data are difficult to manipulate, they can serve as an objective benchmark for carbon emissions data reported by firms. In collaboration with [Climate Trace](#), we use plant-level output and carbon emissions data from high-resolution satellite images and a machine learning algorithm. We benchmark these data against carbon data reported by plants annually.

We provide the first empirical test of a rate-based emissions trading system, situated in the largest carbon market in the world. We estimate output and emissions changes at high-efficiency plants relative to low-efficiency plants. Given coal’s continued dominance in China’s power sector, we also examine whether the ETS alters the spatial allocation of air pollution, incorporating an atmospheric dispersal model.

Within China, we have three main findings. First, initially high-efficiency plants increased electricity generation and satellite-measured CO_2 emissions by about 5% relative to initially low-efficiency plants, consistent with the prediction for rate-based systems. Second, we find no evidence of a differential reduction in emission intensity, the margin the carbon market was intended to move, though those estimates are imprecise. Third, using the HYSPLIT atmospheric dispersal model, we find that ambient air pollution increased downwind of high-efficiency plants relative to low-efficiency plants, in areas cleaner before the ETS.

We then turn to what the differential responses imply for aggregate emissions. Net additions to output induced by the ETS clearly raise total emissions, while the reallocation of output to high-efficiency plants lowers them. We provide additional results that are consistent with net addition. The differential response is no larger in provinces where demand recovered fastest post-COVID or where coal was scarcest, and it persists when plants are aggregated to their parent companies. These patterns are only suggestive, since essentially every Chinese coal plant is regulated, and that requires a benchmark outside the China ETS. We therefore compare Chinese plants with coal plants in Indonesia and the Philippines, which followed similar power-sector trajectories without introducing a carbon market. China’s high-efficiency plants increased carbon emissions relative to them, while the estimate for its low-efficiency plants is small and insignificant. A bounding exercise suggests a conservative lower bound on the net change among the Chinese plants we observe, an increase of about 14 million tons of CO_2 a year. The comparison is imperfect, our pre-period is

short and our two comparison countries had different Pandemic policies from China’s. We interpret the cross-country comparison as more informative on the direction of effects rather than on their precise magnitude.

This study contributes to the broader literature on market-based approaches to climate change. First, we introduce readily-scalable data for research on tensions between the environment and development. A common theme is that poor monitoring undercuts the effectiveness of environmental policies in developing countries (Dufflo et al., 2013, 2018). Findings on emission markets in developed countries do not generalize to settings where monitoring is relatively weak.² These data are particularly useful where monitoring is both challenging and emissions are expected to grow most rapidly.

Second, we advance the literature on the efficacy of market-based instruments under economic and political constraints. Rate-based systems are particularly appealing to rapidly growing economies that fear impinging on their development trajectories. In the same second-best spirit, the US has not been able to adopt a nationwide cap-and-trade program and has instead turned to subsidizing green industries and purchases, often with domestic content requirements (Bistline et al., 2023).

Third, China’s national ETS is the culmination of policy experimentation at more local levels. Such local experimentation is a common pathway to the introduction of other signature national policies (Wang and Yang, 2025). We build on previous work analyzing pilot ETS markets since 2014 (Almond and Zhang, 2021; Cui et al., 2021, 2023; Zhang et al., 2020; Chen and Lin, 2021). That literature finds that carbon prices need not approach their social costs to change behavior. Finally, because the national ETS’s structure and incentives did not mirror the pilot programs, it is difficult to extrapolate from that literature.

2 China’s National ETS

China’s national ETS regulates 2,257 power plants that emitted 26,000 tons or more of CO_2 in any year 2013-2018, covering approximately 40% of China’s total carbon emissions. In every compliance

²Greenstone et al. (2023) is an exception that studied a particulate matter emissions market in India and built a continuous emissions monitoring system to measure plant-level particulate matter. To monitor a carbon market with *high-quality* carbon data, however, is more complex and to our knowledge has not been done in developing countries.

period, each generation unit of a regulated plant receives a free allocation of emission allowances based on the government’s assigned benchmark (maximum allowed emission intensity) and its end-of-period power generation.

The assigned benchmarks are fuel specific and generally set below the average initial emission intensities across regulated facilities. For conventional coal-fired units, the assigned benchmark is $0.877\text{ }tCO_2/MWh$ for units with rated power of at least 400 MW, and $0.979\text{ }tCO_2/MWh$ below that (EDF, 2021). Appendix B gives the benchmark schedule and the classification rule.

China’s program shares an important feature of cap-and-trade programs in the introduction of a national carbon market for allowance trading. A permit must be held for each ton of CO_2 emitted, and these permits are traded to achieve a market-clearing price. The key difference is how emission allowances are allocated. In cap-and-trade programs, a facility’s allowances usually do not depend on its end-of-period output. Under a rate-based system, the number of allowances is endogenous to the facility’s chosen output.

Theoretical literature A rate-based system implicitly subsidizes the intended output (Fischer, 2001; Fischer and Newell, 2008). Theoretical analysis of China’s rate-based ETS (Goulder et al., 2022, 2023) suggests that the implicit output subsidy provides differential incentives to firms.

More efficient firms with initial emission intensity below the benchmark receive more allowances than they need for compliance. They have incentives to *increase* production and sell the extra allowances. Less efficient firms with initial intensity above the benchmark can cut the cost of buying allowances by reducing production or reducing emission intensity. Reducing output is implicitly taxed, however, because it also reduces the allowances received. These firms therefore rely more on reducing emission intensity (Goulder et al., 2022, 2023).

Challenges in monitoring Under China’s National ETS, monitoring, reporting and verifying (MRV) is delegated to local governments, which is another challenge to the program’s effectiveness. Regulated facilities must submit annual carbon reports, and provincial governments then assign third-party companies to verify them. Conflicts of interest may arise because the promotion of local officials depends largely on economic performance.

Carbon prices Figure A.2 shows the carbon prices in the national program since July 2021. The prices started relatively low at \$6 per ton of CO₂ and gradually increased to more than \$14. Although prices are still much lower than in many other existing ETS frameworks, Chinese firms were found to have responded in their CO₂ emissions at low prices of \$6 per ton during the regional pilot programs (Cui et al., 2021; Zhang et al., 2020; Chen and Lin, 2021).

High-efficiency and low-efficiency plants A coal-fired plant typically has multiple generation units, which use sub-critical, super-critical or ultra-supercritical technology, in ascending order of efficiency. The emission-intensity benchmark applies to generation units, and a unit’s efficiency depends mainly on its technology. Figure A.8 shows the distribution of unit-level emission intensities relative to the benchmark in 2019: about 91% of sub-critical units lie above the benchmark, while only 37% of super-critical and ultra-supercritical units do. Most Chinese plants use a single technology: 54% are sub-critical only and 35% are super-critical or ultra-supercritical only. We focus on these “specialized” plants. We define *high-efficiency* plants as super-critical or ultra-supercritical and *low-efficiency* plants as sub-critical. We classify plants by predetermined technology rather than by measured intensity, because measured intensity is itself a potential outcome of the policy. We also check the robustness of classifying plants instead by their baseline position relative to the benchmark.

3 Data and Methods for Power Generation and Carbon Emissions

We collaborate with Climate Trace NGO to adopt a two-step approach based on high-resolution and high-frequency satellite images using the methodology in Couture et al. (2024): a machine learning algorithm predicts electricity generation from water vapor plumes, which is then combined with plant-specific emission intensities to estimate emissions.

3.1 Step 1: Plant-level electricity generation

Direct measurement of CO₂ emissions via satellite imagery is challenging, as flue gas is generally less visible at plants that are more efficient or have air pollution controls. The methodology in Couture et al. (2024) instead uses visible water vapor plumes, which are robust indicators of electricity

generation. Machine learning models are trained on labeled satellite imagery to predict a plant’s hourly capacity factor, which we aggregate to daily and annual plant-level generation. The satellites pass over between 7:30 and 11:30am local time, which we return to below. We describe the imagery, the training data, the image-selection criteria, and the machine learning model in Appendix A.

The methodology in Couture et al. (2024) is applied by Climate Trace to Chinese power plants in 2019-2023 to estimate annual electricity generation at the plant level. We analyze annual measures for comparison with ground-truth data on generation.³

We validate the satellite generation against full-day generation reported by publicly listed companies, which report plant-level annual generation and own about a quarter of the plants in our data. The satellite levels are lower than the reported levels, but the trends agree, and the same holds for Indian plants (Appendix Figures A.3a, A.3b and A.4). A possible explanation is the satellites’ morning overpass times. If Chinese plants ramp up strongly in the afternoon or evening, when the satellites do not observe them, and such ramps are less common in the training sample, the image-based data will understate levels.

A related question is whether the morning overpass biases our difference-in-differences estimates, for example if low-efficiency plants increasingly run as evening peak-shaving units while high-efficiency plants run as baseload. Our estimates compare the *change* at high-efficiency plants with the *change* at low-efficiency plants. The plant and year fixed effects absorb any level bias that is constant or that changes similarly for both plant types. Only an increase in the bias specific to low-efficiency plants after the ETS could generate a spurious differential. For plants that report full-day generation, the satellite data and reported data change similarly over 2019-2021 for both plant types (Appendix Figure A.11), so the proportional level gap is common to both. Solar generation, which drives the peak-shaving pattern behind this concern, was a small part of China’s generation in our period, 3.1% to 6.2% over 2019-2023.⁴ In addition, we have two measures that do not depend on the overpass window, plant-reported CO_2 and ground-monitor pollution, and we estimate the same specification with each alongside the satellite measure.

³Climate Trace releases monthly source-level estimates only from 2021, when the number of usable satellite images increased substantially. Our design requires pre-ETS outcomes. Climate Trace advised that they are more confident in the annual estimates than in the monthly estimates before 2021, and we follow their recommendation.

⁴Data source: Ember, Yearly Electricity Data, ember-energy.org/data/yearly-electricity-data

3.2 Step 2: Plant-level carbon emissions

Plant-level CO_2 emissions are estimated by combining two separately estimated measures: 1) the electricity generation predicted from Step 1; and 2) the estimated emission factor. The emission factor is estimated using the type of fuel and the type of generation technology in an engineering model (Couture et al., 2024). We refer to it as the *engineering measure* of the carbon intensity. This measure captures changes in emission intensity from fuel-switching or technology upgrades. However, it does not detect intensity improvements that occur *within* a given fuel-technology combination, such as incremental gains in combustion efficiency or improvements in heat rate at a plant whose technology class remains unchanged.

We also collect data on reported carbon emissions from power plants. Regulated plants have been required to submit annual carbon reports since 2019, two years before the national ETS. In the report, firms report annual CO_2 emissions at the plant level. These data are publicly available. We assess whether the trends in carbon emissions before and after the ETS are consistent across the satellite and plant-reported data.

We construct a *reported measure* of carbon intensity by using reported CO_2 emissions divided by satellite-based generation. Because the reported measure can detect changes in combustion efficiency or fuel quality, we use it as the primary measure for our intensity analysis and report the engineering measure as a robustness check.

3.3 Summary statistics

Our main analysis sample is a balanced panel of 413 coal-fired plants from 2019 to 2023. We exclude plants in the pilot ETS regions, which had faced a regional carbon price since 2013-2014, so the 2021 national launch is not a clean change in treatment there.⁵

Appendix Table A.1 reports plant-level summary statistics. Column (1) uses the satellite data: plants average 1,060 MW of capacity, a 41% annual utilization rate, 4,100 GWh of generation, and 3,671 kilotons of CO_2 a year. Column (2) uses the data reported by plants, the 242 plants that report consistently over the five years. The reporting sample has a slightly larger capacity on average and reports 4,755 kilotons of CO_2 a year.

⁵The excluded pilot sample is small: in the satellite data, there are 44 plants in pilot regions versus 413 plants in non-pilot regions.

We note two features of the satellite sample. The first is representativeness. Appendix Table A.2 compares the 1,094 generating units at our 413 sample plants with the 1,708 units in data from Global Energy Monitor that are not in our sample. Sample units are on average larger and more often super-critical or ultra-supercritical. On the most relevant dimension for the ETS policy, these two datasets are nearly identical: roughly three-quarters of units in both are above the ETS benchmark. The second is the confidence of the satellite estimates, which is provided by Climate Trace, and most Chinese coal plants are rated at medium-to-high confidence. We conduct robustness tests by restricting the analysis sample to plants with lower generation uncertainty (higher confidence) in the satellite data.

4 Data and Methods for Ambient Air Pollution

4.1 Ambient pollution data

We obtain measures of ambient air pollution from two sources. First, we compile pollution data in 2017-2024 from the real-time monitoring and reporting system of China’s Ministry of Ecology and Environment. Hourly concentration readings of $PM_{2.5}$, PM_{10} , SO_2 , NO_2 and O_3 are available from about 2,000 pollution monitors spread across China. Second, to complement ground monitor pollution measures, we also use high-resolution ($0.01^\circ \times 0.01^\circ$) satellite-derived estimates of monthly average $PM_{2.5}$ in 2018-2022 (Shen et al., 2024).

4.2 Atmospheric dispersal model

To consider transport of air pollution in the environment, we incorporate an atmospheric dispersal model following the recent literature (Grainger and Ruangmas, 2018; Hernandez-Cortes and Meng, 2023). We generate particle trajectories from emitting power plants using the Hybrid Single Particle Lagrangian Integrated Trajectory (HYSPLIT) model, developed by the US National Oceanographic and Atmospheric Administration (NOAA). For meteorology we use NOAA’s Global Data Assimilation System, specifically the GDAS1 dataset.

We feed the location of the plant, the average stack height of 150 meters, the date and hour into HYSPLIT to generate forward trajectories from power plants. For each plant, we generate four trajectories per day at 8am, 12pm, 4pm and 8pm. We let each trajectory run for 72 hours. Each

trajectory consists of hourly locations of an air parcel. After obtaining the trajectories, we match the pollution monitors to the trajectories. Specifically, we define a monitor as “on a trajectory” of a plant at a specific hour if the trajectory intersects a 1 km circle centered on the monitor.

We then construct a balanced panel data from matched pollution monitors on the trajectory of each plant during 2017-2024. We exclude monitors that are only on a trajectory for less than 1% of the time. Using pollution readings from these monitors, we generate monthly pollution measures for each plant by taking a weighted average of the hourly pollution readings in a month of monitors matched to that plant. The weights are the number of times that a monitor is on the plant’s trajectories, applying the exponential decay following [Hernandez-Cortes and Meng \(2023\)](#).

We repeat this exercise with satellite-derived monthly $PM_{2.5}$, matching grid cells rather than monitors: a cell counts as matched when a plant’s trajectory passes through it, and the same 1% threshold and distance-decay weighting then apply.

4.3 Summary statistics

Figure [A.9a](#) shows the locations of high-efficiency plants and low-efficiency plants separately. To compare characteristics of affected areas, we compare counties where the majority of all trajectories passing during our period are from a high-efficiency plant with counties affected by low-efficiency plants (Figure [A.9b](#)). Appendix Table [A.3](#) reports this comparison at three thresholds for the majority rule, 60%, 70% and 80%, measuring county pollution as the 2019 average across ground monitors within 50 km of the county. Counties downwind of high-efficiency plants were cleaner before the ETS, and together they hold fewer people, 132 million against 317 million at the 70% threshold.

5 Empirical Strategy

Our primary analysis compares coal-fired power plants within China by their initial generation efficiency, before and after the ETS. Although almost all coal plants are regulated, efficiency before the ETS determines the incentives they face under the policy. As discussed in Section [2](#), theory predicts that plants with initial intensities below the benchmark increase generation, while plants above the benchmark rely more on reducing emission intensity to comply. We test these predictions

by comparing the two types of plants. We classify plants by their predetermined generation technology: high-efficiency plants use super-critical or ultra-supercritical technology exclusively, and low-efficiency plants use sub-critical technology exclusively.

Our DID coefficient $\hat{\alpha}$ in equation (1) estimates the *differential* response between high- and low-efficiency plants, not the aggregate effect of the ETS. Both groups are regulated, but they face different incentives under the rate-based design. This differential is the parameter of interest and the testable implication of the rate-based design. We discuss what it implies for aggregate emissions in Section 7.

$$Y_{it} = \alpha H_i * Post_t + \gamma_i + \theta_t + X_{it}\beta + \epsilon_{it} \quad (1)$$

where i denotes the power plant and t denotes time in year. The dependent variable Y_{it} denotes operation, electricity generation, carbon emissions, and emission intensity of plant i in year t . The key independent variable is the interaction term $H_i * Post_t$. H_i equals 1 for high-efficiency plants and zero for low-efficiency plants. The policy indicator $Post_t$ equals 1 starting 2021, and zero otherwise. We include plant fixed effects γ_i to absorb plant-specific factors and calendar year dummies θ_t to capture annual changes common to all plants. We also include additional controls X_{it} , including plant-by-year capacity and province-by-year GDP.

The key identification assumption is that outcomes would have followed their baseline trends by initial technology prior to the ETS. We examine pre-existing trends in the event-study analysis and find them modest and roughly parallel, suggesting that any anticipatory responses were similar across plant types. Our pre-ETS period is short because of when the satellite data start. As a robustness check, we also use a SDID specification that matches the pre-ETS trends of outcomes between the two types of plants. The same caveat of only two pre-ETS years applies: the unit weights are determined by the 2019 and 2020 levels and the 2019 to 2020 change, so the SDID should be interpreted as a pre-trend adjustment.

The remaining concern for inference centers on other changes around 2021 that could differentially affect the two plant types. We have read through contemporaneous policies and identified three potential confounders, which we address below: China’s coal structural adjustment policy, the rebound in electricity demand following the COVID-19 shock, and the coal supply and power crisis

in the autumn of 2021.

First, the coal structural adjustment policy encourages the retirement of sub-critical units and the approval of new large high-efficiency capacity. Because our main sample is a balanced panel of 413 plants that all generate throughout 2019-2023, plant retirements and new entrants do not contribute to the estimated differential. In addition, estimating the main specification with plant capacity as the outcome yields small and insignificant estimates (Appendix Table A.6), so the differential mainly reflects higher utilization of existing capacity rather than capacity additions.

Second, the recovery of electricity demand after the 2020 COVID-19 shock may have disproportionately favored dispatch of high-efficiency plants even absent the ETS. We examine the demand channel directly by testing whether the estimated effects vary with the post-COVID recovery in GDP across provinces. If the estimated effects reflected differential dispatch by plant type in the post-COVID demand recovery, we would expect larger effects where electricity demand grew more. We use the provincial GDP growth rate (from 2020 to 2021-2023) as a proxy measure for demand growth to interact with our main regressor, controlling for province-by-year fixed effects.

Relatedly, reallocation across plants, if any, is more likely within a parent company. We therefore aggregate plants to their 121 parent companies and compare companies by baseline average efficiency. Because the comparison is between companies, our estimates do not reflect reallocation within a company.

Third, a severe coal supply and power crisis hit China in the autumn of 2021: spot coal prices spiked, several provinces rationed power, and dispatch and pricing responses favored efficient units. The crisis was a supply shock whose severity varied across provinces: a province with ample coal of its own had less need to reshuffle its plants, while a province that depends on coal brought in from elsewhere had more need. If the crisis and the dispatch response drove our results, the differential should therefore be larger where pre-crisis supply was tighter. We test this by interacting our main regressor with province-level pre-crisis coal self-sufficiency, measured by raw coal output divided by coal consumption in 2019.

To complement the within-China design and learn about the implications for aggregate emissions, we compare Chinese coal plants with foreign coal plants that faced no carbon market. We select comparison countries whose macroeconomic and power-sector trajectories are similar to China's through the COVID-19 period: Indonesia and the Philippines are the relatively close Asian matches

to China on GDP, total electricity consumption, and power-sector CO_2 before and after 2020 (Appendix Figure A.5), and neither had a carbon market during this period. We look at data until 2022 for the cross-country comparison because Climate Trace covers the comparison countries only through 2022, and Indonesia launched its own rate-based ETS in 2023. Using satellite-based outcomes for 88 Indonesian and Philippine plants, we estimate two designs. First, a triple difference interacts the high- versus low-efficiency with pre versus post and with China versus the comparison countries, with plant and country-by-year fixed effects. One caveat is that there are only six high-efficiency plants in these comparison countries. Second, we also estimate two split DID designs comparing China’s high-efficiency plants and, separately, China’s low-efficiency plants with the pooled comparison plants. We acknowledge the limitations of these cross-country comparisons: the comparison group includes only two countries, with a two-year pre-period; China’s COVID-19 policies differed from theirs; and the triple difference rests on only six high-efficiency plants abroad. We therefore interpret these results as suggestive of the direction of aggregate emissions rather than as primary findings.

6 Results

6.1 Within China

6.1.1 Electricity generation and carbon emissions

We start by describing plant-level outcomes before and after the national ETS (without any regression adjustment) in Figure 1. Two patterns emerge. First, the trends before 2021 between these two types of plants are relatively parallel. Second, high-efficiency plants increase their electricity generation and carbon emissions, while the generation and carbon emissions of low-efficiency plants remain unchanged after 2021.

Figure 2 presents event-study estimates. Figure 2a shows the estimates of annual utilization rates. The trend prior to 2021 is relatively flat. After the launch of the ETS in 2021, high-efficiency plants increase their utilization rates. Annual electricity generation shows the same pattern in Figure 2b. These findings are consistent with the theoretical predictions for rate-based systems.

We conduct event-study analyses of CO_2 emissions separately for satellite and plant-reported

data. In Figure 2c, satellite emissions show parallel trends before the ETS, and high-efficiency plants increase emissions after 2021. Figure 2d shows similar patterns in the plant-reported data.

Table 1 reports DID estimates for these outcomes. The regression-adjusted estimates are similar. In Panel A, estimates using satellite data suggest that high-efficiency plants increase their annual utilization rate by 2.6 percentage points after the ETS, which is about 6.3% of the pre-ETS mean. They increase their annual electricity generation by about 210 GWh after the ETS and CO_2 emissions by 186 kilotons, both about 5.2% of the pre-ETS mean. In Panel B, estimates using plant-reported data suggest a relatively greater increase in CO_2 emissions among high-efficiency plants. However, the difference between the two estimates is not statistically significant.⁶

Our estimates are robust to alternative specifications. Including province-by-year fixed effects absorbs province-specific demand shocks (Appendix Table A.9). Including firm-by-year fixed effects controls for parent-company-level confounders, comparing plants of different efficiency within the same firm (Appendix Table A.10). SDID estimation, which reweights control units to adjust pre-trends, and allowing plant-specific linear trends both leave the differential positive and significant, with somewhat larger generation and CO_2 estimates than in Table 1 (Appendix Tables A.11 and A.7).

For robustness tests, we classify plants by their baseline position relative to the ETS benchmark rather than by combustion technology gives similar, finding slightly larger estimates in Appendix Table A.12. The estimates are also robust in the satellite subsamples with the lowest 75%, 50% and 33% generation uncertainty (higher confidence of data quality) in Appendix Table A.13.

Finally, we estimate the differential response of emission intensities. As with our other outcomes, this DID identifies only the differential between the two plant types, which would be near zero if both improved intensity to a similar degree. The reported estimates in Appendix Table A.14 Panel A are small and imprecisely estimated. The engineering estimates in Panel B are near zero, though that measure is constructed from fixed fuel-technology factors and cannot detect improvements within a technology class. Appendix Figure A.10(a) plots reported intensity for the two plant types. The low-efficiency intensity is roughly flat while the high-efficiency intensity increases modestly after the

⁶The estimates are in kilotonnes and scale with the level of emissions in each panel. The reported sample has larger plants on average (Appendix Table A.1), and satellite generation is systematically lower than firm-reported generation for the same plants (Appendix Figure A.3). Together these account for the difference between the two panels' pre-ETS means.

ETS, which explains the small and imprecise estimates in Panel A. In the corresponding event study in Appendix Figure A.10(b), all post-ETS estimates are statistically insignificant. At the company level, reported intensity likewise shows no decline for high-average-efficiency relative to low-average-efficiency companies (Appendix Table A.8). Taken together, we find no evidence of a differential reduction in emission intensity. We also note the limitations: the estimates are imprecise, so they do not establish that intensity was unchanged for either type of plant.

6.1.2 Examine potential confounders

On the coal structural adjustment policy, we estimate the main specification with plant capacity as the outcome. The estimate is 25.9 MW without province-by-year fixed effects and 11.9 MW with them, both statistically insignificant against a pre-ETS mean of 1,051 MW (Appendix Table A.6). We therefore find no evidence that high-efficiency plants added units relative to low-efficiency plants within the balanced panel.

On the post-COVID demand recovery, we conduct two exercises. First, in the heterogeneity test, the interaction of our main regressor with the GDP growth rate is small and statistically insignificant for every outcome, while the main effect remains significant (Appendix Table A.4). The differential is therefore not concentrated where GDP recovered faster, which is inconsistent with the demand channel as the primary driver. Second, comparing companies by baseline average efficiency, high-average-efficiency companies increase generation and CO_2 emissions relative to low-average-efficiency companies, with little change in reported emission intensity (Appendix Table A.8). Since the contrast is across firms rather than within them, output shuffled between a firm’s own plants cannot produce it. The flat generation of low-efficiency plants is likewise more consistent with high-efficiency plants expanding output than with firms reshuffling dispatch across their portfolio.

On the 2021 coal supply and power crisis, the interaction of the main regressor with province-level coal self-sufficiency is statistically insignificant for every outcome (Appendix Table A.5). The emergency dispatch and pricing measures applied nationally, but the amount of re-dispatch depended on how binding local supply was, so a crisis-driven differential should have been concentrated in the provinces the crisis hit hardest. We find that this is not the case.

6.1.3 Ambient Air Pollution

Using monitor-based pollution measures that incorporate the HYSPLIT model (Section 4), we conduct an event study comparing pollution concentrations on the trajectories of high-efficiency versus low-efficiency plants before and after the program. In Figure 3, the difference in the concentration of SO_2 was relatively flat before the ETS, but the concentrations of SO_2 increase more in areas affected by high-efficiency plants after the ETS. We find similar increases in concentrations of $PM_{2.5}$ and PM_{10} . Importantly, our pollution data extend back to 2017, providing four pre-policy years. The absence of differential pre-trends over this longer period further supports our identification assumptions. Considering that areas affected by high-efficiency plants are initially cleaner, our findings suggest that the differential response to the program degraded air quality where it had previously been better. For the placebo test, we examine two pollutants whose ambient concentrations are less closely tied to coal plant output: NO_2 and O_3 . Indeed, we find little change.

We report DID estimates in Table 2. Monitor data cover a longer pre-ETS period, which lets us control for plant-specific time trends. Consistent with Figure 3, we estimate that SO_2 concentrations increase in areas affected by high-efficiency plants by 4.4%, $PM_{2.5}$ by 3.2% and PM_{10} by 2.8%, after the ETS. In the placebo test, we find little change in NO_2 and O_3 .

The estimates are robust to alternative stack heights, 120-hour trajectories, and monitors within 5 km of the trajectories (Appendix Table A.15). Using satellite-derived monthly $PM_{2.5}$ instead of monitors, concentrations within 10 km of the trajectories of high-efficiency plants increase by 2.7% (Appendix Table A.17). Assigning exposure by distance instead of HYSPLIT, averaging satellite $PM_{2.5}$ over all grid cells within 5, 10 or 15 km of each plant, concentrations near high-efficiency plants increase by about 2.6% at each radius (Appendix Table A.18). Without the HYSPLIT model, measuring monthly $PM_{2.5}$ as the inverse-distance-weighted average of monitors within 50 km, no pollution measure changes significantly (Appendix Table A.16). The estimates are smaller and the standard errors larger than in Table 2, which suggests noisier exposure measures without the model.

Two points are worth noting for interpreting the pollution results. First, whether an atmospheric dispersal model is essential depends on the spatial density of the exposure data. Ground monitors are sparse, so a radius average without HYSPLIT could cover many non-downwind monitors, while HYSPLIT focuses on the monitors the plume passes. In contrast, the satellite grid is dense, so a

distance buffer already contains the downwind cells. Appendix Table A.18 shows similar results using satellite-measured $PM_{2.5}$ in nearby grid cells without using HYSPLIT. Second, what the pollution increases imply for total population exposure depends on the counterfactual. If the extra generation at high-efficiency plants displaced generation that low-efficiency plants would otherwise have supplied, pollution would have gone down where more people live. Because our evidence on generation and carbon emissions above is more consistent with net expansion than with reallocation, the pollution increases we document are more likely net additions to exposure, and they fell on counties that were cleaner before the ETS and that together hold fewer people (Appendix Table A.3). Our design identifies where the pollution increases go rather than the aggregate change in national population-weighted exposure.

6.2 Comparison with other Asian countries

We compare Chinese coal plants with the Indonesian and Philippine plants. In the triple-difference estimates in Appendix Table A.19, China’s high-efficiency plants increase generation by 298 GWh and CO_2 emissions by 264 kilotons relative to China’s low-efficiency plants following the ETS, compared to the Asian comparison plants. The split DID in Appendix Table A.20 show the same pattern. Relative to the pooled comparison plants, China’s high-efficiency plants increase utilization by 2.1 percentage points, generation by 333 GWh, and CO_2 emissions by 297 kilotons, while China’s low-efficiency plants show small and insignificant changes. Appendix Figures A.6 and A.7 show the underlying raw means: within high-efficiency plants, China increases above the comparison plants after 2021, while within low-efficiency plants, China and the comparison plants move similarly. These patterns are more consistent with predictions of the rate-based ETS. A China-wide demand shock such as the post-COVID recovery would move both Chinese plant types relative to the comparison plants.

Pooling both plant types, the split DID suggests that Chinese plants in our sample increase CO_2 by about 134 kilotons per plant relative to the comparison plants. We interpret this as pointing toward net expansion rather than output shifted away from low-efficiency plants. The low-efficiency estimate is imprecise, so a decline at those plants cannot be ruled out. To bound how far such a decline could offset the increase, we take both estimates at the low end of their 95% confidence intervals. The high-efficiency estimate is precise, so even the low end of its interval is 160 kilotons

per plant, or 26 million tons a year across the 163 high-efficiency plants. The low end for the low-efficiency plants is a decline of 49 kilotons per plant, or 12 million tons a year across 250 plants. The net is an increase of about 14 million tons a year.

7 Summary and Discussion

National governments have adopted market-based approaches to curtail greenhouse gas emissions. Emissions trading has recently spread to the faster-growing developing economies that are expected to dominate future global emissions. Beyond theoretical predictions, we know little empirically about how well these programs work.

Using satellite data, we find that high-efficiency coal plants increased CO_2 emissions by about 5% relative to low-efficiency plants after the ETS, with no evidence of a differential reduction in emission intensity. Relative to coal plants in Indonesia and the Philippines, which have similar post-COVID trajectories and no carbon market, the increase is concentrated among China's high-efficiency plants, while China's low-efficiency plants move similarly to the foreign plants. Emissions from high-efficiency plants also worsened downwind air quality in historically cleaner areas.

Our within-China estimates identify differential responses by plant efficiency. Whether the program changes total power-sector emissions therefore depends on whether the additional high-efficiency generation was net-new or was existing demand that low-efficiency plants would otherwise have served. Two of our results are more consistent with net addition: the differential does not vary with provincial post-COVID demand growth or with coal self-sufficiency, and it survives a comparison between companies rather than within them; output reshuffling within a firm cannot generate it. Neither is decisive on its own, since both are within-China comparisons with no unregulated benchmark. Only the cross-country comparison provides such a benchmark. China's high-efficiency plants rise well above it while its low-efficiency plants show little difference from the comparison plants. If those countries approximate China's no-ETS counterfactual, and even taking both estimates at the low end of their confidence intervals, the net change across the plants we observe is still an increase of about 14 million tons a year. We acknowledge three limitations that prevent us from providing an estimate of an aggregate effect. First, our sample covers about half of China's coal generation. Second, the comparison group includes only two countries, with a two-year

pre-period, and China’s COVID-19 policies differed from theirs. Third, the triple difference uses only six high-efficiency plants in comparison countries.

Overall, the evidence above is more consistent with net expansion than with reallocation, but it does not identify net expansion or the net effect on aggregate power-sector emissions. China’s power-sector CO_2 rose in absolute terms over 2021-2023 (Appendix Figure A.5), which shows the sector was not contracting but does not distinguish expansion from reallocation. The generous benchmarks encouraged output, and we find no evidence of an offsetting reduction in intensity, which is the mechanism by which a rate-based system fails to curb emissions.

Our findings apply to the rate-based ETS under its current, not-yet-stringent benchmarks. When benchmarks are loose, even inefficient plants comply at low cost, which depresses the carbon price and weakens the incentive to cut intensity. The output expansion among high-efficiency plants is a feature of rate-based design that would not arise under cap-and-trade, where allocations do not depend on output. Two implications follow. Benchmarks have to tighten before the carbon price can do much work, and only an absolute cap removes the reward for producing more. EDF China’s August 2025 [Carbon Pricing Newsletter](#) indicated that China was exploring such a transition.

The findings carry cautionary implications for emerging economies adopting rate-based systems. Rate-based designs may encourage output expansion by efficient producers and suppress carbon prices through generous benchmarks. Both factors work against emissions reduction. Stricter benchmarks can temper these forces, and whether any rate-based system will reduce emissions is an open question.

References


- Almond, D. and S. Zhang (2021). Carbon-trading pilot programs in china and local air quality. *AEA Papers and Proceedings* 111, 391–95.
- Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021, December). Synthetic difference-in-differences. *American Economic Review* 111(12), 4088–4118.
- Bistline, J., N. Mehrotra, and C. Wolfram (2023, May). Economic implications of the climate provisions of the inflation reduction act. Working Paper 31267, National Bureau of Economic Research.
- Chen, X. and B. Lin (2021). Towards carbon neutrality by implementing carbon emissions trading scheme: Policy evaluation in china. *Energy Policy* 157, 112510.
- Couture, H. D., M. Alvara, J. Freeman, A. Davitt, H. Koenig, A. Rouzbeh Kargar, J. O’Connor, I. Söldner-Rembold, A. Ferreira, J. Jeyaratnam, J. Lewis, C. McCormick, T. Nakano, C. Dalisay, C. Lewis, G. Volpato, M. Gray, and G. McCormick (2024). Estimating carbon dioxide emissions from power plant water vapor plumes using satellite imagery and machine learning. *Remote Sensing* 16(7).
- Cui, J., C. Wang, Z. Wang, J. Zhang, and Y. Zheng (2023). Carbon leakage within firm ownership networks. *Available at SSRN 4514971*.
- Cui, J., C. Wang, J. Zhang, and Y. Zheng (2021). The effectiveness of china’s regional carbon market pilots in reducing firm emissions. *Proceedings of the National Academy of Sciences* 118(52).
- Duflo, E., M. Greenstone, R. Pande, and N. Ryan (2013). Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from india. *Quarterly Journal of Economics* 128(4), 1499–1545.
- Duflo, E., M. Greenstone, R. Pande, and N. Ryan (2018). The value of regulatory discretion: Estimates from environmental inspections in india. *Econometrica* 86(6), 2123–2160.
- EDF (2021). 2019-2020 implementation plan for national carbon emission trading total allowances setting and allocation (power generation industry). Technical report, EDF.

- Fischer, C. (2001). Rebating environmental policy revenues: output-based allocations and tradable performance standards. Technical report, Resource for the Future.
- Fischer, C. and R. G. Newell (2008). Environmental and technology policies for climate mitigation. *Journal of Environmental Economics and Management* 55(2), 142–162.
- Friedlingstein, P., M. O’Sullivan, M. W. Jones, R. M. Andrew, J. Hauck, P. Landschützer, C. Le Quéré, H. Li, I. T. Lujckx, A. Olsen, G. P. Peters, W. Peters, J. Pongratz, C. Schwingshackl, S. Sitch, J. G. Canadell, P. Ciais, R. B. Jackson, S. R. Alin, A. Arneeth, V. Arora, N. R. Bates, M. Becker, N. Bellouin, C. F. Berghoff, H. C. Bittig, L. Bopp, P. Cadule, K. Campbell, M. A. Chamberlain, N. Chandra, F. Chevallier, L. P. Chini, T. Colligan, J. Decayeux, L. M. Djeutchouang, X. Dou, C. Duran Rojas, K. Enyo, W. Evans, A. R. Fay, R. A. Feely, D. J. Ford, A. Foster, T. Gasser, M. Gehlen, T. Gkritzalis, G. Grassi, L. Gregor, N. Gruber, O. Gürses, I. Harris, M. Hefner, J. Heinke, G. C. Hurtt, Y. Iida, T. Ilyina, A. R. Jacobson, A. K. Jain, T. Jarníková, A. Jersild, F. Jiang, Z. Jin, E. Kato, R. F. Keeling, K. Klein Goldewijk, J. Knauer, J. I. Korsbakken, X. Lan, S. K. Lauvset, N. Lefèvre, Z. Liu, J. Liu, L. Ma, S. Maksyutov, G. Marland, N. Mayot, P. C. McGuire, N. Metzl, N. M. Monacci, E. J. Morgan, S.-I. Nakaoka, C. Neill, Y. Niwa, T. Nützel, L. Olivier, T. Ono, P. I. Palmer, D. Pierrot, Z. Qin, L. Resplandy, A. Roobaert, T. M. Rosan, C. Rödenbeck, J. Schwinger, T. L. Smallman, S. M. Smith, R. Sospedra-Alfonso, T. Steinhoff, Q. Sun, A. J. Sutton, R. Séférian, S. Takao, H. Tatebe, H. Tian, B. Tilbrook, O. Torres, E. Tourigny, H. Tsujino, F. Tubiello, G. van der Werf, R. Wanninkhof, X. Wang, D. Yang, X. Yang, Z. Yu, W. Yuan, X. Yue, S. Zaehle, N. Zeng, and J. Zeng (2025). Global carbon budget 2024. *Earth System Science Data* 17(3), 965–1039.
- Goulder, L. H., X. Long, J. Lu, and R. D. Morgenstern (2022). China’s unconventional nationwide co2 emissions trading system: Cost-effectiveness and distributional impacts. *Journal of Environmental Economics and Management* 111, 102561.
- Goulder, L. H., X. Long, C. Qu, and D. Zhang (2023). China’s nationwide co2 emissions trading system: A general equilibrium assessment.
- Grainger, C. and T. Ruangmas (2018). Who wins from emissions trading? evidence from california. *Environmental and resource economics* 71, 703–727.

- Greenstone, M. (2024, May). The economics of the global energy challenge. *AEA Papers and Proceedings* 114, 1–30.
- Greenstone, M., G. He, R. Jia, and T. Liu (2022). Can technology solve the principal-agent problem? evidence from china’s war on air pollution. *AER: Insights* 4(1), 54–70.
- Greenstone, M., R. Pande, A. Sudarshan, and N. Ryan (2023). Can pollution markets work in developing countries? experimental evidence from india.
- Hernandez-Cortes, D. and K. C. Meng (2023). Do environmental markets cause environmental injustice? evidence from california’s carbon market. *Journal of Public Economics* 217, 104786.
- Karplus, V., S. Zhang, and D. Almond (2018). Quantifying coal power plant responses to tighter so2 emissions standards in china. *Proceedings of the National Academy of Sciences* 115(27), 7004–7009.
- Shen, S., C. Li, A. van Donkelaar, N. Jacobs, C. Wang, and R. V. Martin (2024, 05). Enhancing global estimation of fine particulate matter concentrations by including geophysical a priori information in deep learning. *ACS ES&T Air* 1(5), 332–345.
- Wang, S. and D. Y. Yang (2025). Policy experimentation in china: The political economy of policy learning. *Journal of Political Economy* 0(0), 000–000.
- Zhang, W., J. Li, G. Li, and S. Guo (2020). Emission reduction effect and carbon market efficiency of carbon emissions trading policy in china. *Energy* 196, 117117.


Figures

Figure 1: Chinese coal plants: raw figures



Notes: Raw (unadjusted) averages for two types of Chinese coal plants, high-efficiency and low-efficiency. Figure 1a plots average annual utilization rates for the two types of plants separately, using data based on satellite imagery. Figure 1b plots average annual electricity generation, using data based on satellite imagery. Figure 1c plots average annual CO2 emissions, using data based on satellite imagery. Figure 1d plots average annual CO2 emissions, using data reported by plants.


Figure 2: Chinese coal plants: first-order outcomes




Notes: Event-study estimates for two types of Chinese coal plants, high-efficiency and low-efficiency. Figure 2a plots estimates of annual utilization rate, using data based on satellite imagery. Figure 2b plots estimates of annual electricity generation, using data based on satellite imagery. Figure 2c plots estimates of CO2 emissions, using data based on satellite imagery. Figure 2d plots estimates of CO2 emissions, using data reported by plants. The bars indicate the 95 percent confidence intervals. Standard errors are clustered at the plant level.

Figure 3: Chinese coal plants: ambient air pollution


(a) SO₂




(b) PM_{2.5}




(c) PM₁₀



(d) NO₂



(e) O₃



Notes: Event-study estimates of pollution concentrations on the trajectories of high- versus low-efficiency plants. The pollution concentration data are from pollution monitors and are matched to the trajectories of the power plants in our sample using HYSPLIT. We examine five pollutants: SO_2 , $PM_{2.5}$, PM_{10} , NO_2 and O_3 . The bars indicate the 95 percent confidence intervals. Standard errors are clustered at the plant level.

Tables

Table 1: First-order effects of the China ETS

	(1)	(2)	(3)
	Utilization rate	Electricity generation	CO2 emissions
Panel A: Satellite data			
$H \times Post$	0.026*** (0.007)	210.36*** (69.60)	186.20*** (62.44)
Pre-ETS mean	0.41	4024	3602
Observations	2,065	2,065	2,065
R ²	0.03	0.35	0.35
Panel B: Reported data			
$H \times Post$			350.02*** (115.23)
Pre-ETS mean			4621
Observations			1,210
R ²			0.17
Plant FE	Y	Y	Y
Year-by-Month FE	Y	Y	Y
Controls	Y	Y	Y

Notes: DID estimates on utilization rate, electricity generation, and CO2 emissions, high-efficiency vs. low-efficiency Chinese plants. Panel A uses data based on satellite imagery. Panel B uses data reported by plants, where only CO2 emissions data are made publicly available. These estimations includes plant fixed effects and year fixed effects. Controls include plant-by-year generation capacity and province-by-year GDP. Standard errors in parentheses are clustered at the plant level. * p<.10, ** p<.05, *** p<.01.

Table 2: Pollution effects of the China ETS (using HYSPLIT)

	(1)	(2)	(3)	(4)	(5)
	ln(SO ₂)	ln(PM2.5)	ln(PM10)	ln(NO ₂)	ln(O ₃)
$H \times Post$	0.044** (0.021)	0.032** (0.015)	0.028* (0.016)	0.000 (0.015)	0.009 (0.013)
Observations	42,290	42,290	42,290	42,290	42,290
R ²	0.78	0.80	0.79	0.81	0.77
Plant FE	Y	Y	Y	Y	Y
Year-by-Month FE	Y	Y	Y	Y	Y
Weather controls	Y	Y	Y	Y	Y
Plant specific time trends	Y	Y	Y	Y	Y

Notes: DID estimates of pollution concentrations on the trajectories of high- versus low-efficiency plants, before and after the ETS. The pollution concentration data are from pollution monitors and are matched to the trajectories of the power plants in our sample using HYSPLIT. Weather controls include temperature, dew point, wind speed and precipitation. Plant-specific time trends are also included. Standard errors in parentheses are clustered at the plant level. * p<.10, ** p<.05, *** p<.01.

A Satellite-based estimation of plant-level generation

The water vapor plumes used to predict generation come mainly from natural draft cooling towers, which release heat during generation, and from wet flue gas desulfurization, which produces plumes saturated with water. Machine learning models are trained on images from NASA’s Landsat-8, ESA’s Sentinel-2A/B, and Planet Labs’ Planet Scope. Planet Scope revisits daily at 3-meter resolution, Sentinel-2A/B every 10 days at 10-60 meters, and Landsat-8 every 16 days at 30-100 meters.

Satellite imagery is matched with plant-level hourly generation data, available only for plants in the US, Europe, and Australia. Each image is labeled with the capacity factor, the plant’s generation that hour divided by its maximum capacity, and an on/off indicator, where on means a capacity factor above 5%.

Plume visibility depends on cloud cover and weather. Couture et al. (2024) therefore keeps only images with less than 20% cloud coverage and applies empirically derived temperature and humidity thresholds, because desulfurization plumes are less visible when it is hot or dry.

Two types of models are trained: gradient-boosted decision trees, which use statistical features from image patches centered on emission structures, and convolutional neural networks, which combine features across patches (Couture et al., 2024). For cross-validation, training plants are rotated into three groups: half for training, a quarter for validation, and a quarter for testing. Out-of-sample plants in Turkey and India provide external validation.

The hourly capacity factor predictions are aggregated into rolling 30-day averages. Multiplying average hourly capacity factor by the plant’s maximum capacity yields daily plant-level electricity generation. To compare with annual data reported by plants, annual electricity generation is calculated by aggregating daily values over the year.

B The ETS benchmark and the capacity classification


The table below reproduces the two conventional coal-fired categories and their benchmarks for the power sector (EDF, 2021).

Category name in the plan	Criterion in Annex 1	Benchmark (tCO_2/MWh)
Conventional coal, “300 MW class and above”	rated power $\geq$ 400 MW	0.877
Conventional coal, “300 MW class and below”	rated power $<$ 400 MW	0.979

Note. The category names refer to generation-unit classes. The 400 MW criterion, set out in Annex 1 of the plan, is what assigns units to each category.


Additional Figures

Figure A.1: Carbon emissions for the top four emitters ([Friedlingstein et al., 2025](#))



Notes: This figure is reproduced from Figure 5(b) in [Friedlingstein et al. \(2025\)](#).


Figure A.2: Daily average carbon price in China




Notes: Data on daily carbon prices are downloaded from China's National Carbon Market Information Network. The dataset provides the highest and lowest daily prices; the plotted series is their average, converted to USD using contemporaneous exchange rates.

Figure A.3: Satellite vs. firm-reported data

(a) Cross-section




(b) Time trends




Notes: About a quarter of the power plants in our data are owned by public companies and reported plant-level electricity generation in 2019-2021. Panel (a) is the cross-sectional correlation and panel (b) the time trend of annual electricity generation, for two data sources: public-company annual reports and satellite images.

Figure A.4: Electricity generation in India: 2019-2021




Notes: Average annual electricity generation of Indian coal plants, 2019-2021, from two data sources: plant reports (India's National Power Portal) and satellite images.

Figure A.5: Selection of cross-country comparison plants



Notes: Annual growth rates of GDP and total electricity consumption (World Bank) and power-sector CO₂ emissions (Ember energy database), for China and the comparison countries, Indonesia and the Philippines, which had no carbon market during our sample period.


Figure A.6: China vs. comparison plants, by efficiency type




Notes: Group-by-year means using satellite data. Panel (a) compares China's 163 high-efficiency plants with the 6 high-efficiency plants in Indonesia and the Philippines; panel (b) compares China's 250 low-efficiency plants with the 70 low-efficiency plants in Indonesia and the Philippines. Satellite data for the comparison countries are provided only through 2022 by Climate Trace. The dashed line marks the 2021 launch of China's national ETS.

Figure A.7: China's high- and low-efficiency plants vs. the pooled comparison plants

(a) China high-efficiency plants vs. comparison plants




(b) China low-efficiency plants vs. comparison plants




Notes: Group-by-year means from satellite data. The control group pools all 88 coal plants in Indonesia and the Philippines; China's high-efficiency group has 163 plants and its low-efficiency group has 250. Satellite data for the comparison countries are provided only through 2022 by Climate Trace. The dashed line marks the 2021 launch of China's national ETS.

Figure A.8: Emission intensities relative to the benchmark

(a) Sub-critical generation units




(b) Super-critical and ultra-supercritical generation units




Notes: Distribution of unit-level emission intensity relative to the ETS benchmark. There are two benchmarks for coal-fired generation units, depending on their rated power. The assigned benchmark is $0.877 \text{ tCO}_2/\text{MWh}$ for units with rated power of at least 400 MW, and $0.979 \text{ tCO}_2/\text{MWh}$ for units below 400 MW (EDF, 2021); see Appendix B. Figure A.8a plots the distribution for generation units using subcritical technology, and figure A.8b for generation units using super-critical or ultra-supercritical technology.

Figure A.9: Power plants and counties on their trajectories (using HYSPLIT)

(a) Power plants




(b) Counties affected by power plants




Notes: Panel (a) maps high-efficiency plants (green triangles) and low-efficiency plants (dark gray dots); panel (b) maps the counties on their trajectories. Using HYSPLIT, we define a county as on the trajectory of a high-efficiency plant if 70% of all trajectories passing by this county during our sample period are from a high-efficiency plant and likewise for counties affected by low-efficiency plants.

Figure A.10: Reported emission intensity by efficiency group over time



Notes: Reported emission intensity is reported CO₂ divided by satellite-derived generation (tCO₂/MWh), for the balanced panel of reporting plants. Panel (a) plots the average intensity by efficiency group. Panel (b) plots event-study coefficients on $\text{High} \times \text{year}$ (2020 omitted) from the specification of Table A.14, Panel A, with plant and year fixed effects, capacity, new-unit, and GDP controls, and 95% confidence intervals; standard errors are clustered at the plant level.

Figure A.11: Reported vs. satellite generation, by efficiency type



Notes: Average annual generation of plants owned by publicly listed companies, 2019-2021, from two sources: generation reported in annual reports (full-day) and satellite-derived generation (morning overpass), shown separately for high- and low-efficiency plants.

Additional Tables

Table A.1: Summary statistics

	(1) Satellite data	(2) Reported data
Total capacity (MW)	1,060 (786)	1,197 (654)
Annual utilization rate	0.41 (0.14)	
Annual electricity generation (GWh)	4,100 (3,713)	
Annual CO2 emissions (Kilotonne)	3,671 (3,327)	4,755 (2,502)
Number of power plants	413	242

Notes: In the plant-by-year balanced panel data for our analysis time period, there are 413 coal-fired power plants in the satellite data and 242 plants in the data reported by the plants. In the data reported by plants, only data on capacity and annual CO2 emissions are publicly available. We show sample means and standard deviations in parentheses.

Table A.2: Summary statistics of generation units

	Uncovered units	Satellite sample	Difference
Units (N)	1,708	1,094	
Mean capacity (MW)	304.0	370.2	66.3***
Super-/ultra-supercritical share	0.23	0.33	0.10***
Mean CO ₂ intensity (tCO ₂ /MWh)	0.997	0.991	−0.006*
Share above ETS benchmark	0.775	0.746	−0.029*

Notes: Column 1 covers the coal-fired generating units in the Global Energy Monitor inventory that are not in the satellite sample; column 2 covers the 1,094 units at the 413 plants in the satellite analysis sample.

Table A.3: Summary statistics of counties on the trajectories of power plants

	Counties affected by high-efficiency plants	Counties affected by low-efficiency plants		
	Mean	Mean	Difference	P-Value
<i>Panel A: 60% threshold</i>				
PM _{2.5}	34.97	42.01	-7.04	0.00
PM ₁₀	59.72	75.28	-15.56	0.00
SO ₂	9.55	13.17	-3.62	0.00
NO ₂	26.93	30.97	-4.03	0.00
O ₃	60.23	61.74	-1.51	0.01
Total Population (million)	279.41	556.98		
Number of Counties	509	1,177		
Number of Counties with Monitors	335	753		
<i>Panel B: 70% threshold</i>				
PM _{2.5}	33.59	39.64	-6.05	0.00
PM ₁₀	59.34	71.87	-12.52	0.00
SO ₂	9.71	13.62	-3.91	0.00
NO ₂	26.08	30.22	-4.14	0.00
O ₃	58.11	59.89	-1.78	0.04
Total Population (million)	131.78	317.30		
Number of Counties	246	692		
Number of Counties with Monitors	158	454		
<i>Panel C: 80% threshold</i>				
PM _{2.5}	34.46	38.31	-3.85	0.04
PM ₁₀	64.49	69.98	-5.49	0.25
SO ₂	10.44	13.46	-3.02	0.00
NO ₂	25.88	30.10	-4.21	0.00
O ₃	58.44	58.20	0.24	0.84
Total Population (million)	46.70	155.13		
Number of Counties	90	345		
Number of Counties with Monitors	64	224		

Notes: Summary statistics of counties on the trajectories of power plants. Using HYSPLIT, we define a county as on the trajectory of a high-efficiency plant if the majority of all trajectories passing by this county during our sample period are from a high-efficiency plant and likewise for counties affected by low-efficiency plants. We report summary stats using three thresholds to define the majority: 60%, 70% and 80%. Pollution concentrations are 2019 annual averages across all ground monitors within 50 km of the county centroid, in micrograms per cubic metre. Counties with no monitor within 50 km are excluded from reporting pollution data, and the last row of each panel reports the number of counties that report pollution data. We show sample means for the two groups of counties separately, with the P-value of their differences.

Table A.4: Heterogeneity test by post-COVID demand growth

	(1)	(2)	(3)
	Utilization	Generation	CO2
$H \times Post$	0.025*** (0.007)	150.925** (62.185)	132.400** (55.705)
$H \times Post \times GDP\ growth$	0.001 (0.001)	2.115 (6.947)	2.240 (6.215)
Observations	2,065	2,065	2,065
R ²	0.91	0.99	0.99

Notes: Demand growth is the percent change in province GDP (National Bureau of Statistics) from the 2020 level to the 2021–2023 mean. All specifications include plant and province-by-year fixed effects and control for generation capacity. Standard errors in parentheses are clustered at the plant level. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.5: Heterogeneity test by coal self-sufficiency

	(1)	(2)	(3)
	Utilization	Generation	CO2
$H \times Post$	0.030*** (0.010)	183.982** (90.107)	160.357** (80.708)
$H \times Post \times Coal\ self-sufficiency$	-0.006 (0.007)	-47.315 (72.877)	-40.254 (65.396)
Observations	2,065	2,065	2,065
R ²	0.91	0.99	0.99

Notes: Self-sufficiency is the province's pre-crisis (2019) ratio of raw coal output to coal consumption (National Bureau of Statistics); a value of one means the province produces exactly what it consumes, above one a province that produces more coal than it burns and below one a province that relies on coal brought in from other provinces and was therefore more exposed to the 2021 supply and price pressure. All specifications include plant and province-by-year fixed effects and control for generation capacity. The estimation sample contains 19 net-importing and 5 net-exporting provinces. Standard errors in parentheses are clustered at the plant level. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.6: Capacity as an outcome

	(1)	(2)
	Capacity	Capacity
$H \times Post$	25.878 (16.113)	11.937 (9.122)
Province $\times$ year FE	No	Yes
Pre-ETS mean	1051.3	1051.3
Observations	2,065	2,065

Notes: DID estimates with plant capacity (MW) as the outcome, comparing high- to low-efficiency plants; column (2) adds province-by-year fixed effects. Standard errors clustered at the plant level. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.7: Robustness: controlling for plant-specific linear time trends

	(1)	(2)	(3)
	Utilization	Generation	CO2
$H \times Post$	0.018** (0.007)	300.014*** (89.010)	266.410*** (79.634)
Control for plant-specific trends	Yes	Yes	Yes
Observations	2,065	2,065	2,065
R ²	1.00	1.00	1.00

Notes: The main emissions DID controlling for plant-specific linear time trends. Standard errors clustered at the plant level. * p<.10, ** p<.05, *** p<.01.

Table A.8: Company-level analysis: high- vs. low-average-efficiency companies

	(1)	(2)	(3)
	Generation	CO2	Reported intensity
$High-ave-eff \times Post$	641.629** (317.907)	568.392** (285.196)	0.103 (0.081)
$High-eff\ share \times Post$	481.961** (214.648)	422.875** (192.654)	0.129 (0.117)
Observations	605	605	165

Notes: Plant-level satellite data aggregated to 121 parent companies, 2019-2023, with company and year fixed effects and controls for company capacity and province GDP. The two rows are from separate regressions: the first defines companies as high-average-efficiency if they have more than 50% of high-efficiency capacity in the baseline; the second uses the continuous share of high-efficiency capacity. Reported intensity measure is only available for the 33 companies that made their carbon reports available throughout. Standard errors clustered at the company level. * p<.10, ** p<.05, *** p<.01.

Table A.9: Robustness check: including province-by-year FE

	(1) Utilization rate	(2) Electricity generation	(3) CO2
Panel A: Satellite data			
$H \times Post$	0.025*** (0.007)	150.12** (62.58)	131.55** (56.06)
Observations	2,065	2,065	2,065
R ²	0.24	0.48	0.48
Panel B: Reported data			
$H \times Post$			259.24** (113.56)
Observations			1,210
R ²			0.35
Plant FE	Y	Y	Y
Year FE	Y	Y	Y
Controls	Y	Y	Y
Province-by-Year FE	Y	Y	Y

Notes: DID estimates on utilization rate, electricity generation, and CO2 emissions, high-efficiency vs. low-efficiency Chinese plants. Panel A uses satellite data; Panel B uses plant-reported data (only CO2 is public). Plant and year fixed effects; controls include plant-by-year capacity and province-by-year GDP. Province-by-year FE are also controlled for. Standard errors in parentheses are clustered at the plant level.
* p<.10, ** p<.05, *** p<.01.

Table A.10: Robustness check: include parent firm-by-year FE

	(1) Utilization rate	(2) Electricity generation	(3) CO2
Panel A: Satellite data			
$H \times Post$	0.015* (0.008)	209.59** (92.62)	187.85** (83.12)
Observations	2,065	2,065	2,065
R ²	0.44	0.53	0.53
Panel B: Reported data			
$H \times Post$			224.51* (125.77)
Observations			1,210
R ²			0.44
Plant FE	Y	Y	Y
Year FE	Y	Y	Y
Controls	Y	Y	Y
Parent firm-by-Year FE	Y	Y	Y

Notes: DID estimates on utilization rate, electricity generation, and CO2 emissions, high-efficiency vs. low-efficiency Chinese plants. Panel A uses satellite data; Panel B uses plant-reported data (only CO2 is public). Plant and year fixed effects; controls include plant-by-year capacity and province-by-year GDP. Parent firm-by-year FE are also controlled for. Standard errors in parentheses are clustered at the plant level. * p<.10, ** p<.05, *** p<.01.

Table A.11: First-order effects of the China ETS: SDID

	(1) Utilization rate	(2) Electricity generation	(3) CO2 emissions
Panel A: Satellite data			
$H \times Post$	0.023*** (0.006)	329.36*** (85.71)	291.54*** (76.73)
Observations	2,065	2,065	2,065
Panel B: Reported data			
$H \times Post$			393.77*** (112.76)
Observations			1,210
Plant FE	Y	Y	Y
Year-by-Month FE	Y	Y	Y
Controls	Y	Y	Y

Notes: SDID estimates (Arkhangelsky et al., 2021) on utilization rate, electricity generation, and CO2 emissions, high-efficiency vs. low-efficiency Chinese plants. Panel A uses data based on satellite imagery. Panel B uses data reported by plants, where only CO2 emissions data are made publicly available. These estimations includes plant fixed effects and year fixed effects. Controls include plant-by-year generation capacity and province-by-year GDP. Standard errors in parentheses are clustered at the plant level. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.12: Main DID by position relative to the ETS benchmark

	(1) Utilization	(2) Generation	(3) CO2
$Below-benchmark \times Post$	0.031*** (0.009)	262.544** (109.712)	232.366** (98.508)
Observations	2,065	2,065	2,065
R ²	0.03	0.35	0.35

Notes: The main DID re-estimated classifying plants by their capacity-weighted CO2 intensity relative to the size-specific ETS benchmark, rather than by combustion technology. Standard errors clustered at the plant level. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.13: Robustness: satellite subsamples with lowest uncertainty

<i>Panel A: Electricity generation</i>				
	(1)	(2)	(3)	(4)
	Full sample	75% lowest uncertainty	50% lowest uncertainty	33% lowest uncertainty
$H \times Post$	210.360*** (69.600)	191.887*** (72.729)	237.917*** (87.157)	325.966*** (110.086)
Observations	2,065	1,545	1,030	685
Plants	413	309	206	137
<i>Panel B: CO2 emissions</i>				
	(1)	(2)	(3)	(4)
	Full sample	75% lowest uncertainty	50% lowest uncertainty	33% lowest uncertainty
$H \times Post$	186.199*** (62.438)	170.094*** (65.246)	210.468*** (78.171)	288.107*** (98.742)
Observations	2,065	1,545	1,030	685
Plants	413	309	206	137

Notes: The main DID re-estimated on high-confidence subsamples at three thresholds, defined by Climate TRACE's plant-level generation-uncertainty measure: plants in the lowest 75%, lowest 50%, and lowest 33% of the plant-level uncertainty distribution. Standard errors clustered at the plant level. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.14: Emission intensity

	Emission intensity	
	(1) DID	(2) SDID
Panel A: Reported measure		
$H \times Post$	0.029 (0.035)	0.026 (0.028)
Observations	615	615
R ²	0.02	
Panel B: Engineering measure		
$H \times Post$	-0.0004 (0.0006)	-0.0006 (0.0005)
Observations	2,005	2,005
R ²	0.01	
Plant FE	Y	Y
Year FE	Y	Y
Controls	Y	Y

Notes: Emission-intensity estimates for high-efficiency vs. low-efficiency Chinese plants. Panel A uses the reported measure (our primary measure) and Panel B the engineering measure. Column (1) is DID; column (2) is SDID ([Arkhangelsky et al., 2021](#)). These estimations includes plant fixed effects and year fixed effects. Controls include plant-by-year generation capacity and province-by-year GDP. Standard errors in parentheses are clustered at the plant level. * p<.10, ** p<.05, *** p<.01.

Table A.15: Ambient pollution at pollution monitors: Robustness checks (with HYSPLIT)

	(1) ln(SO ₂)	(2) ln(PM2.5)	(3) ln(PM10)	(4) ln(NO ₂)	(5) ln(O ₃)
Panel A: stack height is 200 meters					
$H \times Post$	0.041** (0.021)	0.030* (0.016)	0.027* (0.016)	0.001 (0.016)	0.006 (0.013)
Observations	42,289	42,289	42,289	42,289	42,289
R ²	0.78	0.80	0.79	0.81	0.77
Panel B: trajectories run for 120 hours					
$H \times Post$	0.044** (0.021)	0.032** (0.015)	0.028* (0.016)	-0.000 (0.015)	0.009 (0.013)
Observations	42,242	42,242	42,242	42,242	42,242
R ²	0.78	0.80	0.79	0.81	0.77
Panel C: monitors within 5 km from trajectories are matched					
$H \times Post$	0.044** (0.021)	0.033** (0.015)	0.028* (0.016)	0.003 (0.015)	0.008 (0.013)
Observations	42,285	42,285	42,285	42,285	42,285
R ²	0.78	0.80	0.79	0.81	0.77
Plant FE	Y	Y	Y	Y	Y
Year-by-Month FE	Y	Y	Y	Y	Y
Weather controls	Y	Y	Y	Y	Y
Plant specific time trends	Y	Y	Y	Y	Y

Notes: Robustness checks for the pollution DID. Pollution data are from monitors matched to plant trajectories using HYSPLIT. In Panel A, we assume that the average stack height is 200 meters. In panel B, the trajectories run for 120 hours. Panel C uses pollution monitors within 5 km from the trajectories. Weather controls include temperature, dew point, wind speed, and precipitation. Plant-specific time trends are also included. Standard errors in parentheses are clustered at the plant level. * p<.10, ** p<.05, *** p<.01.

Table A.16: Ambient pollution at pollution monitors from plants (without HYSPLIT)

	(1) ln(SO ₂)	(2) ln(PM2.5)	(3) ln(PM10)	(4) ln(NO ₂)	(5) ln(O ₃)
$H \times Post$	0.032 (0.034)	0.017 (0.025)	0.015 (0.025)	-0.013 (0.024)	-0.009 (0.022)
Observations	32,524	32,523	32,524	32,524	32,524
R ²	0.68	0.77	0.75	0.75	0.77
Plant FE	Y	Y	Y	Y	Y
Year-by-Month FE	Y	Y	Y	Y	Y
Weather controls	Y	Y	Y	Y	Y
Plant specific time trends	Y	Y	Y	Y	Y

Notes: DID estimates of pollution concentrations at monitors, without HYSPLIT. We measure monthly PM2.5 from a plant by the weighted average of pollution readings from monitors within a 50-km radius of a plant, using inverse distance weighting (distance from a plant to a monitor). Weather controls include temperature, dew point, wind speed, and precipitation. Plant-specific time trends are also included. Standard errors in parentheses are clustered at the plant level. * p<.10, ** p<.05, *** p<.01.

Table A.17: Satellite measures of PM2.5 (with HYSPLIT)

	ln(PM2.5)	
	(1) 10 km	(2) 30 km
$H \times Post$	0.027*** (0.010)	0.026** (0.010)
Observations	27,011	26,717
R ²	0.68	0.68
Plant FE	Y	Y
Year FE	Y	Y
Weather controls	Y	Y
Plant specific time trends	Y	Y

Notes: DID estimates using monthly satellite-derived PM2.5 as an alternative pollution measure, incorporating HYSPLIT. Weather controls include temperature, dew point, wind speed, and precipitation. Plant-specific time trends are also included. Standard errors in parentheses are clustered at the plant level. * p<.10, ** p<.05, *** p<.01.

Table A.18: Satellite measures of PM2.5 (without HYSPLIT)

	(1)	(2)	(3)
	≤ 5 km	≤ 10 km	≤ 15 km
$H \times Post$	0.026***	0.026***	0.026***
	(0.010)	(0.010)	(0.010)
Observations	27,011	27,011	27,011

Notes: DID estimates using monthly satellite-derived PM2.5, assigning exposure by distance instead of HYSPLIT: grid cells within the stated distance of a high-efficiency plant are compared with cells within the same distance of a low-efficiency plant. Standard errors in parentheses are clustered at the plant level. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.19: Cross-country triple difference: China vs. Indonesia and the Philippines

	(1)	(2)	(3)
	Utilization	Generation	CO2
$H \times Post$	0.001*	5.224**	4.821**
	(0.000)	(2.527)	(2.098)
$H \times Post \times China$	0.026***	297.542***	264.415***
	(0.007)	(89.495)	(80.285)
Observations	1,956	1,956	1,956

Notes: Triple-difference estimates interacting the high- vs. low-efficiency contrast with pre vs. post ETS and with China vs. the comparison countries (Indonesia and the Philippines). Satellite data, balanced 2019-2022. Efficiency is technology-based throughout: super-critical or ultra-supercritical vs. sub-critical in China, and the corresponding combustion technology from the Global Energy Monitor tracker for comparison-country plants. China has 163 high-efficiency and 250 low-efficiency plants. Among the 76 comparison plants whose combustion technology is classified in the Global Energy Monitor, there are 6 high-efficiency and 70 low-efficiency plants. All specifications include plant and country-by-year fixed effects. Standard errors in parentheses are clustered at the plant level. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.20: Split DID: China's high- and low-efficiency plants vs. the pooled comparison plants

	(1)	(2)	(3)
	Utilization	Generation	CO2
<i>Panel A: High-efficiency plants</i>			
<i>China × Post</i>	0.021*** (0.004)	332.801*** (78.124)	297.187*** (70.125)
Observations	1,004	1,004	1,004
<i>Panel B: Low-efficiency plants</i>			
<i>China × Post</i>	-0.006 (0.005)	30.035 (43.572)	27.951 (39.008)
Observations	1,352	1,352	1,352

Notes: Two separate DIDs using satellite-based outcomes, balanced over 2019-2022, with plant and year fixed effects. Panel A treats China's high-efficiency plants; Panel B treats China's low-efficiency plants. In both, the control group pools all 88 coal plants in Indonesia and the Philippines; China has 163 high-efficiency and 250 low-efficiency plants. Standard errors in parentheses are clustered at the plant level. * p<.10, ** p<.05, *** p<.01.