The Pollution Moat: Strategic Emissions under the Clean Air Act

Christine Zhuowei Huang Amit Kumar Mayank Kumar*

October 8, 2026 [Updated Version Here]

Abstract

Clean Air Act nonattainment imposes higher costs on entrants than incumbents, which may induce incumbents to strategically increase pollution to deter entry and strengthen market power. Consistent with this, during the 1997 PM2.5 standards' nonattainment-determination window, ex-ante near-nonattainment-threshold firms increased emissions relative to those slightly away. The increase was stronger at PM2.5-intensive and near-monitor plants, and where incumbents had greater incentives to deter entry. We identify two mechanisms: within-plant cross-media substitution toward air and cross-plant reallocation within multi-plant firms, neither of which raises aggregate emissions. Following nonattainment, plant openings decline, and local emissions correlate positively with firm profits and valuations.

JEL Classification: D22, L13, L51, Q52, Q58

Keywords: Clean Air Act, Air Pollution, Firm Behavior, Entry Barriers, Market Power

*Christine Zhuowei Huang is at Leibniz Institute for Financial Research SAFE; email: huang@safe-frankfurt.de. Amit Kumar is at Singapore Management University; email: amitkumar@smu.edu.sg. Mayank Kumar is at Questrom School of Business, Boston University; email: mkumar2@bu.edu. All errors are our own.

Environmental regulations often follow an area-based monitor-and-intervene approach. Under the Clean Air Act (CAA), for example, the environmental protection agency (EPA) monitors ambient air pollution across geographically delineated areas and designates areas as nonattainment when local pollution exceeds prescribed thresholds, triggering additional pollution-control requirements for firms operating within them. This area-based approach creates a novel externality: a firm's emissions can influence its area's nonattainment designation and, in turn, affect the regulatory costs faced by other firms in the area. Moreover, the CAA includes grandfathering provisions under which nonattainment imposes more stringent and costly pollution-control requirements on new plants than on existing ones. The combination of this asymmetric requirement and the area-based externalities may incentivize incumbent firms to adopt a pollution moat strategy. Under this strategy, incumbents increase emissions to raise the likelihood of nonattainment in their areas and, in turn, increase potential entrants' regulatory costs, strengthening their competitive advantage and market power.[1] This paper presents evidence that firms responded to the introduction of the CAA's fine particulate matter (PM2.5) standards by adopting this strategy and subsequently experienced higher profits and market valuations.

PM2.5 standards were first established under the National Ambient Air Quality Standards (NAAQS) in 1997. Areas where ambient PM2.5 levels violate the standards are subsequently designated as nonattainment and become subject to additional pollution-control requirements. Crucially, existing plants are subject to reasonably available control technology (RACT) requirements, which largely involve retrofitting, whereas new major plants and modifications are subject to the lowest achievable emission rate (LAER) standards, the most stringent and often the costliest emission-control requirements (Becker and Henderson (2000), EPA-commissioned Report

[1] Unlike a product-market strategy, which may target entrants' production, a pollution moat is an input-market strategy that targets entrants' costs. Specifically, incumbents raise local pollution-control costs, which influence entrants' location choices and deter entry into incumbents' counties. Entrants either incur higher production costs by locating in an incumbent's county or choose another location with higher expected net benefits, which depend on location-specific factors such as pollution-control costs, labor, and resource availability. We show later that incumbents increased emissions more in counties with greater future growth opportunities in PM2.5-related industries, consistent with incumbent's location-based incentive to deter entry.

(1980), and 42 U.S.C. §7501(3)). The asymmetric costs are central to the pollution moat strategy.

Two features of nonattainment designations facilitate this strategy. First, under the 1997 PM2.5 standards, nonattainment imposes control requirements only on plants directly emitting PM2.5 or its two precursors, SO_2 and NO_X (EPA (2007)). Thus, the PM2.5-based strategy remains viable even when other NAAQS pollutants, such as ozone, are also being monitored simultaneously.[2] Second, because PM2.5 nonattainment triggers new pollution-control requirements for 5 to 10 years from the date of designation regardless of prior attainment status, incumbents retain incentives to induce subsequent PM2.5 nonattainment even in areas already in nonattainment under other standards.

The timing of PM2.5 monitoring and nonattainment designations made the strategy feasible. Most PM2.5 monitors were deployed in 1999 and 2000, making monitoring locations and ambient PM2.5 levels publicly observable. Although nonattainment designations were scheduled for 2005, they were based on measurements from 2001 to 2003. Consequently, firms could assess their counties' proximity to the regulatory thresholds using monitor data available from 1999 onward and anticipate that their emissions during 2001–2003 would contribute to the PM2.5 levels determining eventual nonattainment designations.

We implement a difference-in-differences (DD) strategy exploiting these features. Because additional pollution is more likely to trigger nonattainment in areas with ex-ante pollution levels near the threshold, we define treatment based on proximity of a county's average PM2.5 concentrations during 1999–2000 (before the statutory monitoring period) to the 15 $\mu g/m^3$ threshold. Counties within $\pm 1 \mu g/m^3$ of the threshold are treated, and those in the 10–13 and 17–20 $\mu g/m^3$ ranges are controls. We

[2] For example, nonattainment under the 1997 8-hour ozone standards triggers requirements targeting volatile organic compounds (VOCs) and NO_X, which do not fully overlap with those targeted under the PM2.5 standard. Moreover, the CAA ozone standards, unlike PM2.5 standards, tier the nonattainment areas by severity and impose fees on both existing and new plants in high-severity zones after attainment deadlines, with fees proportional to the quantity of pollutants emitted (EPA (2010)). Targeting both old and new plants with these fees substantially undermines the existing plants' incentives to strategically trigger ozone nonattainment as a barrier to entry for new plants.

draw treated units from both sides of the threshold to reflect firms' ex ante uncertainty over whether their emissions would push the subsequent three-year county average past the threshold, an uncertainty that even the firms in counties ex-ante just above the threshold face. Control units are also drawn from both sides of the threshold because a simple above-versus-below comparison may confound the treatment effect with the effects of any local characteristics correlated with PM2.5 levels. Finally, we use two alternative treatment definitions for all estimations exploiting the threshold: a wider treatment window of $\pm 2 \mu g/m^3$ and a continuous threshold proximity measure.

We estimate a two-way fixed effects (TWFE) DD specification using plant-by-chemical emissions data from the Toxics Release Inventory (TRI). Plants inherit the treatment or control status of the county in which they operate. We find that treated facilities increased air emissions by approximately 8 percent during the 2001–2003 monitoring period relative to controls. The result is robust to both alternative specifications. Incorporating plant, chemical, and year fixed effects in these specifications restricts identification to within-plant and within-chemical variation and accounts for temporal shifts in chemical inputs across production processes and common annual shocks.

The emission increase was larger for chemicals whose emissions are strongly associated with PM2.5 emissions (17.8 percent) and for plants in 4-digit NAICS industries in the top decile of the PM2.5-emissions-share distribution (18.8 percent). By contrast, emissions did not increase for low-relevance chemicals or plants in low-PM2.5-share industries.

A placebo test mitigates the concern that the emission increases during the monitoring period might reflect some contemporaneous industry trends or unobserved shocks in 2001. We compare the emissions at plants in counties without PM2.5 monitors against the same control group and find no significant increase in emissions.

Three further findings indicate that the emissions increase was consistent with a strategic motive to raise the probability of nonattainment. First, plants within five miles of PM2.5 monitors increased emissions significantly, whereas those farther away did

not. This pattern is consistent with a response aimed at affecting measured PM2.5 concentrations, as near-monitor plants are more likely to influence monitor readings. Second, plants in markets (defined as county×3-digit NAICS) with above-median ex-ante concentration in both labor and emissions increased emissions, whereas those in markets with below-median concentration did not. This pattern is consistent with incumbents' greater ability to internalize anticipated post-nonattainment cost increases, coordinate emission responses, and stronger incentives to protect rents associated with concentrated markets. Third, plants in counties with above-median PM2.5-industry growth opportunities increased emissions relative to controls, whereas those in below-median counties did not.[3] This differential pattern is consistent with entry deterrence being more valuable to incumbents in areas with stronger growth prospects. The latter two findings also suggest that firms' responses were more pronounced where the benefits of deterring entry were greater.

We identify two mechanisms underlying the emissions increase. First, plants shifted the proportions of emissions across media. Total quantities of chemicals released or managed did not increase, but the share released directly into the air rose, while the shares managed through off-site transfers and on-site waste-management activities declined. Second, multi-plant firms reallocated emissions across plants toward treated counties. Emissions at treated standalone establishments did not increase relative to controls, whereas emissions in treated plants of multi-plant firms that own plants in both treated and control counties increased by approximately 13 percent relative to controls. At the same time, total emissions at the parent-firm level for these multi-plant firms did not increase relative to those of multi-plant firms that operated at least one plant in a control county and none in a treated county. These patterns together suggest that multi-plant firms reallocated emissions toward treated counties, where additional emissions were more likely to affect nonattainment designations, without increasing aggregate emissions.

[3] We measure county-level growth opportunities by mapping industry-level market-to-book ratios to counties, following a procedure similar to Di Maggio et al. (2022). This measure captures growth opportunities as perceived by financial markets and accounts for local technology and resource endowments. See Section A.5.2 for details on its construction.

The pollution moat strategy appears to benefit incumbents following nonattainment designation. First, counties that become nonattainment have 16 to 18 percent fewer new plant openings over 5- or 10-year horizons, consistent with nonattainment creating a persistent barrier to local entry. Second, Compustat data suggest that incumbent firms benefit from these barriers. Among firms with plants in counties that subsequently become nonattainment, greater contemporaneous emissions aggregated across those counties are associated with higher profits and market-to-book ratios, controlling for aggregate emissions across their plants in attainment counties. These associations persist over five- and ten-year horizons, indicating that nonattainment designations generate persistent financial benefits for incumbent firms.

Our findings are inconsistent with an explanation based on production front-loading, in which firms shift future production into the monitoring period in anticipation of higher pollution-control costs. First, neither plant-level nor parent-firm-level aggregate emissions increased, contrary to what production front-loading would predict. Second, because production front-loading incentives depend on the county's expected nonattainment probability, they should be similar within a county across standalone and multi-plant firms and across plants near and far from monitors. Contrary to this, emission increases were greater along both these dimensions. Our results also do not support an alternate strategy where incumbents increase emissions to disadvantage their less-efficient local peers. Because incumbents in concentrated markets face fewer local competitors, this alternative predicts weaker emission responses in ex-ante concentrated markets, which runs counter to our findings discussed previously.

Contribution and literature: Overall, we present evidence consistent with firms adopting a pollution moat strategy, exploiting regulatory externalities through their emission decisions to gain market power. Our findings highlight a novel view of pollution as a strategic firm choice that can shape the competitive environment. More broadly, they suggest that environmental policies that mitigate incentives for such behavior may achieve pollution control while limiting industry concentration.

The frequency and direction of strategic emission adjustments depend on firms' underlying motives. Zou (2021) documents that firms reduce emissions on days when nearby pollution monitors are active and increase them on non-monitored days, presumably to conceal their true pollution levels. In contrast, incumbent firms in our setting increase emissions over a longer horizon to raise the likelihood of nonattainment designation, which depends on whether three-year average pollution exceeds regulatory thresholds, thereby deterring future entry.

We contribute most directly to the well-established literature studying the effects of CAA nonattainment on local manufacturing and economic activities (Henderson (1996), Becker and Henderson (2000), Greenstone (2002), Greenstone (2003), List et al. (2003), Becker (2005), Greenstone, List, and Syverson (2012)).[4] The novelty of our paper lies in examining firms' emissions decisions before nonattainment designation, specifically during the statutory monitoring period when emissions can themselves influence the designation. We find that firms increase emissions during this period, consistent with an effort to raise the probability of nonattainment. We also offer a novel interpretation of the post-nonattainment declines in plant openings and employment, which the literature largely attributes to higher regulatory costs (Becker and Henderson (2000), Greenstone (2002)). Our findings of increased emissions before nonattainment and higher profits and market valuations afterward suggest that at least some of these declines arise from incumbent firms' efforts to deter entry and strengthen their market power.

We also add to the corporate finance literature that examines the determinants of firm emissions, including environmental pressure (Duchin, Gao, and Xu (2025)), financial constraints (Xu and Kim (2022)), local political preferences (Bisetti et al. (2026)), parent-firm liability protections (Akey and Appel (2021)), and investor type (Bellon (2025), Naaraayanan, Sachdeva, and Sharma (2026)). We highlight a novel strategic motive for firms' emissions: firms increase emissions to raise entry barriers

[4] Related studies examine nonattainment effects on labor reallocation (Walker (2011)), housing markets (Chay and Greenstone (2005)), and human capital (Isen, Rossin-Slater, and Walker (2017)). Aldy et al. (2022) survey the literature.

and subsequently realize higher profits and market-to-book ratios through continued production at existing plants in nonattainment counties.

More broadly, our findings relate to the public choice theory of regulation, which views stricter regulations as entry barriers that may lead to market power (see, e.g., Djankov et al. (2002, p.3), Bertrand and Kramarz (2002) and Klapper, Laeven, and Rajan (2006)). We also highlight that environmental regulations can generate market power for incumbents. This effect, however, does not stem from stricter regulation per se, but rather from incumbents strategically exploiting regulatory externalities to subject potential entrants to heightened compliance standards. Finally, our mechanism differs from classic regulatory capture, in which firms actively influence policy design to advance their interests (Stigler (1971), p. 3). Instead, our mechanism operates through firms strategically exploiting a pre-existing regulatory structure.

I. Institutional Background

A. The Clean Air Act and Nonattainment Requirements

Air pollution is federally regulated under the Clean Air Act (CAA), which authorizes the Environmental Protection Agency (EPA) to establish NAAQS for the major air pollutants. PM2.5 was added to the NAAQS in 1997. These standards specify maximum allowable concentrations of pollutants in ambient air based on public health and welfare considerations. Areas that meet the standards are designated as attainment areas, while those that fail to meet them are designated as nonattainment areas for the respective pollutants.

To reduce pollution, nonattainment designations impose stringent requirements in the affected areas governing important aspects of manufacturing processes and technological choices including equipment design and choice of machinery, solvents, fuels, and scrubbers. States are required to submit State Implementation Plans (SIPs) to the EPA specifying the control strategies they will use to bring the area into compliance with the relevant standards. The EPA reviews SIPs to assess their scientific and technical sufficiency. If a SIP is deemed deficient, the EPA may require the state to

revise it. Once approved, SIP provisions become federally enforceable. If necessary, the EPA itself may develop and enforce a Federal Implementation Plan (FIP).

The regulatory requirements under nonattainment are substantially different for existing facilities and new sources of pollution. Existing facilities may be required to adopt RACT, which considers the technological and economic feasibility of reducing emissions. In contrast, new major stationary sources and major modifications to existing sources are subject to Nonattainment New Source Review (NNSR), which requires compliance with the LAER (EPA (n.d.)). Unlike RACT, LAER does not consider the economic feasibility of emission reductions and is considered the most stringent and costliest pollution-control technology (Becker and Henderson (2000), McAllister (2007), EPA-commissioned Report (1980), and 42 U.S.C. §7503(a)(2)). New sources may also need to obtain emission offsets to compensate for the additional pollution they generate. These requirements can substantially increase the costs of establishing new plants in nonattainment areas.

The nonattainment SIP requirements are specific to both the underlying standards, which the EPA revises over time, and the associated pollutants. Thus, nonattainment under the 1997 PM2.5 or the 1997 8-hour ozone standards triggers their separate specific requirements, and plants in affected areas must comply with applicable requirements associated with each nonattainment designation. For example, the requirements under the 1997 PM2.5 standards target any plants that directly emit PM2.5 or its two precursors, SO_2 and NO_X (EPA (2007)), whereas the ozone standards target VOCs and NO_X.

Nonattainment requirements are also specific to each designation episode. For example, a nonattainment designation announced in 2005 for violating the 1997 PM2.5 standards triggers its own SIP requirements and a new deadline of five to ten years to bring the area into attainment (EPA (2007)), even if the area was already designated nonattainment under another standard.

A.1. 1997 PM2.5 Standards and Nonattainment Designation

The EPA established the PM2.5 standards in 1997, including an annual standard of 15 $\mu g/m^3$, based on the three-year average of annual mean PM2.5 concentrations, and a 24-hour standard of 65 $\mu g/m^3$, based on the three-year average of the annual 98th percentile of daily PM2.5 concentrations. An area was designated as nonattainment if it violated either standard. In practice, however, the annual standard, rather than the 24-hour standard, was the binding constraint: every nonattainment designation involved a violation of the annual standard, and only two involved violations of both (EPA (n.d.)). We therefore design our empirical strategy using the annual standard, an approach also used in Bishop, Ketcham, and Kuminoff (2023).

Because PM2.5 was a new pollutant added to the NAAQS, new monitors were needed to be deployed nationwide. This was completed during 1999 and 2000. Although initially the EPA intended to use data collected from 2000 to 2002 to determine nonattainment status, the final designations were based on data from 2001 to 2003 (Bishop, Ketcham, and Kuminoff (2023)). This shift in the monitoring period was caused by both delays in deploying the monitoring network and litigation over the EPA's authority to establish the 1997 NAAQS. The Supreme Court ultimately upheld that authority in February 2001 (CRS (2010), footnote 149).[5] The final monitoring period motivates our design choice of using 2001 to 2003 as the "Post" period in our DD design (see Section II).

II. Empirical Research Design

We employ a DD design that compares plant-chemical-level air emissions between counties ex ante close to the PM2.5 annual threshold and those slightly farther from it during the statutory pre-designation monitoring period. We classify counties as treated or control based on the distance of their average ambient PM2.5 concentrations

[5] The legality of EPA's authority to set NAAQS and the exclusion of implementation costs from the NAAQS standard-setting process were challenged in *American Trucking Associations, Inc. v. EPA* (175 F.3d 1027 (D.C. Cir. 1999)). On May 14, 1999, the D.C. Circuit ruled against EPA, creating uncertainty about the enforceability of the 1997 NAAQS. Firms may therefore have delayed strategic emission responses until the legal uncertainty was resolved. On February 27, 2001, the Supreme Court reversed the D.C. Circuit's ruling and upheld the EPA's authority to set and enforce NAAQS (121 S. Ct. 903).

during 1999–2000 from the annual PM2.5 standard of 15 $\mu g/m^3$. We obtain annual mean PM2.5 concentrations from the EPA's Air Quality System (AQS).[6] The AQS contains PM2.5 data for 951 sites across 584 counties in 1999 and 1,108 sites across 687 counties in 2000, consistent with the phased rollout of PM2.5 monitors discussed in Section I. Averaging the PM2.5 data over the two years for treatment assignment mitigates the influence of transitory fluctuations in pollution concentrations, while accommodating the phased deployment of PM2.5 monitors.

We consider two treatment definitions based on the distance between each county's average PM2.5 concentration over 1999–2000 and the 15 $\mu g/m^3$ threshold. The first defines treated counties as those with PM2.5 concentrations within ± 1 $\mu g/m^3$ of the threshold (14–16 $\mu g/m^3$), while the second uses a wider bandwidth of ± 2 $\mu g/m^3$ (13–17 $\mu g/m^3$). Importantly, we draw treated counties from both sides of the regulatory threshold, because ex-ante average concentrations above the threshold do not necessarily result in subsequent nonattainment designations, which are based on ambient PM2.5 concentrations measured during 2001–2003. Firms in counties just above the threshold still have an incentive to increase emissions and raise the probability of subsequent nonattainment.

The control group consists of counties with 1999–2000 average PM2.5 concentrations between 10 and 13 $\mu g/m^3$ or between 17 and 20 $\mu g/m^3$ for both specifications. We draw control counties from both sides of the regulatory threshold to mitigate confounding the treatment effect from the unobserved county characteristics correlated with local PM2.5 levels. For example, any county characteristic that systematically increases with PM2.5 levels will confound the threshold (treatment) effect that is estimated from comparing counties just above the threshold with those just below.

As an alternative to the binary treatment, we construct a continuous measure of proximity to the 15 $\mu g/m^3$ threshold, defined as $5 - |\text{PM2.5} - 15|$, using the same sample of counties as in the binary treatment design: those with 1999–2000 average PM2.5

[6] We restrict the AQS sample to "PM2.5 - Local Conditions" (parameter code 88101) and observations corresponding to the 1997 annual standard. For each county-year, we retain the highest reported annual mean concentration across monitoring observations and then average these values across 1999 and 2000. For counties with data available in only one of the two years, we use the available year's concentration.

concentrations between 10 and 20 $\mu g/m^3$. Higher values indicate that a county's PM2.5 concentration is closer to the regulatory threshold and, consequently, that additional emissions by local incumbent firms have a greater marginal effect on raising the county's probability of nonattainment.

We define 1997–2000 as the pre-treatment period and 2001–2003 as the post-treatment period in our DD design. The latter corresponds to the EPA's monitoring period for determining nonattainment status and aligns with the timing of monitor deployment and a landmark Supreme Court case, both discussed in detail in Section I.A.1.

A. Regression Specification

We estimate the following specification for establishment-chemical-level analysis:

$$y_{ikct} = \alpha_0 + \beta_1 \text{Treat}_{ikc} \times \text{Post}_t + \text{Controls} + \mu_i + \delta_k + \gamma_t + \epsilon_{ikct}, \quad (1)$$

where y_{ikct} is the outcome for chemical k at plant i in county c in year t. $Treat_{ikc}$ equals one if plant i is located in county c, whose average PM2.5 concentration in 1999–2000 falls within ±1 (primary or narrow band) or ±2 (wider band) of the 15 $\mu g/m^3$ threshold, and zero if it falls within either the 10–13 or 17–20 $\mu g/m^3$ range. $Post_t$ is one for years 2001–2003 and zero for years 1997–2000. β is the coefficient of interest. It captures the average change in chemical release before and after the event by plants in the treated counties relative to the control. *Controls* are the natural logarithms of county-level population and per capita income. μ_i, δ_k, and γ_t denote establishment, chemical, and year fixed effects, respectively.

We also use the following continuous-treatment DD specification:

$$y_{ikct} = \alpha_0 + \beta_1 \text{Threshold Proximity}_{ikc} \times \text{Post}_t + \text{Controls} + \mu_i + \delta_k + \gamma_t + \epsilon_{ikct}. \quad (2)$$

Here $Threshold\ Proximity_{ikc}$ is defined as $5 - |\text{PM2.5} - 15|$, where PM2.5 denotes county c's average PM2.5 concentration during 1999–2000. Since the counties included in the sample have ex-ante average PM2.5 concentrations between 10 and 20 $\mu g/m^3$, this measure ranges from 0 to 5, with higher values indicating that the county's pre-treatment PM2.5 concentration is closer to the 1997 annual threshold. The coefficient β_1

captures the difference in the post- versus pre-period outcome associated with a one-unit increase in ex-ante proximity of the county's average PM2.5 levels over 1999–2000 to the threshold of 15 $\mu g/m^3$.

III. Data and Summary Statistics

Our primary data come from the TRI dataset, which reports the quantities of individual toxic chemicals released annually into air, surface water, and land by plants that meet its reporting requirements. The database also identifies plant locations and parent companies, allowing us to examine emissions across plants owned by publicly listed firms. We exclude plants in the utilities sector (NAICS 22), including coal-fired and other power plants. We also exclude federally owned plants and government-owned, contractor-operated plants. Our TRI panel runs from 1997 to 2003. In addition, we obtain information on plants' waste management practices from TRI. The data provide information on the quantities of each chemical that plants recycle, use for energy recovery, treat, or transfer off-site for further management or disposal.

Our monitor-level ambient PM2.5 concentration data come from the EPA's Air Quality System (AQS). We use the coordinates of AQS monitors to calculate their distances from TRI plants. We also obtain PM2.5 emissions data (in tons) from the 2008 National Emissions Inventory (NEI, version 3) to examine the associations between the quantities of PM2.5 emissions and TRI chemical releases from plants. NEI data are compiled every three years and provide the best estimates of emissions based on most updated tools, methods, and data, including data from facility-level continuous emissions monitoring systems.

The data on 2005 PM2.5 nonattainment designations under the 1997 NAAQS come from the EPA's Green Book. Firm-level financial and accounting information comes from Compustat, and stock price data are from CRSP. We obtain Compustat-TRI links from Hsu, Li, and Tsou (2023). We draw county-level socioeconomic and industry characteristics from public sources, including population estimates from the U.S. Census Bureau, establishment counts and employment from County Business Patterns

(CBP) and the Quarterly Census of Employment and Wages (QCEW), and per capita income from the Bureau of Economic Analysis (BEA).

Table I reports summary statistics for the key variables. The statistics are shown separately for treated and control groups (columns (1) to (3) and (4) to (6), respectively) over the 1997–2003 sample period. Treated and control establishments are broadly comparable in observed characteristics. On average, annual toxic chemical air emissions are 11,160 pounds in treated counties and 10,120 pounds in control counties. Air emissions account for approximately 29% of total releases, while water and land releases account for less than 4% in both groups. Off-site transfers and on-site waste management activities together account for approximately 67% in both groups. Treated and control counties also have similar population and per capita income, which we include as control variables in our regression specifications.

[Insert Table I Here]

IV. Results

We assess our hypothesized pollution moat strategy in three stages. First, we examine firms' emissions responses during the statutory monitoring period. Second, we investigate the underlying mechanisms. Third, we assess whether the strategy enables incumbent firms to gain market power following PM2.5 nonattainment.

A. Strategic Emission Responses to the 1997 PM2.5 Standards

Under the hypothesized pollution moat strategy, incumbent firms have incentives to raise the nonattainment probability under PM2.5 standards so that future entrants in PM2.5-relevant industries face stricter and costlier LAER standards, if they locate in incumbent's county. Because additional emissions are more likely to influence nonattainment probability in counties that have ex-ante PM2.5 levels closer to the emission threshold, plants in these counties should increase their emissions more than plants in counties with ex-ante PM2.5 levels slightly farther from the threshold. We therefore compare emissions at plants in treated counties with ex-ante PM2.5 levels

within $\pm 1 \mu g/m^3$ of the threshold with emissions at plants in control counties, which have ex-ante PM2.5 levels of 10–13 and 17–20 $\mu g/m^3$.

[Insert Table II here]

We estimate equation (1) using the plant-by-chemical panel from TRI data and report the results in Table II. Column (1) shows that plants in treated counties increased emissions by 8% relative to plants in control counties farther from the threshold. Widening the treatment band to 2 μg/m^3 on either side of the threshold also yields a statistically significant increase of approximately 6%, as shown in column (2). Column (3) confirms this pattern using a continuous distance measure. A one-unit increase in threshold proximity increases emissions by 2.2%. Unlike the binary treatment specification, the continuous specification captures more closely how the marginal effect of additional emissions on the probability of nonattainment increases as counties' ex-ante PM2.5 levels approach the threshold. This also assures that the findings are robust to alternative treatment band definitions.

[Insert Figure 4 here]

Figure 4, Panel A, shows the event study plot for the emission outcome in column (1). Relative to the 1999 gap in emissions between treated and control plants, the estimated differences during the pre-treatment period are small and statistically insignificant. The difference begins to rise in 2000 and increases further during 2001–2003, with the largest effects observed in 2002 and 2003. The absence of discernible pre-trends across the treated and control plants alleviate concerns about violation of parallel-trends assumption.

A.1. Emission Responses for PM2.5-Relevant Chemicals

Under our hypothesis, firms strategically increase emissions to increase nonattainment probability. If so, the emission increases should be more pronounced for chemicals whose releases are more closely linked to PM2.5 emissions. To test this, we first examine whether the emissions response varies with the relevance of individual TRI chemicals to PM2.5 emission. To measure how relevant a TRI chemical is for PM2.5 emissions, we estimate separate cross-sectional regressions of county-aggregated total

quantity of primary PM2.5 emissions reported in the 2008 NEI data on county-aggregated emissions of each TRI chemical k in the same year. The specification is as follows:

$$\text{PM2.5 Quantity}_c = \alpha_k + \beta_k \, \text{Emission}_{kc} + \delta_{s(c)} + \epsilon_{kc}. \tag{3}$$

Here, c indexes counties, k indexes TRI chemicals, and δ denotes state fixed effects. Standard errors are clustered at the state level. A chemical k has very high relevance if β_k is positive and significant at the 1% level, high relevance if positive and significant at the 5% level, and low relevance if statistically insignificant at the 10% level.[7]

[Insert Table III here]

We estimate the emissions response separately for each of the three PM2.5 relevance groups using the specification in equation (1). Table III presents the results. Treated plants increased emissions of very-high-relevance chemicals by 17.8% ($= e^{0.16} - 1$) and high-relevance chemicals by 11.4%, as reported in columns (1) and (4), respectively. Both estimates are statistically significant at the 1% level. By contrast, emissions of low-relevance chemicals increased by only 6.1%, significant at the 10% level (column (7)). Using the wider treatment band yields increases of 14.1% ($= e^{0.13} - 1$) and 9.9% for very-high- and high-relevance chemicals, respectively (columns (2) and (5)), but no statistically significant increase for low-relevance chemicals (column (8)). The continuous DD specification also reveals a similar pattern. A 1 μg/m^3 decrease in distance to the threshold is associated with a 4.7% increase in emissions of very-high-relevance chemicals (column (3), significant at the 1% level) and a 3.3% increase for high-relevance chemicals (column (6), significant at the 5% level), while the estimated response for low-relevance chemicals is statistically insignificant (column (9)).

The event study plot in Panel B of Figure 4, corresponding to column (1) of Table III, shows that the emission difference between treated and control plants was close to

[7] Because our aim is to identify time-invariant associations between TRI chemicals and PM2.5 emissions, the outcome variable is county-level primary PM2.5 emissions from the 2008 NEI, rather than ambient PM2.5 concentrations from AQS. We use the 2008 NEI data for its updated emissions estimation methods and improvements in source-specific PM2.5 estimates relative to earlier inventories, such as NEI 2005. Zou (2021) uses the 2011 NEI data for a similar purpose in the context of PM10 in a 2001–2013 sample.

zero during the pre-treatment period but increased significantly after treatment. The plot also reveals no discernible divergence in pre-trends.

Overall, these heterogeneous emission responses across TRI chemicals with varying PM2.5 relevance support the interpretation that incumbent firms strategically targeted emission increases toward chemicals more relevant to PM2.5 emissions.

A.2. Emission Responses by PM2.5-Intensive Plants

Following the preceding argument, the emission increases should be more pronounced among plants in industries that contribute more to PM2.5 emissions. To examine this prediction, we rank 4-digit NAICS industries operating in the sample counties by their shares of aggregate PM2.5 emissions across these counties, as measured in the 2008 NEI data. We classify industries in the top decile as PM2.5-intensive and the remaining industries as low-intensity. We then estimate the baseline specification in equation (1) separately for plants in these two industry groups.

[Insert Table IV here]

Table IV presents the results. Plants in PM2.5-intensive industries in treated counties increased emissions by approximately 18.8% ($= e^{0.17} - 1$) relative to plants in the same industries in control counties (column (1)). By contrast, plants in less-intensive industries exhibited no statistically significant emission response (column (4)). Widening the treatment band to ±2 around the regulatory threshold yields a similar pattern. Emissions increased by 15.1% ($= e^{0.14} - 1$) among plants in PM2.5-intensive industries, with no significant increase among plants in other industries. The continuous DD specification also supports this pattern. A 1 μg/m^3 decrease in distance to the threshold is associated with a 5.4% increase in emissions among plants in PM2.5-intensive industries (column (3)), but has no statistically significant association with emissions among plants in less-intensive industries (column (6)).

Figure 4, Panel C, shows the event study plot corresponding to column (1) of Table IV. It suggests that the treated and control units emitted similar quantities of TRI chemicals during pre-treatment period, but treated plants increased their emissions significantly in the post-treatment period. Moreover, it does not indicate

any discernible divergence in pre-trends. Overall, consistent with our hypothesized pollution moat strategy, incumbent plants in PM2.5-intensive industries increased emissions more to induce subsequent PM2.5 nonattainment.

A.3. A Placebo Test: Emission Response in Counties without PM2.5 Monitors

If firms increased emissions during the 2001–2003 statutory monitoring period to raise the probability of PM2.5 nonattainment, plants in counties without PM2.5 monitors should exhibit no such response, as their emissions cannot affect monitored nonattainment outcomes. We therefore conduct a placebo test by redefining treated counties as those without monitors while retaining the same control group. Moreover, because only 627 counties had PM2.5 monitors by 2000, unmonitored counties vary widely in their local industrial activity. We therefore construct an alternative placebo sample of unmonitored counties with substantial local manufacturing presence in PM2.5-related industries. Specifically, the sample comprises unmonitored PM2.5 hotspot counties, defined as those in the top 5% of the employment-share distribution in at least one of the 12 PM2.5-related industries.[8] Table A.1 in Online Appendix presents the results of the placebo tests using specification (1). We find no significant emission increase when we pool all TRI chemicals, both in the sample of all unmonitored counties and in the subset of hotspot unmonitored counties (columns (1) and (2)). The results remain statistically insignificant when restricting the sample to chemicals whose PM2.5 relevance is very high (columns (3) and (4)) or high (columns (5) and (6)). These null results in unmonitored counties indicate that higher emissions in monitored counties are unlikely to be driven by some common trend or unobserved shocks.

A.4. Emission Response Heterogeneity by Monitor Proximity

Proximity to PM2.5 monitors generates variation in strategic incentives even among plants within the same county, which share the same ex-ante PM2.5 level and

[8] We begin with identifying three-digit NAICS industries that each account for at least 0.5% of national PM2.5 emissions in the 2008 NEI data. The NAICS codes are: 211, 212, 221, 311, 321, 322, 324, 325, 327, 331, 486, and 928. We then calculate each county's employment share in each of these industries using average employment over 1999–2000, rank counties by these shares, and identify those in the top 5% for each industry. Zou (2021) uses a similar procedure to identify PM10 hotspot counties.

nonattainment probability. Because emissions dissipate with distance, additional emissions have a larger marginal effect on nonattainment probability when released by plants located closer to monitors. We test this prediction by splitting plants into those within five miles of a PM2.5 monitor and those farther away. Table V shows that treated plants closer to monitors increased emissions by approximately 10.5% relative to control plants in the same proximity group (column (1)), while the effect is small and statistically insignificant for plants farther away (column (2)). Columns (3) and (4) reveal a similar pattern under the wider treatment band. These findings support our strategic-incentive mechanism and help rule out county-level confounding factors, such as changes in local labor supply following the dot-com bubble, which are unlikely to vary systematically over such short distances.

[Insert Table V here]

A.5. Emission Response Heterogeneity by Entry Deterrence Benefits

If incumbent firms increase emissions to deter entry, their response should be stronger when the benefits of entry deterrence are greater. Because these benefits are difficult to measure directly, as future entrants' location choices and location-specific expected net benefits are unobservable, we use two proxies: ex-ante local industry concentration and counties' future growth opportunities in PM2.5-related industries. We explain the rationale for these proxies and examine heterogeneity in emission responses along these two dimensions below.

A.5.1 Ex-ante Market Concentration: Ceteris paribus, market concentration shapes incumbents' net benefits from deterring entry. Incumbents in more concentrated markets can better internalize the expected increase in pollution-control costs associated with RACT requirements following nonattainment. They also have greater concentration-related rents to protect and face lower costs of coordinating emission responses, which mitigates free-riding. Therefore, the emission response should be stronger among firms in more concentrated markets. We test this prediction next.

We measure concentration in both labor and emissions markets in the pre-event year 2000, defining each market as a county×3-digit NAICS cell. Because plant-level employment is not publicly observable, we construct a labor market Herfindahl-Hirschman index (HHI) using CBP data, following Meiselbach et al. (2022).[9] For air emissions market concentration, we calculate the HHI using each plant's share of total air emissions (across all TRI chemicals) in its market. We classify markets as highly concentrated if their HHI in 2000 is above the median, with the classification held fixed over the sample period.

[Insert Table VI here]

We estimate emission responses separately in the markets with high and low concentration using equation (1) and present the results in Table VI. Under both measures, the estimated emissions response is more than twice as large in ex-ante more concentrated markets. Treated plants in concentrated labor markets increased emissions by 16% ($= e^{0.15} - 1$), compared with a statistically insignificant 4.7% in less concentrated markets (columns (1) and (2)). For emission-market concentration, the corresponding increases are 15.2% ($= e^{0.15} - 1$, significant at the 1% level) and 6.4% (significant at the 10% level) in columns (3) and (4), respectively. The wider treatment band yields a similar pattern in columns (5) to (8). Table A.2 in Online Appendix reports similar findings when concentration is measured alternatively using the combined emissions share of the four largest emitters in each county×3-digit-NAICS industry. Overall, these results suggest that firms increase emissions more in concentrated markets, where incumbents have more rents to protect and can better internalize the higher pollution-control costs associated with nonattainment.

A.5.2 Future PM2.5-Industry Growth Opportunities: Our second proxy for incumbents' benefits of deterring entry is based on future growth opportunities in local

[9] For each county×3-digit-NAICS industry cell, n_i is the number of establishments in each of 12 employment-size categories i: 1–4, 5–9, 10–19, 20–49, 50–99, 100–249, 250–499, 500–999, 1,000–1,499, 1,500–2,499, 2,500–4,999, and 5,000 or more employees. All establishments in each category are assumed to have employment equal to the category midpoint (e_i). For the open-ended top category, we calculate average employment per establishment as total employment in the county-by-industry cell minus $\sum_i n_i \times e_i$ for all smaller categories, divided by n_{5000+}. Finally, for each cell, we calculate estimated total employment E as $\sum_i n_i \times e_i$ and HHI as $\sum_i n_i \times (e_i/E)^2$.

PM2.5-related industries. Incumbents in counties with greater growth opportunities have more to protect from potential entry than those with fewer opportunities. We construct a forward-looking measure based on industry-level market-to-book ratios from the financial markets, which capture expected growth opportunities. Following a procedure similar to Di Maggio et al. (2022), we calculate the market-capitalization-weighted average market-to-book ratio of public firms in each PM2.5-related four-digit NAICS industry and aggregate these ratios to the county level using each county's employment shares in the corresponding industries. The resulting proxy is agnostic to local costs and benefits and captures each county's exposure to industries that financial markets expect to grow. We construct this measure in the pre-event year 2000, classify counties as high- or low-growth based on the cross-sectional median across sample counties, and hold this classification fixed throughout the sample period.

[Insert Table VII here]

We estimate emissions responses among the plants in PM2.5-intensive industries (defined in Section IV.A.2) separately for plants in counties with high and low future growth opportunities in PM2.5 industries. We use the specification in equation (1) and report the results in Table VII. Columns (1) and (4) show that treated plants in high-growth-opportunity counties increased emissions by 25% ($= e^{0.22} - 1$) relative to plants in control counties, whereas the corresponding estimate for low-growth-opportunity counties is only 1.3% and statistically insignificant. Widening the treatment band to ±2 yields a similar pattern (columns (2) and (5)), as does the continuous DD specification (columns (3) and (6)). These findings, together with the heterogeneous results across market concentration, suggest that incumbents increase emissions more where the potential gains from deterring future entrants are greater.

B. Mechanism: How Firms Increased Emissions

Our results so far indicate that firms in counties ex ante closer to the PM2.5 regulatory threshold increased emissions during the statutory PM2.5 monitoring period relative to firms in counties slightly farther from the threshold. The natural question is how. Did firms increase emissions by bringing future production forward to the

statutory monitoring period to take advantage of the lower pollution-control costs before stricter nonattainment requirements take effect, a form of production front-loading? We examine this question next and identify two mechanisms through which firms increased emissions.

B.1. Within-Plant Cross-Media Emission Substitution Toward Air

If firms increased emissions through production front-loading, the total quantity of chemical handled should also increase, as higher production would require greater input use and generate more waste. We therefore examine the effects on the total chemical quantity handled at the plant level, as well as on plants' chemical handling and management practices. We measure the total chemical quantity handled (*Tot Release*) as the sum of on-site releases, off-site transfers, and on-site waste management, including treatment, recycling, and energy recovery.

[Insert Table VIII here]

Table VIII presents the results from estimating equation (1). The total quantity of chemicals handled did not increase at treated plants relative to control plants (columns (1) and (2)). Instead, treated plants increased the share released directly into the air by approximately 1.0–1.1 percentage points (columns (3) and (4)), while reducing the share managed on-site or transferred off-site by approximately 1.4 percentage points (columns (7) and (8)). The share released into water and land also increased, but by a much smaller amount of approximately 0.33–0.37 percentage points (columns (5) and (6)). Together, these findings suggest that higher air emissions stem from changes in chemical handling practices rather than increased production. Specifically, plants appear to engage in release-medium management, shifting chemicals toward direct air releases rather than on-site treatment, recycling, or off-site transfers. This behavior is consistent with our hypothesized pollution moat strategy: firms strategically alter their chemical management practices to raise air emissions and, in turn, increase the likelihood of future county nonattainment.

B.2. Cross-Plant Emission Reallocation Within Multi-plant Firms

Prior studies document cross-plant emission reallocation in several settings (Bisetti et al. (2026), Xu and Kim (2022), Ben-David et al. (2021)). Under our hypothesized strategy, multi-plant firms are particularly well positioned to shift emissions from other counties to those closer to the regulatory threshold, thereby influencing nonattainment probabilities without increasing firm-wide production or emissions. We consider this possibility next by examining emissions at standalone plants and at plants owned by multi-plant firms operating in both treated and control counties. We use equation (1) and report the results in Table IX. Columns (1) and (2) show no change in emissions at treated standalone plants relative to control plants. In contrast, plants belonging to multi-plant firms increased air emissions by approximately 13.6% ($= e^{0.13} - 1$) in the narrow treatment band (column (3), statistically significant at the 5% level) and 9% in the wider band (column (4), statistically significant at the 10% level).

[Insert Table IX here]

We next examine whether the emission increases at plants of multi-plant firms reflect cross-plant reallocation or higher parent-firm-wide emissions. To do so, we aggregate air emissions across plants to the parent-firm level. We classify firms operating in both treated and control counties as treated parents and firms operating in control counties but not treated counties as control parents. The latter lack an incentive to reallocate emissions toward treated counties because they have no plants there. Column (5) compares these parent firms, controlling for the total number of plants and including parent-firm and year fixed effects. We find no statistically significant increase in aggregate emissions among treated parents relative to control parents. Together, these findings contradict increases in parent-firm-wide emissions and suggest cross-plant reallocation as a mechanism through which multi-plant firms increased emissions in treated counties.

C. Did the Pollution Moat Strategy Benefit Incumbents?

The preceding results suggest that firms increased emissions where pollution was more likely to increase future nonattainment probability and where restricting com-

petition was more valuable. This strategy nevertheless exposes incumbents to RACT requirements, which raise incumbents' pollution-control costs above those they face under attainment. Firms would therefore pursue it only if nonattainment creates entry barriers that benefit incumbents. We examine two outcomes during the post-nonattainment period to evaluate whether incumbents benefit from this strategy: new plant openings and incumbent firms' profitability and market valuations.

C.1. Plant Openings

Nonattainment designation may discourage new plants openings in the county, because it imposes stricter permitting and emissions-control requirements on new major facilities. To examine this prediction, we estimate county-level DD regressions comparing counties designated as nonattainment in 2005 with counties that remain in attainment. Plant opening is measured as the number of TRI plants reporting for the first time in a county-year. The sample period for this analysis extends to five or ten years after the end of the statutory monitoring period in 2003. We choose five- and ten-year post-measurement horizons based on the initial duration of pollution-control requirements under the nonattainment provisions (EPA (2007)). We use the following specification:

$$\begin{aligned} y_{ct} = \alpha_0 + \beta_1 \text{Non Atn}_c \times \mathbb{1}[\text{Post}]_t + \beta_2 \text{Non Atn}_c \times \mathbb{1}[\text{Post} + K \text{ years}]_t \\ + \text{Controls} + \mu_c + \gamma_t + \epsilon_{ct}. \end{aligned} \tag{4}$$

Here, *Non Atn*$_c$ indicates counties designated as PM2.5 nonattainment in 2005. The first post indicator, *Post*, covers 2001–2003, while the second, *Post* + *k years*, covers 2004 through 2003 + K, where K = 5 or 10. The specification includes county and year fixed effects and the county-level controls used in the baseline regressions.

[Insert Table X here]

Table X shows that counties that became nonattainment substantially experienced fewer plant openings. Relative to attainment counties, the number of new plants is approximately 16% (= $1 - e^{-0.17} - 1$) lower over the five-year horizon and 18% (= $1 - e^{-0.20} - 1$) lower over the ten-year horizon (columns (1) and (3)). The negative coefficients on *Post* in these columns indicate that plant openings also declined during

the statutory PM2.5 monitoring period. This pattern is consistent with potential entrants deciding against entry in anticipation of nonattainment and higher RACT compliance costs. Finally, the Poisson pseudo-maximum likelihood estimates in columns (2) and (4) indicate similar reductions. Overall, the decline in new plant openings is consistent with nonattainment acting as a barrier to entry.

C.2. Firm Profits and Market-to-Book Ratios

We next examine whether incumbent firms benefit from these potential nonattainment-related barriers. As discussed previously (footnote 1), strategic nonattainment may distort location choices of potential entrants, raising their operational costs, which may yield profit advantages for incumbents. Financial markets may recognize such advantages and raise incumbent firms' market valuations. We test these predictions by linking Compustat firm-level profitability and valuation data to plant-level air emissions in nonattainment and attainment counties using the TRI-Compustat linkages from Hsu, Li, and Tsou (2023). We measure each firm's emissions intensity separately for nonattainment and attainment counties annually as total TRI air emissions from plants in each county category scaled by firm sales. To capture emissions only from pre-PM2.5-nonattainment plants, we only consider the plants that were already operational by 2004 (based on each plant's first year of TRI reporting). We use the following regression specification:

$$\begin{aligned} y_{it} = \alpha \; &+ \; \beta_1 \, \text{Emission}_{NonAtn,\,it} \times \mathbb{1}[\text{Post} + K \text{ years}]_t \; + \; \beta_2 \, \text{Emission}_{NonAtn,\,it} \times \mathbb{1}[\text{Post}]_t \\ &+ \; \beta_3 \, \text{Emission}_{Atn,\,it} \times \mathbb{1}[\text{Post} + K \text{ years}]_t \; + \; \beta_4 \, \text{Emission}_{Atn,\,it} \times \mathbb{1}[\text{Post}]_t \\ &+ \text{Controls} + \mu_i + \gamma_t + \epsilon_{it}. \end{aligned} \tag{5}$$

The specification includes separate indicators for *Post* and [*Post* + *K years*], allowing the relationship between firm outcomes and emissions exposure to differ between 2001–2003 (monitoring period) and the subsequent five- and ten-year periods. The two dependent variables are (i) the natural logarithm of profit, defined as operating income after depreciation and amortization, and (ii) the market-to-book ratio. The specification includes firm and year fixed effects and controls for firm size, sales, and property, plant, and equipment.

[Insert Table XI here]

Table XI presents the results. Higher emissions intensity in contemporaneous non-attainment counties is associated with higher profits and higher market-to-book ratio over the five years following the monitoring period in columns (1) and (2) and ten years in columns (3) and (4). In contrast, emissions intensity in attainment counties is not significantly associated with either outcome. The results are similar when controlling for industry-specific time trends using 2-digit NAICS industry×year fixed effects (Online Appendix, Table A.3). The positive associations of profits and market valuation with nonattainment emission intensity during the post-nonattainment period, driven by continued production at their pre-nonattainment plants, are consistent with incumbent firms benefiting from the pollution moat strategy.

V. Conclusion

We examine the market power consequences of the nonattainment provisions of the U.S. Clean Air Act (CAA), one of the world's most consequential environmental regulatory frameworks. Under the CAA, nonattainment designations trigger additional pollution-control requirements for all firms operating in affected areas. Similar to grandfathering provisions common across regulatory regimes, these requirements impose disproportionately higher compliance costs on new entrants, who face stringent Lowest Achievable Emission Rate (LAER) standards. On the other hand, existing plants face less demanding Reasonably Available Control Technology (RACT) requirements. We show that incumbent firms strategically exploit this statutory asymmetry by increasing emissions during the pre-designation statutory monitoring period, when higher emissions can raise the likelihood of county nonattainment and expose potential entrants to higher compliance costs. We term this strategy a pollution moat: incumbents increase pollution strategically to create entry barriers and strengthen their market power.

Using plant-by-chemical-level emissions data from the TRI and a difference-in-differences design that exploits counties' proximity to the PM2.5 nonattainment threshold, we show that plants in counties closer to the threshold increased air

emissions during the statutory monitoring period relative to plants in counties slightly farther from it. The increase was larger for chemicals more closely linked to PM2.5 emissions. It was also larger at plants in PM2.5-intensive industries, at plants closer to monitors, and at plants in more concentrated markets or in counties with stronger expected growth in PM2.5-related industries. We find no evidence that higher production drove the increase in air emissions. Instead, our findings suggest that it stemmed from firms shifting emissions from other media toward the air and multi-plant firms reallocating emissions across facilities toward plants located in counties near the regulatory PM2.5 threshold. The strategy pays off in the post-designation period. Counties experience fewer new plant openings over the subsequent five- to ten-year horizons. Moreover, greater contemporaneous emission intensity in nonattainment counties correlates with higher firm profits and market-to-book ratios.

Overall, we highlight that incumbent firms may exploit regulations that impose heavier compliance burdens on entrants than on incumbents to strengthen their market power. Designing policies that mitigate incentives for such strategic behavior may achieve pollution reductions while limiting industry concentration.

References

Akey, Pat and Ian Appel, 2021, The limits of limited liability: Evidence from industrial pollution, *The Journal of Finance* 76, 5–55.

Aldy, Joseph E, Maximilian Auffhammer, Maureen Cropper, Arthur Fraas, and Richard Morgenstern, 2022, Looking back at 50 years of the Clean Air Act, *Journal of Economic Literature* 60, 179–232.

Becker, Randy and Vernon Henderson, 2000, Effects of air quality regulations on polluting industries, *Journal of Political Economy* 108, 379–421.

Becker, Randy A, 2005, Air pollution abatement costs under the Clean Air Act: Evidence from the PACE survey, *Journal of Environmental Economics and Management* 50, 144–169.

Bellon, Aymeric, 2025, Does private equity ownership make firms cleaner? The role of environmental liability risks, *The Review of Financial Studies* 38, 2517–2556.

Ben-David, Itzhak, Yeejin Jang, Stefanie Kleimeier, and Michael Viehs, 2021, Exporting pollution: Where do multinational firms emit CO2?, *Economic Policy* 36, 377–437.

Bertrand, Marianne and Francis Kramarz, 2002, Does entry regulation hinder job creation? Evidence from the French retail industry, *The Quarterly Journal of Economics* 117, 1369–1413.

Bisetti, Emilio, Stefan Lewellen, Arkodipta Sarkar, and Xiao Zhao, 2026, Smokestacks and the Swamp, *The Review of Financial Studies*, hhag039.

Bishop, Kelly C, Jonathan D Ketcham, and Nicolai V Kuminoff, 2023, Hazed and confused: The effect of air pollution on dementia, *Review of Economic Studies* 90, 2188–2214.

Chay, Kenneth Y and Michael Greenstone, 2005, Does air quality matter? Evidence from the housing market, *Journal of Political Economy* 113, 376–424.

CRS, 2010, Particulate Matter (PM2.5): Implementation of the 1997 National Ambient Air Quality Standards (NAAQS), *Congressional Research Service*, RL32431, https://www.everycrsreport.com/reports/RL32431.html.

Di Maggio, Marco, Amir Kermani, Rodney Ramcharan, Vincent Yao, and Edison Yu, 2022, The pass-through of uncertainty shocks to households, *Journal of Financial Economics* 145, 85–104.

Djankov, Simeon, Rafael La Porta, Florencio Lopez-de-Silanes, and Andrei Shleifer, 2002, The regulation of entry, *The Quarterly Journal of Economics* 117, 1–37.

Duchin, Ran, Janet Gao, and Qiping Xu, 2025, Sustainability or greenwashing: Evidence from the asset market for industrial pollution, *The Journal of Finance* 80, 699–754.

EPA, 2007, Final Clean Air fine particle implementation rule for implementation of 1997 PM2.5 standards fact sheet, *Environmental Protection Agency*, March, https://www3.epa.gov/pmdesignations/1997standards/documents/Mar07/factsheet.htm.

EPA, 2010, Guidance on developing fee programs required by Clean Air Act Section 185 for the 1-hour Ozone NAAQS, *Environmental Protection Agency*, January 5, https://dep.nj.gov/wp-content/uploads/boss/isg/2010/21910guidance-on-fee-programs-required-by-the-caa-section-185.pdf.

EPA, n.d.(a), Nonattainment NSR Basic Information, *U.S. Environmental Protection Agency*, https://www.epa.gov/nsr/nonattainment-nsr-basic-information.

EPA, n.d.(b), PM2.5 Design Value Notes, *U.S. Environmental Protection Agency*, https://www3.epa.gov/airquality/greenbook/pm25_notes.html.

EPA-commissioned Report, 1980, RACT determination for five industry categories in Florida, *PEDCo Environmental, Inc.*, A report under U.S. Environmental Protection Agency Contract No 68 02 3173.

Greenstone, Michael, 2002, The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 Clean Air Act amendments and the census of manufactures, *Journal of Political Economy* 110, 1175–1219.

Greenstone, Michael, 2003, Estimating regulation-induced substitution: The effect of the Clean Air Act on water and ground pollution, *American Economic Review* 93, 442–448.

Greenstone, Michael, John A List, and Chad Syverson, 2012, The effects of environmental regulation on the competitiveness of US manufacturing, Working Paper, University of Chicago.

Henderson, J Vernon, Sept. 1996, Effects of air quality regulation, *American Economic Review* 86, 789–813. https://ideas.repec.org/a/aea/aecrev/v86y1996i4p789-813.html.

Hsu, Po-Hsuan, Kai Li, and Chi-Yang Tsou, 2023, The pollution premium, *The Journal of Finance* 78, 1343–1392.

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, 848–902.

Klapper, Leora, Luc Laeven, and Raghuram Rajan, 2006, Entry regulation as a barrier to entrepreneurship, *Journal of Financial Economics* 82, 591–629.

List, John A, Daniel L Millimet, Per G Fredriksson, and W Warren McHone, 2003, Effects of environmental regulations on manufacturing plant births: Evidence from a propensity score matching estimator, *Review of Economics and Statistics* 85, 944–952.

McAllister, Lesley Krista, 2007, Putting persuasion back in the equation: Compliance in cap and trade programs, *Pace Environmental Law Review* 24, 299.

Meiselbach, Mark K, Matthew D Eisenberg, Ge Bai, Aditi Sen, and Gerard F Anderson, 2022, Labor market concentration and worker contributions to health insurance premiums, *Medical Care Research and Review* 79, 198–206.

Naaraayanan, S Lakshmi, Kunal Sachdeva, and Varun Sharma, Aug. 2026, The real effects of environmental activist investing, *The Review of Financial Studies*, hhag073.

Stigler, George J., 1971, The theory of economic regulation, *The Bell Journal of Economics and Management Science* 2, 3–21.

Walker, W Reed, 2011, Environmental regulation and labor reallocation: Evidence from the Clean Air Act, *American Economic Review* 101, 442–447.

Xu, Qiping and Taehyun Kim, 2022, Financial constraints and corporate environmental policies, *The Review of Financial Studies* 35, 576–635.

Zou, Eric Yongchen, 2021, Unwatched pollution: The effect of intermittent monitoring on air quality, *American Economic Review* 111, 2101–2126.

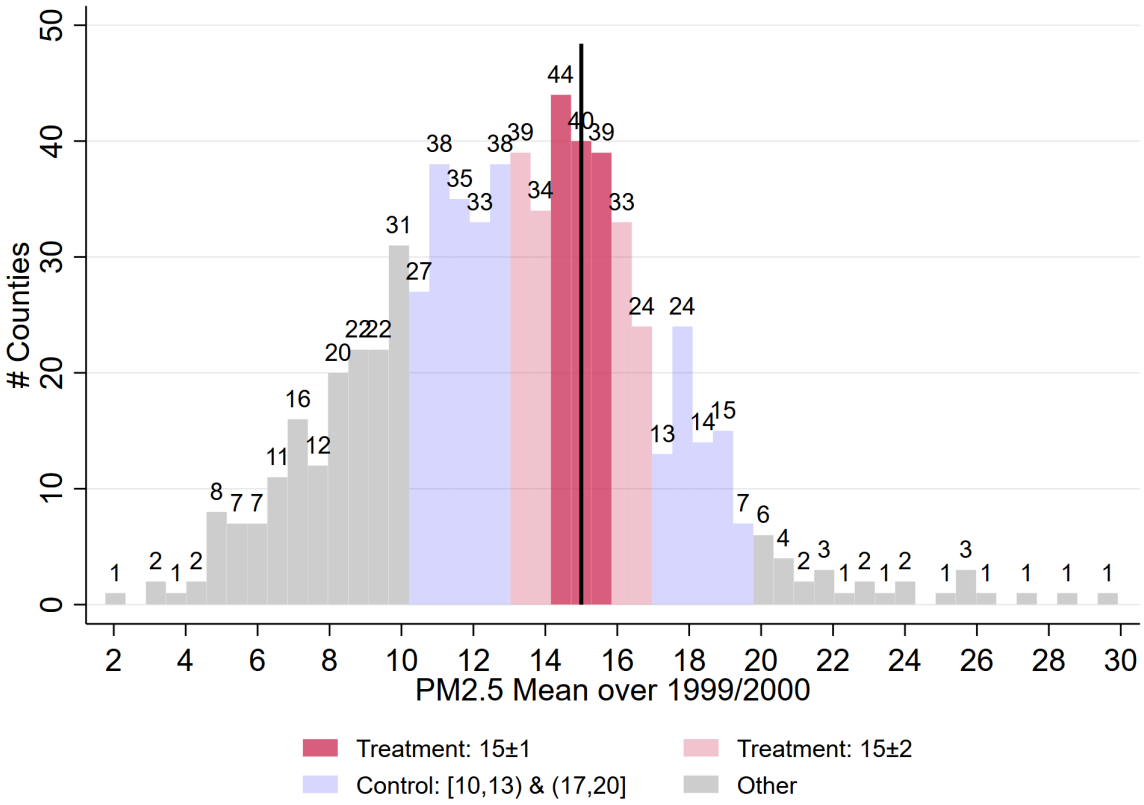


Figure 1. Distribution of County-Level PM2.5 Concentrations. This figure shows the distribution of average county-level PM2.5 concentrations over 1999–2000. The vertical line indicates the 15 $\mu g/m^3$ threshold. The darker red denotes the treated counties with average PM2.5 concentrations between 14 and 16 $\mu g/m^3$. The lighter red extends the treatment band to 13–17 $\mu g/m^3$. Blue indicates control counties with average PM2.5 concentrations in the ranges 10–13 and 17–20 $\mu g/m^3$.

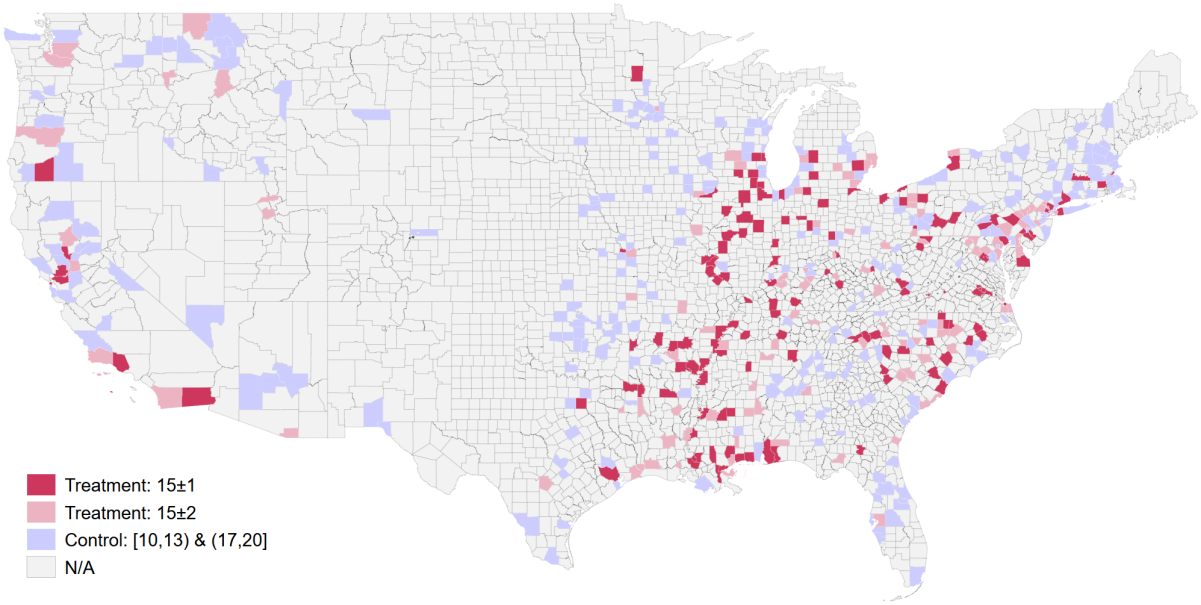


Figure 2. Geographic Distribution of Treated and Control Counties. This figure shows the treated and control counties in the contiguous U.S. based on county-level PM2.5 concentrations. Treated counties are those with average PM2.5 concentrations between 14 and 16 (13 and 17) $\mu g/m^3$, as indicated by the dark (light) red shading. Control counties have average PM2.5 concentrations in the ranges 10–13 and 17–20 $\mu g/m^3$, as indicated by the blue shading.

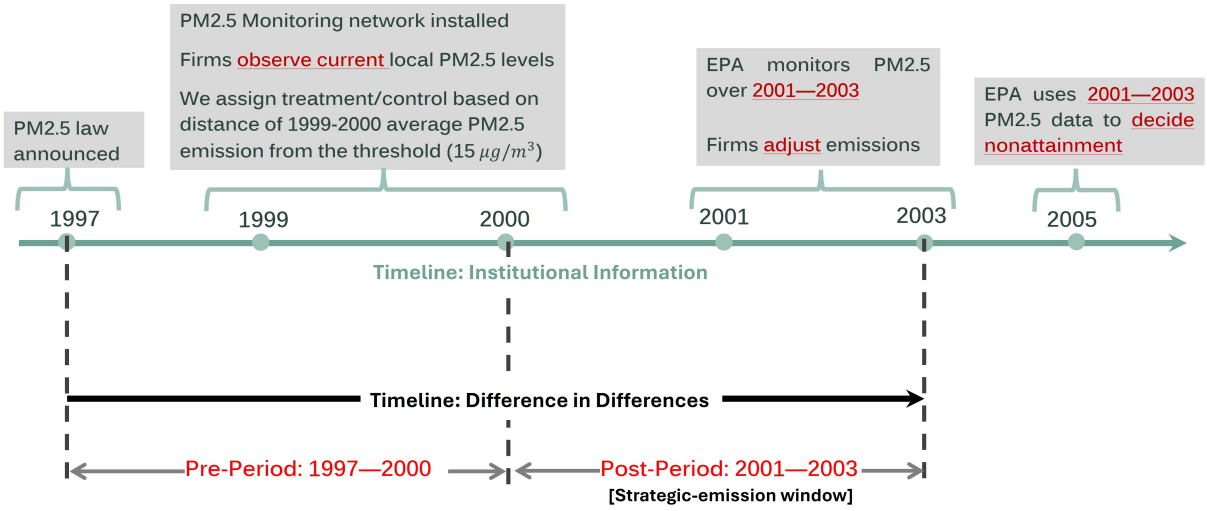


Figure 3. Event Timeline and Empirical Design. This figure shows the PM2.5 standards' implementation timeline and the corresponding empirical design.

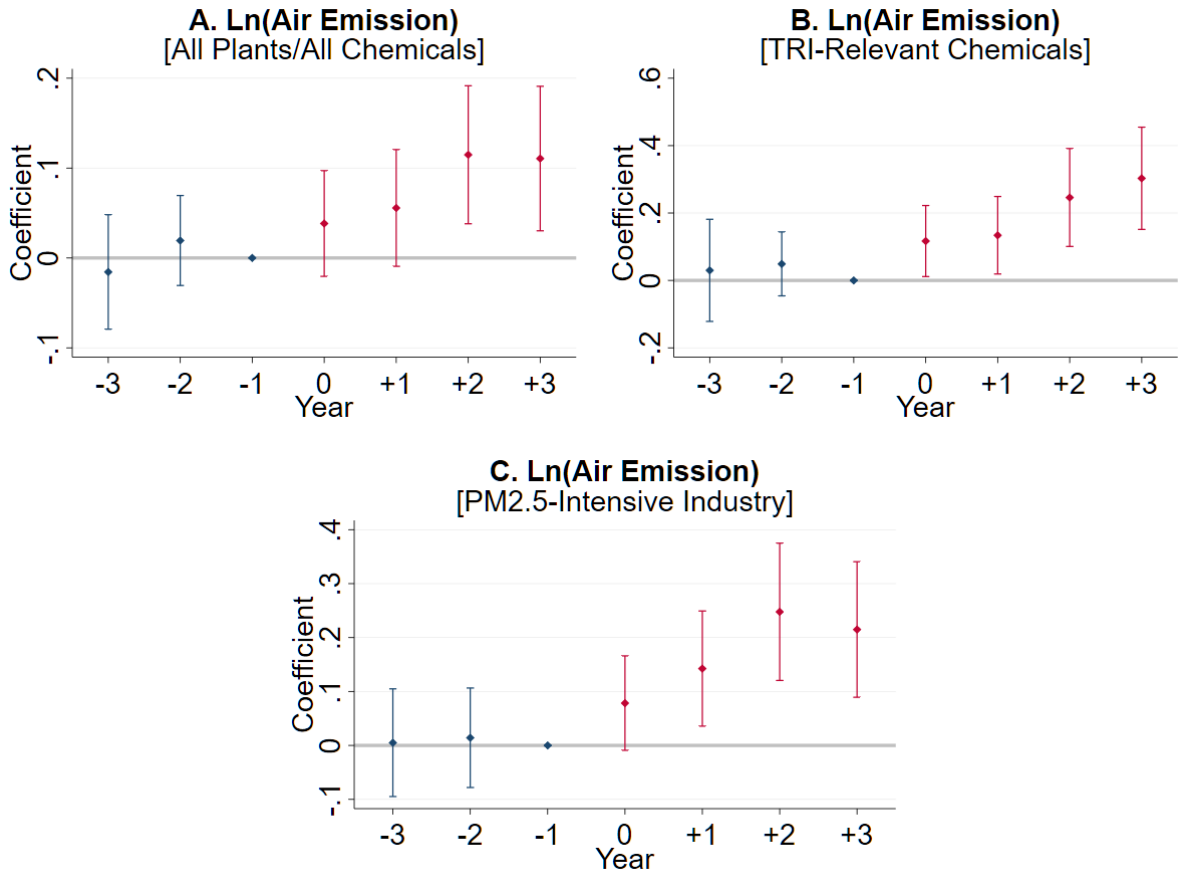


Figure 4. Event-study Plots. This figure plots the estimated coefficients and their 95% confidence intervals in three panels corresponding to the specifications in column (1) of Tables II, III, and IV, respectively. The outcome and sample are denoted in respective panels. Controls are the natural logarithms of county-level population and per capita income. The horizontal axis denotes years relative to 2000 ($\tau = 0$). Confidence intervals are based on standard errors clustered at the county level.

Table I
Summary Statistics

This table reports summary statistics for treated and control observations. The top five variables are measured at the establishment-chemical level over the 1997–2003 sample period. Variable item numbers refer to EPA's (TRI) Basic Plus Data Files, File Type 1A. *Treated* refers to establishments in counties with average PM2.5 concentrations between 14 and 16 $\mu g/m^3$ in 1999–2000, while *Control* refers to establishments in counties with concentrations between 10 and 13 $\mu g/m^3$ or between 17 and 20 $\mu g/m^3$. *Tot Air Emission* (*'000*) is annual air releases of toxic chemicals (Item 115, in thousands of pounds). *Tot Releases* (*'000*) is total releases in thousands of pounds, defined as the sum of on-site releases (Item 218), off-site transfers (Item 262), and on-site waste management (Item 267). *Air/Output %* is air emissions as a percentage of total releases. *Water & Land/Output %* is surface-water discharges plus land releases as a percentage of total releases. *Release Mgmt/Output %* is off-site transfers plus on-site waste management as a percentage of total releases. The bottom two variables are county-level population and per capita income, whose natural logarithms are included as *Controls* in the regressions.

	Treated			Control		
	(1)	(2)	(3)	(4)	(5)	(6)
	N	Mean	SD	N	Mean	SD
Tot Air Emission ('000)	104,092	11.16	76.05	142,147	10.12	92.49
Tot Releases ('000)	104,092	378.43	7529.41	142,147	249.36	4478.17
Air/Tot Releases (%)	86,478	29.03	39.61	116,332	29.81	40.16
Water & Land/Tot Releases (%)	86,478	3.78	17.00	116,332	3.61	16.57
Release Mgmt/Tot Releases (%)	86,478	67.20	41.08	116,332	66.57	41.50
County Population ('000)	929	280.54	417.90	1,700	305.89	402.64
Per Capita Income ('000)	929	27.46	7.47	1,700	27.50	7.48

Table II
Emissions During PM2.5 Measurement Window

All columns report estimates from the DD regression. *Ln*(*Tot Air Emission*) is the natural logarithm of establishment-chemical-level air emissions. *Treat* equals one if a county's average PM2.5 concentration in 1999–2000 falls within the specified treatment band around the 15 $\mu g/m^3$ threshold—either 15±1 (14–16 $\mu g/m^3$) or 15±2 (13–17 $\mu g/m^3$)—and zero if it falls within either control ranges of 10–13 or 17–20 $\mu g/m^3$. *Threshold Proximity* is measured as $5 - |PM2.5 - 15|$ using the county's average PM2.5 concentration in 1999–2000 for counties with average PM2.5 concentrations between 10 and 20 μg/m^3. Higher values indicate greater proximity to the threshold. *Post* is one for years 2001–2003 and zero for years 1997–2000. Controls are the natural logarithms of county-level population and per capita income. All models include establishment, chemical, and year fixed effects. Standard errors are clustered by county and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Outcome:	Ln (Tot Air Emission)		
	(1)	(2)	(3)
Treatment:	15 ± 1	15 ± 2	Cont.
Treat×Post	0.080**	0.058**	
	(0.031)	(0.028)	
Threshold Proximity×Post			0.022**
			(0.011)
FE: Establishment	Y	Y	Y
FE: Chemical	Y	Y	Y
FE: Year	Y	Y	Y
Adjusted R^2	0.69	0.69	0.69
Observations	161,630	201,007	201,007

Table III
Heterogeneity: PM2.5-Relevant TRI Chemicals

All columns report estimates from the DD regression. *PM2.5 Relevance* is measured using chemical-level regressions of county PM2.5 emissions on the air emissions of each TRI chemical in 2008, controlling for state fixed effects and clustering standard errors by state. A chemical is classified as *Very High* (*High*) if its estimated coefficient is positive and statistically significant at the 1% (5%) level, and *Low* if not statistically significant at the 10% level. *Ln*(*Tot Air Emission*) is the natural logarithm of establishment-chemical-level air emissions. *Treat* equals one if a county's average PM2.5 concentration in 1999–2000 falls within the specified treatment band around the 15 $\mu g/m^3$ threshold—either 15±1 (14–16 $\mu g/m^3$) or 15±2 (13–17 $\mu g/m^3$)—and zero if it falls within either control ranges of 10–13 or 17–20 $\mu g/m^3$. *Threshold Proximity* is measured as $5 - |PM2.5 - 15|$ using the county's average PM2.5 concentration in 1999–2000 for counties with average PM2.5 concentrations between 10 and 20 μg/m^3. Higher values indicate greater proximity to the threshold. *Post* is one for years 2001–2003 and zero for years 1997–2000. Controls are the natural logarithms of county-level population and per capita income. All models include establishment, chemical, and year fixed effects. Standard errors are clustered by county and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Outcome:	Ln (Tot Air Emission)								
PM2.5 Relevance of Chemical:	Very High			High			Low		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment:	15 ± 1	15 ± 2	Cont.	15 ± 1	15 ± 2	Cont.	15 ± 1	15 ± 2	Cont.
Treat×Post	0.164***	0.132***		0.114***	0.099**		0.061*	0.045	
	(0.050)	(0.046)		(0.044)	(0.039)		(0.034)	(0.030)	
Threshold Proximity×Post			0.047***			0.033**			0.020
			(0.017)			(0.014)			(0.012)
FE: Establishment	Y	Y	Y	Y	Y	Y	Y	Y	Y
FE: Chemical	Y	Y	Y	Y	Y	Y	Y	Y	Y
FE: Year	Y	Y	Y	Y	Y	Y	Y	Y	Y
Adjusted R^2	0.68	0.68	0.68	0.68	0.69	0.69	0.71	0.71	0.71
Observations	19,462	24,411	24,411	34,883	43,543	43,543	101,006	125,413	125,413

Table IV
Heterogeneity: PM2.5-Intensive Industries

All columns report estimates from the DD regression. *Industry's PM2.5 Share* measures each four-digit NAICS industry's share of total PM2.5 emissions in the 2008 NEI data. It is *High* if the industry is in the top decile of the PM2.5 share distribution and *low* otherwise. *Ln*(*Tot Air Emission*) is the natural logarithm of establishment-chemical-level air emissions. *Treat* equals one if a county's average PM2.5 concentration in 1999–2000 falls within the specified treatment band around the 15 $\mu g/m^3$ threshold—either 15±1 (14–16 $\mu g/m^3$) or 15±2 (13–17 $\mu g/m^3$)—and zero if it falls within either control ranges of 10–13 or 17–20 $\mu g/m^3$. *Threshold Proximity* is measured as $5 - |PM2.5 - 15|$ using the county's average PM2.5 concentration in 1999–2000 for counties with average PM2.5 concentrations between 10 and 20 μg/m^3. Higher values indicate greater proximity to the threshold. *Post* is one for years 2001–2003 and zero for years 1997–2000. Controls are the natural logarithms of county-level population and per capita income. All models include establishment, chemical, and year fixed effects. Standard errors are clustered by county and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Outcome:	Ln (Tot Air Emission)					
Industry's PM2.5 Share:	High			Low		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment:	15 ± 1	15 ± 2	Cont.	15 ± 1	15 ± 2	Cont.
Treat×Post	0.173***	0.141***		0.004	-0.015	
	(0.047)	(0.043)		(0.033)	(0.030)	
Threshold Proximity×Post			0.054***			-0.004
			(0.016)			(0.013)
FE: Establishment	Y	Y	Y	Y	Y	Y
FE: Chemical	Y	Y	Y	Y	Y	Y
FE: Year	Y	Y	Y	Y	Y	Y
Adjusted R^2	0.65	0.65	0.65	0.75	0.75	0.75
Observations	71,186	89,768	89,768	90,327	111,106	111,106

Table V
Heterogeneity: Monitor Proximity

All columns report estimates from the DD regression. *Dist. to Monitor* (*mi.*) is measured as the average distance from an establishment to its nearest EPA PM2.5 monitor in 1999 and 2000. We identify PM2.5 monitors using EPA parameter code 88101 under the 1997 annual PM2.5 standard. Columns (1) and (3) report results for establishments with an average distance to the nearest monitor of 5 miles or less, while columns (2) and (4) report results for those more than 5 miles away. *Ln*(*Tot Air Emission*) is the natural logarithm of establishment-chemical-level air emissions. *Treat* equals one if a county's average PM2.5 concentration in 1999–2000 falls within the specified treatment band around the 15 $\mu g/m^3$ threshold—either 15±1 (14–16 $\mu g/m^3$) or 15±2 (13–17 $\mu g/m^3$)—and zero if it falls within either control ranges of 10–13 or 17–20 $\mu g/m^3$. *Post* is one for years 2001–2003 and zero for years 1997–2000. Controls are the natural logarithms of county-level population and per capita income. All models include establishment, chemical, and year fixed effects. Standard errors are clustered by county and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Outcome:	Ln (Tot Air Emission)			
Treatment:	15 ± 1		15 ± 2	
	(1)	(2)	(3)	(4)
Dist. to Monitor (mi.):	≤ 5	> 5	≤ 5	> 5
Treat×Post	0.105**	0.048	0.085**	0.025
	(0.045)	(0.037)	(0.040)	(0.036)
FE: Establishment	Y	Y	Y	Y
FE: Chemical	Y	Y	Y	Y
FE: Year	Y	Y	Y	Y
Adjusted R^2	0.68	0.70	0.68	0.70
Observations	89,746	69,250	108,119	89,575

Table VI
Heterogeneity: Ex-Ante Market Concentration

All columns report estimates from the DD regression. Columns (1), (2), (5), and (6) split the sample by *Emission HHI*, while columns (3), (4), (7), and (8) split the sample by *Labor HHI*. Concentration is *High* for establishments in county–three-digit NAICS industries with above-median pre-event HHI, and *Low* otherwise; the classification is fixed over time. *Emission HHI* is the sum of squared establishment shares of total TRI air emissions within a county-industry in 2000. *Labor HHI* is the sum of squared establishment employment shares within a county-industry, calculated from County Business Patterns employment-size categories in 2000. *Ln*(*Tot Air Emission*) is the natural logarithm of establishment-chemical-level air emissions. *Treat* equals one if a county's average PM2.5 concentration in 1999–2000 falls within the specified treatment band around the 15 $\mu g/m^3$ threshold—either 15±1 (14–16 $\mu g/m^3$) or 15±2 (13–17 $\mu g/m^3$)—and zero if it falls within either control ranges of 10–13 or 17–20 $\mu g/m^3$. *Post* is one for years 2001–2003 and zero for years 1997–2000. Controls are the natural logarithms of county-level population and per capita income. All models include establishment, chemical, and year fixed effects. Standard errors are clustered by county and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Outcome:	Ln (Tot Air Emission)							
Treatment:	15 ± 1				15 ± 2			
HHI Market:	Labor		Emission		Labor		Emission	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Concentration:	High	Low	High	Low	High	Low	High	Low
Treat×Post	0.150**	0.047	0.149**	0.064*	0.129**	0.011	0.164***	0.032
	(0.061)	(0.043)	(0.069)	(0.033)	(0.050)	(0.042)	(0.061)	(0.031)
FE: Establishment	Y	Y	Y	Y	Y	Y	Y	Y
FE: Chemical	Y	Y	Y	Y	Y	Y	Y	Y
FE: Year	Y	Y	Y	Y	Y	Y	Y	Y
Adjusted R^2	0.68	0.70	0.77	0.67	0.68	0.70	0.76	0.67
Observations	55,194	89,028	27,760	130,710	71,089	108,921	35,163	161,735

Table VII
Heterogeneity: County Growth Opportunities

All columns report estimates from the DD regression for establishments in *PM2.5-relevant industries*, defined as industries in the top decile of the PM2.5-share distribution. *County PM2.5 Ind Growth Opportunities* is the county-level market-to-book ratio in 2000, calculated as an average of industry-level market-to-book ratios weighted by local employment shares across PM2.5-relevant industries in the county. Industry-level market-to-book ratios are calculated as the market-capitalization-weighted average of firm-level market-to-book ratios among public firms in each four-digit NAICS industry. The growth opportunities measure is *High* if the establishment's county has an above-median value and *Low* otherwise. *Ln*(*Tot Air Emission*) is the natural logarithm of establishment-chemical-level air emissions. *Treat* equals one if a county's average PM2.5 concentration in 1999–2000 falls within the specified treatment band around the 15 $\mu g/m^3$ threshold—either 15±1 (14–16 $\mu g/m^3$) or 15±2 (13–17 $\mu g/m^3$)—and zero if it falls within either control ranges of 10–13 or 17–20 $\mu g/m^3$. *Threshold Proximity* is measured as $5 - |PM2.5 - 15|$ using the county's average PM2.5 concentration in 1999–2000 for counties with average PM2.5 concentrations between 10 and 20 μg/m^3. Higher values indicate greater proximity to the threshold. *Post* is one for years 2001–2003 and zero for years 1997–2000. Controls are the natural logarithms of county-level population and per capita income. All models include establishment, chemical, and year fixed effects. Standard errors are clustered by county and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Outcome:	Ln (Tot Air Emission)					
County PM2.5 Ind Growth Opportunities:	High			Low		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment:	15 ± 1	15 ± 2	Cont.	15 ± 1	15 ± 2	Cont.
Treat×Post	0.229***	0.170***		0.009	0.063	
	(0.048)	(0.045)		(0.095)	(0.097)	
Threshold Proximity×Post			0.065***			0.021
			(0.018)			(0.033)
FE: Establishment	Y	Y	Y	Y	Y	Y
FE: Chemical	Y	Y	Y	Y	Y	Y
FE: Year	Y	Y	Y	Y	Y	Y
Adjusted R^2	0.65	0.65	0.65	0.70	0.69	0.69
Observations	53,827	68,271	68,271	17,312	21,463	21,463

Table VIII
Mechanism: Cross-Media Emission Substitution Toward Air

All columns report estimates from the DD regression. All outcome variables are measured at the establishment-chemical level. *Ln(Tot Releases)* is the natural logarithm of total output, defined as the sum of on-site releases, off-site transfers, and on-site waste management. *Air/Tot Releases %* measures the air emissions as a percentage of total output. *Water & Land/Tot Releases %* measures the surface-water discharges plus land releases as a percentage of total output. *Release Mgmt/Tot Releases %* measures the off-site transfers plus on-site waste management as a percentage of total output. *Treat* equals one if a county's average PM2.5 concentration in 1999–2000 falls within the specified treatment band around the 15 $\mu g/m^3$ threshold—either 15±1 (14–16 $\mu g/m^3$) or 15±2 (13–17 $\mu g/m^3$)—and zero if it falls within either control ranges of 10–13 or 17–20 $\mu g/m^3$. *Post* is one for years 2001–2003 and zero for years 1997–2000. Controls are the natural logarithms of county-level population and per capita income. All models include establishment, chemical, and year fixed effects. Standard errors are clustered by county and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Outcome:	Ln (Tot Releases)		$\frac{\text{Air}}{\text{Tot Releases}}$ %		$\frac{\text{Water \& Land}}{\text{Tot Releases}}$ %		$\frac{\text{Release Mgmt}}{\text{Tot Releases}}$ %	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment:	15 ± 1	15 ± 2	15 ± 1	15 ± 2	15 ± 1	15 ± 2	15 ± 1	15 ± 2
Treat×Post	0.002	-0.026	1.045***	1.091***	0.370*	0.325*	-1.415***	-1.415***
	(0.034)	(0.027)	(0.368)	(0.324)	(0.218)	(0.182)	(0.426)	(0.364)
FE: Establishment	Y	Y	Y	Y	Y	Y	Y	Y
FE: Chemical	Y	Y	Y	Y	Y	Y	Y	Y
FE: Year	Y	Y	Y	Y	Y	Y	Y	Y
Adjusted R^2	0.59	0.59	0.60	0.60	0.48	0.49	0.58	0.58
Observations	201,432	250,947	201,432	250,947	201,432	250,947	201,432	250,947

Table IX
Mechanism: Cross-Plant Emission Reallocation Toward Treated Plants Within Multi-Plant Firms

All columns report estimates from the DD regression. Columns (1) and (2) report results for standalone establishments, while columns (2) and (4) report results for establishments whose parent firms have establishments in both treated and control counties. *Ln*(*Tot Air Emission*) is the natural logarithm of establishment-chemical-level air emissions. *Treat* equals one if a county's average PM2.5 concentration in 1999–2000 falls within the specified treatment band around the 15 $\mu g/m^3$ threshold—either 15±1 (14–16 $\mu g/m^3$) or 15±2 (13–17 $\mu g/m^3$)—and zero if it falls within either control ranges of 10–13 or 17–20 $\mu g/m^3$. *Post* is one for years 2001–2003 and zero for years 1997–2000. Controls are the natural logarithms of county-level population and per capita income. All models include establishment, chemical, and year fixed effects. Standard errors are clustered by county and reported in parentheses. Column (5) reports the parent-firm-level specification. *Treat* equals one if the parent firm operates at least one establishment in a treated county (13–17 μg/m^3) and at least one establishment in a control county (10–13 or 17–20 μg/m^3), and zero if the parent firm operates at least one establishment in a control county but none in treated counties. The specification controls for the natural logarithm of the number of establishments operated by the parent firm and includes parent-firm and year fixed effects. Standard errors are clustered at the parent-firm level and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Outcome:	Ln (Tot Air Emission)				
Establishment Type:	Standalone		Parent in both Trt/Ctrl		Parent Firm Agg
	(1)	(2)	(3)	(4)	(5)
Treatment:	15 ± 1	15 ± 2	15 ± 1	15 ± 2	-
Treat×Post	0.023	0.028	0.128**	0.087*	0.002
	(0.039)	(0.036)	(0.051)	(0.047)	(0.063)
Panel	Est×Chem	Est×Chem	Est×Chem	Est×Chem	Parent Firm
FE: Establishment	Y	Y	Y	Y	N
FE: Chemical	Y	Y	Y	Y	N
FE: Parent	N	N	N	N	Y
FE: Year	Y	Y	Y	Y	Y
Adjusted R^2	0.74	0.73	0.64	0.65	0.89
Observations	86,745	108,765	46,566	59,474	6,381

Table X
New Plant Entry Post-Nonattainment

All columns report estimates from the DD regression for new plant entry in counties with average PM2.5 concentrations between 10 and 20 $\mu g/m^3$ in 1999–2000. *New Estab* is the number of new plants in a county-year, measured as the number of TRI plants first observed reporting in that county-year. *Ln*(*New Estab*) is the natural logarithm of the number of new plants. Columns (1) and (3) report OLS estimates using *Ln*(*New Estab*) as the outcome, while columns (2) and (4) report Poisson pseudo-maximum likelihood (PPML) estimates using *New Estab*. *Non Atn* equals one if the county was designated as NAAQS nonattainment in 2005, and zero otherwise. $\mathbb{1}[Post]$ indicates the post-treatment observation period used in the baseline specification, equal to one for 2001–2003 and zero otherwise. $\mathbb{1}[Post + K\ years]$ indicates the subsequent post-treatment period, equal to one for 2004 through $2003 + K$ and zero for 1997–2003. Columns (1) and (2) set $K = 5$, while columns (3) and (4) set $K = 10$. Controls are the natural logarithms of county-level population and per capita income. All models include county and year fixed effects. Standard errors are clustered by county and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Horizon Years (K):	$K = 5$		$K = 10$	
	(1)	(2)	(3)	(4)
Outcome:	Ln (New Estab)	New Estab	Ln (New Estab)	New Estab
[Model]	[OLS]	[PPML]	[OLS]	[PPML]
Non Atn×$\mathbb{1}$[Post+K years]	-0.172***	-0.211***	-0.203***	-0.261***
	(0.061)	(0.080)	(0.060)	(0.073)
Non Atn×$\mathbb{1}$[Post]	-0.141**	-0.303***	-0.148***	-0.302***
	(0.057)	(0.079)	(0.057)	(0.079)
FE: County	Y	Y	Y	Y
FE: Year	Y	Y	Y	Y
Adjusted R^2	0.52	0.45	0.50	0.43
Observations	3,040	5,856	4,023	8,364

Table XI
Firm Profits and Valuation Post-Nonattainment

All columns report estimates from firm-level regressions. $Emission_{NonAtn}$ ($Emission_{Atn}$) measures a firm's TRI air emissions from plants located in nonattainment (attainment) counties, scaled by firm sales. The calculation includes only plants whose first year of TRI reporting is 2004 or earlier. $\mathbb{1}[Post]$ indicates the post-treatment observation period used in the baseline specification, equal to one for 2001–2003 and zero otherwise. $\mathbb{1}[Post+\ K\ years]$ indicates the subsequent post-treatment period, equal to one for 2004 through $2003 + K$ and zero for 1997–2003. Columns (1) and (2) set $K = 5$, while columns (3) and (4) set $K = 10$. $Ln(Profit)$ is the natural logarithm of a firm's operating income after depreciation and amortization. M/B is the market-to-book ratio. Controls are the natural logarithms of total assets, sales, and property, plant, and equipment. All models include firm and year fixed effects. Standard errors are clustered by firm and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Horizon Years (K):	$K = 5$		$K = 10$	
	(1)	(2)	(3)	(4)
Outcome:	Ln(Profit)	M/B	Ln(Profit)	M/B
$Emission_{NonAtn} \times \mathbb{1}[Post+K\ years]$	0.223***	7.183***	0.257***	7.791***
	(0.056)	(1.378)	(0.060)	(1.236)
$Emission_{NonAtn} \times \mathbb{1}[Post]$	0.067	1.758**	-0.045	2.184**
	(0.245)	(0.743)	(0.315)	(0.982)
$Emission_{Atn} \times \mathbb{1}[Post+K\ years]$	-0.078	0.354	-0.082	0.004
	(0.062)	(0.792)	(0.062)	(0.489)
$Emission_{Atn} \times \mathbb{1}[Post]$	-0.027	-0.042	-0.039	-0.182
	(0.052)	(0.311)	(0.059)	(0.247)
FE: Firm	Y	Y	Y	Y
FE: Year	Y	Y	Y	Y
Adjusted R^2	0.93	0.28	0.93	0.07
Observations	5,031	4,943	7,316	7,255

Online Appendix to

"The Pollution Moat: Strategic Emissions under the Clean Air Act"

Christine Zhuowei Huang Amit Kumar Mayank Kumar*

* Christine Zhuowei Huang is at Leibniz Institute for Financial Research SAFE; email: huang@safe-frankfurt.de. Amit Kumar is at Singapore Management University; email: amitkumar@smu.edu.sg. Mayank Kumar is at Questrom School of Business, Boston University; email: mkumar2@bu.edu.

A. Additional Regression Results

Table A.1
Placebo

All columns report estimates from the DD regression. *Ln*(*Tot Air Emission*) is the natural logarithm of establishment-chemical-level air emissions. *PM2.5 Relevance* is measured using chemical-level regressions of county PM2.5 emissions on the air emissions of each TRI chemical in 2008, controlling for state fixed effects and clustering standard errors by state. A chemical is classified as *Very High* (*High*) if its estimated coefficient is positive and significant at the 1% (5%) level. *Treat* equals one if an establishment is located in a county without a PM2.5 monitor in 1999–2000 in columns (1), (3), and (5), and in a hotspot county without a PM2.5 monitor in 1999–2000 in columns (2), (4), and (6). *Hotspot County* is a county in the top 5% of counties by the employment share of at least one of the 12 PM2.5-related three-digit NAICS industries, based on average employment shares in 1999–2000. *Post* is one for years 2001–2003 and zero for years 1997–2000. Controls are the natural logarithms of county-level population and per capita income. All models include establishment, chemical, and year fixed effects. Standard errors are clustered by county and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Outcome:	Ln (Tot Air Emission)					
	All		Very High PM2.5 Relevant Chemicals		High PM2.5 Relevant Chemical	
	(1)	(2)	(3)	(4)	(5)	(6)
Placebo Counties:	All Unmonitored	Hotspot Unmonitored	All Unmonitored	Hotspot Unmonitored	All Unmonitored	Hotspot Unmonitored
Treat×Post	0.034	0.032	0.054	0.039	0.037	0.032
	(0.025)	(0.031)	(0.051)	(0.066)	(0.036)	(0.047)
FE: Establishment	Y	Y	Y	Y	Y	Y
FE: Chemical	Y	Y	Y	Y	Y	Y
FE: Year	Y	Y	Y	Y	Y	Y
Adjusted R^2	0.72	0.71	0.69	0.68	0.78	0.77
Observations	216,542	144,745	22,795	16,529	47,325	32,994

Table A.2
Heterogeneity by Emission Market Concentration using the Four-Plant Concentration (CR4) Ratio

All columns report estimates from the DD regression. Concentration is *High* for establishments in county–three-digit NAICS industries with above-median pre-event CR4, and *Low* otherwise; the classification is fixed over time. *CR4* is the combined share of total TRI air emissions accounted for by the four largest plants within a county and three-digit NAICS industry, averaged over 1999–2000. *Ln*(*Tot Air Emission*) is the natural logarithm of establishment-chemical-level air emissions. *Treat* equals one if a county's average PM2.5 concentration in 1999–2000 falls within the specified treatment band around the 15 $\mu g/m^3$ threshold—either 15±1 (14–16 $\mu g/m^3$) or 15±2 (13–17 $\mu g/m^3$)—and zero if it falls within either control ranges of 10–13 or 17–20 $\mu g/m^3$. *Post* is one for years 2001–2003 and zero for years 1997–2000. Controls are the natural logarithms of county-level population and per capita income. All models include establishment, chemical, and year fixed effects. Standard errors are clustered by county and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Outcome:	Ln (Tot Air Emission)			
Treatment:	15 ± 1		15 ± 2	
	(1)	(2)	(3)	(4)
Concentration:	High	Low	High	Low
Treat×Post	0.146***	0.013	0.132***	-0.016
	(0.044)	(0.043)	(0.038)	(0.039)
FE: Establishment	Y	Y	Y	Y
FE: Chemical	Y	Y	Y	Y
FE: Year	Y	Y	Y	Y
Adjusted R^2	0.74	0.64	0.74	0.64
Observations	74,456	84,012	95,759	101,155

Table A.3
Long-Term Effects on Firm Profits and Valuation [with Firm and Industry×Year FE]

All columns report estimates from the DD regression. $Emission_{NonAtn}$ ($Emission_{Atn}$) is a firm's TRI air emissions from establishments located in nonattainment (attainment) counties, scaled by firm sales. Only plants whose first year of TRI reporting is 2004 or earlier enter the emission calculations. $\mathbb{1}[Post]$ indicates the post-treatment observation period used in the baseline specification, equal to one for 2001–2003 and zero otherwise. $\mathbb{1}[Post+\ K\ years]$ indicates the subsequent post-treatment period, equal to one for 2004 through $2003+K$ and zero for 1997–2003. Columns (1) and (2) set $K = 5$, while columns (3) and (4) set $K = 10$. $Ln(Profit)$ is the natural logarithm of a firm's operating income after depreciation and amortization. M/B is the market-to-book ratio. Controls are the natural logarithms of total assets, sales, and property, plant, and equipment. All models include firm and two-digit NAICS industry×year fixed effects. Standard errors are clustered by firm and reported in parentheses. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

Horizon Years (K):	$K=5$		$K=10$	
	(1)	(2)	(3)	(4)
Outcome:	Ln(Profit)	M/B	Ln(Profit)	M/B
$Emission_{NonAtn} \times \mathbb{1}[Post+K\ years]$	0.209***	7.384***	0.239***	7.734***
	(0.061)	(1.425)	(0.065)	(1.300)
$Emission_{NonAtn} \times \mathbb{1}[Post]$	0.163	2.080***	0.071	2.527**
	(0.238)	(0.607)	(0.316)	(0.992)
$Emission_{Atn} \times \mathbb{1}[Post+K\ years]$	-0.074	-0.049	-0.075	-0.194
	(0.065)	(0.803)	(0.064)	(0.512)
$Emission_{Atn} \times \mathbb{1}[Post]$	-0.045	-0.085	-0.049	-0.166
	(0.052)	(0.322)	(0.059)	(0.266)
FE: Firm	Y	Y	Y	Y
FE: NAICS 2×Year	Y	Y	Y	Y
Adjusted R^2	0.93	0.27	0.93	0.05
Observations	4,982	4,896	7,243	7,186