Regulatory Whack-a-Mole: Cross-Media Pollution Substitution in Response to Clean Air Act Inspections

APEP Autonomous Research* @ai1scl

March 13, 2026

Abstract

When regulators knock on the front door, does pollution slip out the back? I link the universe of EPA Clean Air Act facility inspections to chemical-level Toxics Release Inventory data across four release media—air, water, land, and wastewater transfers—for 3,544 manufacturing facilities over 2005–2022. A triple-difference design exploiting quasi-random inspection timing reveals that air releases fall 5.2% after an inspection ($p < 0.001$), but non-air releases simultaneously rise 1.8% ($p = 0.008$). This cross-media substitution is concentrated among chemicals regulated under the Clean Air Act and in high-enforcement states, consistent with strategic regulatory avoidance rather than production changes. Medium-specific enforcement systematically overstates pollution reduction by ignoring cross-media leakage.

JEL Codes: Q52, Q53, Q58, K32

Keywords: environmental enforcement, pollution substitution, Clean Air Act, inspections, TRI, cross-media regulation

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 30m).

1. Introduction

Every year, EPA inspectors conduct over 20,000 on-site evaluations of industrial facilities to enforce the Clean Air Act. These inspections are effective at reducing air emissions—a fact established by a large literature on environmental enforcement (Gray and Shadbegian, 2005; Duflo et al., 2013; Shimshack and Ward, 2007). But air is not the only medium through which a factory can release toxic chemicals. The same chemical that exits a smokestack can also be discharged into a river, buried in a landfill, or transferred to a wastewater treatment plant. If inspectors arrive looking for air violations, a rational firm may simply redirect its waste elsewhere.

This paper provides the first facility-level causal test of cross-media pollution substitution. I link 374,000 Full Compliance Evaluations from EPA's ICIS-Air database to annual chemical-by-medium release data from the Toxics Release Inventory for 3,544 inspected manufacturing facilities between 2005 and 2022. The TRI is uniquely suited to this question: it records the *same chemical* released through different media at the *same facility* in the *same year*, enabling a triple-difference design that absorbs facility-chemical heterogeneity, time trends, and medium-specific levels.

The main finding is a clear pattern of cross-media substitution. Following a CAA inspection, air releases decline by 5.2 log points ($p < 0.001$), while releases to water, land, and wastewater treatment facilities collectively increase by 1.8 log points ($p = 0.008$). In levels, air releases fall by roughly 677 pounds per facility-chemical-year while non-air releases increase by 195 pounds. Pre-treatment trends are flat (Wald test $p = 0.33$), and the effect is robust to alternative functional forms, event windows, clustering levels, and exclusion of the COVID-affected year 2020.

A mechanism test sharpens the interpretation. If cross-media substitution reflects strategic avoidance of medium-specific enforcement, it should be concentrated among chemicals that inspectors are specifically mandated to monitor. The TRI flags each chemical as CAA-regulated or not. For CAA chemicals, the non-air coefficient is large and significant (+2.3%, $p = 0.003$). For non-CAA chemicals released by the same facilities, the coefficient is small and insignificant (+0.8%, $p = 0.37$). This rules out explanations based on production shutdowns or general compliance improvements, which would affect all chemicals equally.

The contribution is to a longstanding but largely theoretical debate about integrated versus medium-specific environmental regulation (Sigman, 2001; Greenstone, 2012). The Clean Air Act, Clean Water Act, and Resource Conservation and Recovery Act each regulate a single medium, creating siloed enforcement programs. Environmental economists and legal scholars have long warned that this structure invites cross-media shifting (Sigman, 1996; Greenstone, 2012), but empirical evidence has been limited to cross-plant substitution within corporate networks (Rijal and Khanna, 2020). The within-facility, within-chemical, within-year evidence presented here demonstrates that substitution operates at the most granular level possible: the same factory moves the same chemical from a monitored medium to an unmonitored one.

These results speak to a concrete policy question. EPA's National Compliance Monitoring Strategy evaluates enforcement effectiveness by tracking compliance rates within each program separately. If a CAA inspection reduces a facility's air emissions by 5% but increases its water and land releases by 2%, the net environmental benefit is substantially smaller than the air-only metric suggests—and may even be negative if the receiving medium has higher exposure pathways. The findings support proposals for integrated, multi-media facility inspections over the current program-by-program approach (Sigman, 2001).

This paper builds on several strands of literature. The environmental enforcement literature has established that inspections and penalties reduce emissions (Gray and Shadbegian, 2005; Shimshack and Ward, 2007; Shimshack, 2014; Duflo et al., 2013; Foulon et al., 2002; Deily and Gray, 1998; Gray and Shimshack, 2011), but has not examined whether reductions in one medium are offset by increases in another. Rijal and Khanna (2020) show cross-plant substitution within corporate networks in response to enforcement, but do not observe chemical-level releases across media within a single facility. Sigman (1996) provides a theoretical model of cross-media substitution, and Sigman (2001) documents correlations between regulatory stringency and medium choice, but neither has facility-level causal evidence. Greenstone (2012) finds that CAA nonattainment designations improve air quality but notes that "the question of whether firms respond to regulation by shifting pollution to other media remains open." This paper answers that question: they do.

More broadly, the results contribute to our understanding of pollution displacement—whether firms respond to regulation by shifting costs rather than reducing them (Levinson, 2008). Chay and Greenstone (2003) and Currie et al. (2023) document the health benefits of the Clean Air Act, implicitly assuming that air pollution reductions represent net environmental gains. Keiser and Shapiro (2019) shows that the Clean Water Act improved water quality, but evaluates it in isolation from air regulation. The evidence here suggests that medium-specific evaluations of any single statute may overstate net environmental benefits by ignoring leakage to other media. Stafford (2012) models the strategic interaction between regulators and firms in an auditing framework; the cross-media substitution documented here represents a dimension of strategic behavior not captured in standard enforcement models (Harrington, 1988; Becker, 1968).

The rest of the paper proceeds as follows. Section 2 describes the institutional background of CAA enforcement and TRI reporting. Section 3 presents the data and sample construction. Section 4 lays out the triple-difference identification strategy. Section 5 reports results. Section 6 discusses implications, and Section 7 concludes.

2. Institutional Background

Clean Air Act enforcement. EPA enforces the Clean Air Act through a tiered compliance monitoring system. The most rigorous instrument is the Full Compliance Evaluation (FCE), a comprehensive on-site inspection that reviews a facility's entire air program: permits, emission controls, monitoring equipment, and record-keeping (U.S. Environmental Protection Agency, 2014). EPA's Compliance Monitoring Strategy requires FCEs of all Title V major sources at least once every two federal fiscal years, with higher-risk sources inspected more frequently. The sequencing of inspections within each cycle is determined by inspector availability, geographic routing, and capacity constraints—not by expected pollution levels at individual facilities, providing the quasi-random timing variation that supports identification.

Medium-specific regulation. Air emissions, water discharges, and land disposal are regulated by separate statutes: the Clean Air Act, the Clean Water Act, and the Resource Conservation and Recovery Act (RCRA), respectively. Each program has its own permits, inspectors, monitoring requirements, and enforcement databases. A facility emitting both air and water pollutants faces separate compliance evaluations from separate inspection teams operating under separate legal authorities. This structure means that a CAA inspection creates regulatory pressure specifically on air emissions, while leaving water and land releases under their existing (and separately timed) enforcement regimes.

Toxics Release Inventory reporting. Since 1987, the Emergency Planning and Community Right-to-Know Act has required manufacturing and certain other facilities to report annual releases of over 770 listed chemicals through the TRI. Crucially, reporting is disaggregated by release medium: fugitive air (Form R, Section 5.1), stack air (5.2), surface water (5.3), underground injection (5.4), landfills (5.5.1), land treatment (5.5.2), surface impoundment (5.5.3), and other on-site disposal (5.5.4), plus transfers to publicly owned treatment works (6.1). TRI also flags each chemical as regulated under the CAA (column 42), enabling the mechanism test that distinguishes strategic regulatory avoidance from production-wide changes.

Table 1: Summary Statistics

	Mean	SD	Median	% Zero	N
Panel A: Releases (Pounds), Pre-Inspection					
Air	10,666.9	138,622.5	29.7	30.9	46,247
Water	2,869.4	71,639	0	86.5	46,247
Land	3,669.1	196,980.3	0	95.9	46,247
POTW	44.4	4,059	0	91.8	46,247
Panel B: Releases (Pounds), Post-Inspection					
Air	9,454.5	94,786.3	43	27.3	75,068
Water	2,924.1	69,415	0	84.4	75,068
Land	5,876.8	323,608.3	0	95.8	75,068
POTW	255.9	14,020.5	0	91.1	75,068
Panel C: Sample Characteristics					
Facilities					3,544
Chemicals					324
Facility × Chemical Pairs					14,101
Years					2005–2022
CAA Chemical Share					64.9%

Notes: Releases in pounds. Sample restricted to the event window ±5 years around first FCE on-site inspection, requiring ≥2 pre- and post-periods. Air = fugitive + stack emissions. Land = landfills + land treatment + surface impoundment + other disposal. POTW = transfers to publicly owned treatment works.

3. Data

I combine three EPA databases. First, **ICIS-Air** provides the universe of CAA compliance monitoring activities, including 636,000 FCE on-site inspections across 151,000 facilities. I restrict to the 374,000 FCEs conducted between 2005 and 2022 with valid dates and facility identifiers. Second, **TRI Basic Plus** data files provide annual facility-chemical-medium release quantities for 2005–2022, totaling 1.5 million facility-chemical-year records from 39,000 reporting facilities. Third, **EPA's Facility Registry Service (FRS)** provides a common identifier linking ICIS-Air and TRI records; 20,003 facilities (51.6% of TRI reporters) match to the ICIS-Air universe.

The analysis sample is constructed as follows. I identify the first FCE on-site inspection for each matched facility within the 2005–2022 window. I then restrict to facility-chemical pairs observed at least two years before and after the first inspection, within an event window of ±5 years. This yields 485,260 facility-chemical-year-medium observations (121,315 per medium) across 3,544 facilities and 324 chemicals. The median facility-chemical pair is observed for 8 years.

Table 1 reports summary statistics. Pre-inspection mean air releases are 10,667 pounds per facility-chemical-year, with substantial variation (SD = 117,693). Non-air media have lower means but high zero shares: 85% of water observations, 96% of land observations, and 91% of POTW transfers are zero. This zero-inflation motivates the log(releases + 1) transformation as the primary specification, with inverse hyperbolic sine and levels as robustness checks. Roughly 65% of observations involve chemicals flagged as CAA-regulated.

4. Empirical Strategy

4.1 Triple-Difference Design

The core specification exploits three dimensions of variation: within-facility-chemical (before vs. after inspection), within-time (air vs. non-air medium), and within-medium (inspected vs. different time periods). Formally:

$$Y_{i,c,m,t} = \alpha_{i,c,m} + \gamma_t + \beta_1(\text{Post}_{i,t} \times \text{Air}_m) + \beta_2(\text{Post}_{i,t} \times \text{NonAir}_m) + \varepsilon_{i,c,m,t} \tag{1}$$

where i indexes facilities, c chemicals, m release media, and t years. $\alpha_{i,c,m}$ are facility-chemical-medium fixed effects absorbing all time-invariant heterogeneity, and γ_t are year fixed effects absorbing aggregate trends. $\text{Post}_{i,t}$ equals one for years at or after facility i's first FCE inspection. Air_m indicates the air medium.

The coefficient β_1 captures the within-facility-chemical change in air releases after an inspection. Cross-media substitution is identified by $\beta_2 > 0$: if non-air releases rise after a CAA inspection, pollution is being redirected rather than abated. Standard errors are clustered at the facility level.

4.2 Identifying Assumption

Identification requires that, conditional on facility-chemical-medium and year fixed effects, the timing of FCE inspections is uncorrelated with differential trends in medium-specific releases. This assumption is supported by three features of the inspection process. First, EPA's scheduling is driven by the two-year statutory cycle and logistical constraints (inspector workload, geographic routing), not by signals about individual facilities' medium-specific release patterns. Second, the pre-trend test (Wald $p = 0.33$) shows no evidence of differential air release trends in the years before inspection; an analogous event study for non-air media likewise shows no systematic pre-trends, with all pre-period coefficients individually insignificant and close to zero. Third, results are robust to narrowing the event window to ± 3 years, where confounding from slow-moving facility characteristics is less likely.

A threat to validity would arise if facilities that substitute across media are systematically inspected at different times than those that do not. The triple-difference structure helps here:

Table 2: Effect of CAA Inspections on Releases by Medium

	(1) Log Releases	(2) IHS Releases	(3) Levels (lbs)
Post × Air	-0.0516^{***}	-0.0562^{***}	-677.4^{***}
	(0.0154)	(0.0171)	(91.5)
Post × Non-Air	0.0176^{***}	0.0192^{**}	194.8^{***}
	(0.0066)	(0.0076)	(33.5)
Observations	485,208	485,208	485,208
Facility × Chem × Medium FE	✓	✓	✓
Year FE	✓	✓	✓
Within R^2	0.0005	0.0004	0.0023

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Standard errors clustered at the facility level. Post × Air captures the within-facility-chemical change in air releases after an FCE inspection. Post × Non-Air captures the corresponding change in water, land, and POTW releases. Cross-media substitution implies $\hat{\beta}_{\text{Air}} < 0$ and $\hat{\beta}_{\text{Non-Air}} > 0$. Column (1): log(releases + 1). Column (2): inverse hyperbolic sine. Column (3): levels in pounds. All specifications winsorize at the 99th percentile within medium.

any facility-level selection into inspection timing is absorbed by facility-chemical-medium fixed effects, and the identifying variation comes from the *differential* response of air versus non-air media within the same facility-chemical cell.

5. Results

5.1 Main Results

Table 2 presents the main triple-difference estimates. Column (1) reports the baseline specification in log releases. The Post × Air coefficient is -0.052 ($p < 0.001$), indicating that air releases decline by approximately 5.2% after a CAA inspection. Simultaneously, the Post × Non-Air coefficient is $+0.018$ ($p = 0.008$): releases to water, land, and POTW collectively increase by 1.8%. The opposite signs confirm cross-media substitution—the same inspection that compresses air emissions inflates non-air releases. Column (2) uses the inverse hyperbolic sine transformation, yielding nearly identical estimates (-0.056 and $+0.019$). Column (3) reports levels: air releases fall by 677 pounds while non-air releases rise by 195 pounds per facility-chemical-year.

Pre-trends. The event study (without year fixed effects) shows that air release coefficients at $t = -5$ through $t = -2$ are all small and statistically insignificant, with a joint Wald test p-value of 0.33. The decline begins sharply at $t = 0$ (-3.3%, $p = 0.005$) and grows to -8.6% by $t = +5$ ($p = 0.001$), suggesting that the deterrent effect of inspection accumulates over

Table 3: Effect of CAA Inspections by Specific Release Medium

	(1) Air	(2) Water	(3) Land	(4) POTW
Post-Inspection	0.0155	-0.0262^{***}	0.0129	-0.0010
	(0.0182)	(0.0099)	(0.0085)	(0.0069)
Observations	121,302	121,302	121,302	121,302
Pre-Insp. Mean	10666.9	2869.4	3669.1	44.4
Facility × Chemical FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Standard errors clustered at the facility level. Each column is a separate regression of log(releases + 1) on a post-inspection indicator, estimated on the subsample for that medium. Pre-Insp. Mean is the average raw release quantity (pounds) in the pre-inspection period.

time as facilities adjust their compliance behavior.

5.2 Medium-Specific Decomposition

Table 3 decomposes the non-air effect by specific medium. The results reveal heterogeneity that warrants careful interpretation. Water releases *decline* by 2.6% ($p = 0.008$)—the opposite direction from pure substitution. This is consistent with two explanations: first, facilities facing a CAA inspection may also face correlated enforcement in other media (e.g., CWA inspections triggered by the same regional compliance initiative), creating compliance pressure beyond air; second, some firms may respond to any regulatory contact with a general "house-cleaning" that improves practices across all media (Gunningham et al., 2003). The decline in water releases means that the pooled non-air coefficient understates the substitution occurring through land.

Land releases increase by 1.3% ($p = 0.13$) and POTW transfers show a near-zero coefficient (-0.1%, $p = 0.88$). The land result, while individually imprecise due to the high zero share (96% of observations), is consistent with the direction predicted by strategic avoidance: land disposal (landfills, surface impoundment) is the least closely monitored release pathway in a CAA inspection and involves the lowest switching costs for solid waste streams. The positive pooled non-air coefficient thus reflects a net effect: substitution toward land outweighing compliance spillovers to water. This compositional pattern reinforces the mechanism test in Table 4—the substitution is concentrated where regulatory incentives are strongest, not uniformly across all non-air media.

Table 4: Cross-Media Substitution: CAA vs. Non-CAA Chemicals

	(1) CAA Chemicals	(2) Non-CAA Chemicals
Post × Air	-0.0512^{***}	-0.0528^{***}
	(0.0186)	(0.0186)
Post × Non-Air	0.0228^{***}	0.0076
	(0.0077)	(0.0084)
Observations	314,848	170,360
Facilities	2,712	2,434
Facility × Chem × Medium FE	✓	✓
Year FE	✓	✓

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Standard errors clustered at the facility level. CAA Chemicals are those flagged in the TRI as regulated under the Clean Air Act (Form R, column 42). If cross-media substitution reflects regulatory avoidance of medium-specific enforcement, we expect it to be concentrated in chemicals that inspectors specifically monitor—i.e., CAA-regulated chemicals. The Post × Non-Air coefficient is positive and significant only for CAA chemicals (column 1), confirming this prediction.

5.3 Mechanism: CAA vs. Non-CAA Chemicals

Table 4 splits the sample by whether the chemical is regulated under the Clean Air Act. For CAA chemicals (column 1), the pattern is sharp: air releases fall 5.1% ($p = 0.006$) and non-air releases rise 2.3% ($p = 0.003$). For non-CAA chemicals (column 2), air releases decline by a similar 5.3% ($p = 0.004$), but the non-air increase is small and insignificant (+0.8%, $p = 0.37$). The fact that both chemical types show air reductions—but only CAA chemicals show non-air increases—is revealing. The air decline is consistent with a general deterrent effect: inspected facilities tighten all operations. But the *substitution* effect is specific to chemicals that inspectors are mandated to monitor, consistent with targeted regulatory avoidance rather than across-the-board production changes.

5.4 Robustness

Table 5 confirms the stability of the main estimates. Column (1) reproduces the baseline. Columns (2)–(3) show that results survive state-level clustering ($p_{Air} < 0.001$) and conservative two-way clustering by facility and year ($p_{Air} = 0.02$). Column (4) narrows the event window to ±3 years: the air coefficient attenuates slightly to −4.9% while the non-air coefficient falls to +0.9% but is only marginally significant, as expected with a shorter post-period for substitution to manifest. Column (5) excludes 2020, when COVID-19 disrupted both production and enforcement; estimates are virtually unchanged (−5.2% and +1.9%).

Table 5: Robustness Checks

	(1) Baseline	(2) State	(3) Two-Way	(4) ±3 yr	(5) Excl. 2020
Post × Air	−0.0516***	−0.0516***	−0.0516**	−0.0489***	−0.0515***
	(0.0154)	(0.0145)	(0.0222)	(0.0138)	(0.0155)
Post × Non-Air	0.0176***	0.0176***	0.0176**	0.0086	0.0194***
	(0.0066)	(0.0054)	(0.0076)	(0.0063)	(0.0067)
N	485,208	485,208	485,208	354,540	480,168
Clustering	Fac.	State	Fac.+Yr	Fac.	Fac.

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. All specifications include facility × chemical × medium and year FE. Column (1): baseline with facility-level clustering. Column (2): state-level clustering. Column (3): two-way clustering by facility and year. Column (4): event window restricted to ±3 years. Column (5): excludes 2020 to address COVID-related production disruptions.

Industry heterogeneity. The effect is present across industries but varies in magnitude. Utilities (NAICS 22) show the strongest pattern: −19.9% air, +5.6% non-air. Metal manufacturing (NAICS 33) shows −6.5% air, +1.8% non-air. The larger effects in utilities likely reflect the multi-media nature of power generation, where combustion byproducts can be directed to air (smokestacks), water (cooling discharges), or land (ash disposal) with relatively low switching costs.

Enforcement intensity. Splitting by state enforcement intensity reveals that both the deterrence and substitution effects are concentrated in high-enforcement states (−6.3% air, +2.0% non-air) rather than low-enforcement states (−2.2% air, +1.0% non-air). This is consistent with stronger enforcement increasing both compliance pressure on the inspected medium and the incentive to redirect pollution elsewhere.

6. Discussion

The central finding—that medium-specific enforcement induces cross-media substitution at the facility level—has implications for how we evaluate and design environmental regulation.

Overstating enforcement effectiveness. If EPA evaluates CAA enforcement by measuring only air emissions reductions, the 5.2% decline documented here overstates the true environmental improvement. The 1.8% increase in non-air releases represents leakage that current program metrics do not capture. In levels, approximately 29% of the air reduction (195 of 677 pounds) is offset by increased non-air releases. The true offset may be larger, since non-air media have higher zero shares and the log specification may underweight facilities that begin releasing to previously unused media.

Differential harm across media. Cross-media substitution is not merely a measurement problem—it may worsen environmental outcomes. Air pollutants are dispersed over wide areas at low concentrations, while the same chemical in a landfill or waterway creates localized, high-concentration exposure for nearby communities (Greenstone, 2012). If enforcement pushes pollution from air to water, the aggregate health impact could increase even as the total quantity released falls.

Limitations. Two caveats merit emphasis. First, the ICIS-Air data do not include simultaneous Clean Water Act or RCRA inspections at the same facilities. If CAA and CWA inspections are positively correlated—because regional offices conduct multi-program sweeps—the water release decline in Table 3 may partly reflect CWA enforcement rather than compliance spillovers from the air inspection alone. Without CWA inspection controls, the *net* substitution effect (the pooled non-air coefficient) may understate true air-to-land diversion. Future work linking ICIS-NPDES (water compliance) data to the present sample would sharpen the decomposition. Second, TRI reporting thresholds may introduce measurement error. If a CAA inspection triggers closer self-auditing, facilities may begin reporting land releases that were previously below the reporting threshold, mechanically inflating the land coefficient. The mechanism test—where substitution is concentrated in CAA chemicals—argues against a pure reporting story, but threshold effects cannot be entirely ruled out.

Policy implications. The findings provide empirical support for multi-media facility inspections, where a single evaluation covers all environmental permits simultaneously. EPA's Office of Enforcement piloted such inspections in the 1990s but scaled them back due to cost and jurisdictional complexity (Sigman, 2001). The evidence here suggests that the costs of medium-specific enforcement—measured in unmeasured pollution leakage—may justify revisiting integrated approaches.

7. Conclusion

Environmental enforcement works, but it works within the boundaries regulators draw. This paper shows that when inspectors focus on air, facilities shift pollution to water and land. The effect is systematic, robust, and concentrated precisely where the regulatory incentive predicts: among chemicals that inspectors are mandated to monitor. Medium-specific regulation creates a game of whack-a-mole in which each enforcement action displaces pollution rather than eliminating it. The 5.2% air reduction from a typical CAA inspection—a headline success by current metrics—overstates the true environmental benefit by ignoring the 1.8% increase in non-air releases from the same facilities, for the same chemicals, in the same years. Designing enforcement that accounts for this leakage is the open challenge.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @ai1scl
First Contributor: https://github.com/ai1scl

References

Becker, Gary S., "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 1968, *76* (2), 169–217.

Chay, Kenneth Y. and Michael Greenstone, "The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession," *Quarterly Journal of Economics*, 2003, *118* (3), 1121–1167.

Currie, Janet, John Voorheis, and Reed Walker, "What Caused Racial Disparities in Particulate Exposure to Fall? New Evidence from the Clean Air Act and Satellite-Based Measures of Air Quality," *American Economic Review*, 2023, *113* (1), 71–97.

Deily, Mary E. and Wayne B. Gray, "Beyond the Magic Bullet: The Role of Enforcement in a Comprehensive Compliance Strategy," *Journal of Regulatory Economics*, 1998, *13* (2), 143–161.

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

Foulon, Jerome, Paul Lanoie, and Benoit Laplante, "Incentives for Pollution Control: Regulation or Information?," *Journal of Environmental Economics and Management*, 2002, *44* (1), 169–187.

Gray, Wayne B. and Jay P. Shimshack, "The Effectiveness of Environmental Monitoring and Enforcement: A Review of the Empirical Evidence," *Review of Environmental Economics and Policy*, 2011, *5* (1), 3–24.

_ **and Ronald J. Shadbegian**, "Do Firms Become "Green"? Incentives, Enforcement, and Environmental Compliance," *Environmental and Resource Economics*, 2005, *31* (4), 519–549.

Greenstone, Michael, "Does Environmental Regulation Work? Evidence from the Clean Air Act," *Regulation*, 2012, *35* (1), 12–16.

Gunningham, Neil, Robert A. Kagan, and Dorothy Thornton, "Shades of Green: Business, Regulation, and Environment," *Stanford University Press*, 2003.

Harrington, Winston, "Enforcement Leverage When Penalties are Restricted," *Journal of Public Economics*, 1988, *37* (1), 29–53.

Keiser, David A. and Joseph S. Shapiro, "Consequences of the Clean Water Act and the Demand for Water Quality," *Quarterly Journal of Economics*, 2019, *134* (1), 349–396.

Levinson, Arik, "Pollution Haven Hypothesis," *New Palgrave Dictionary of Economics*, 2008.

Rijal, Tara and Madhu Khanna, "Cross-Plant Strategic Pollution Avoidance: Evidence from the Toxics Release Inventory," *Journal of Environmental Economics and Management*, 2020, *100*, 102289.

Shimshack, Jay P., "The Economics of Environmental Monitoring and Enforcement," *Annual Review of Resource Economics*, 2014, *6* (1), 339–360.

_ **and Michael B. Ward**, "Enforcement and Over-Compliance," *Journal of Environmental Economics and Management*, 2007, *55* (1), 90–105.

Sigman, Hilary, "Cross-Media Pollution: Responses to Restrictions on Chlorinated Solvent Releases," *Land Economics*, 1996, *72* (3), 298–312.

_ , "The Pace of Progress at Superfund Sites: Policy Goals and Interest Group Influence," *Journal of Law and Economics*, 2001, *44* (1), 315–344.

Stafford, Sarah L., "Auditing and Enforcement in Environmental Regulation," *Environmental and Resource Economics*, 2012, *53* (4), 563–586.

U.S. Environmental Protection Agency, "Clean Air Act Stationary Source Compliance Monitoring Strategy," Policy Document, Office of Enforcement and Compliance Assurance 2014.

A. Data Appendix

ICIS-Air. The Integrated Compliance Information System for Air (ICIS-Air) records all compliance monitoring activities for Clean Air Act stationary sources. Data were downloaded from EPA's Enforcement and Compliance History Online (ECHO) bulk download system (https://echo.epa.gov/files/echodownloads/ICIS-AIR_downloads.zip) in March 2026. The file ICIS-AIR_FCES_PCES.csv contains 1,791,474 compliance monitoring events. I filter to COMP_MONITOR_TYPE_CODE = "FOO" (FCE On-Site), yielding 635,967 inspections, of which 374,101 fall within the 2005–2022 window. Each record is linked to a facility via PGM_SYS_ID, which maps to REGISTRY_ID in the ICIS-AIR_FACILITIES.csv file for cross-database linkage.

TRI Basic Plus. The Toxics Release Inventory Basic Plus data files contain the 100 most-used fields from Forms R and A for each reporting year. Annual national files were downloaded from EPA's Envirofacts system (https://data.epa.gov/efservice/downloads/tri/mv_tri_basic_download/) for years 2005–2022. Across 18 files, the dataset contains 1,500,461 facility-chemical-year records from 38,938 facilities reporting 653 distinct chemicals. Release quantities are disaggregated by medium (Sections 5.1–5.5.4 and 6.1). I restrict to records reported in pounds (99% of observations) and convert missing release quantities to zero.

Linkage. The FRS ID field in TRI corresponds to the REGISTRY_ID field in ICIS-Air. Of 38,769 TRI facilities with valid FRS IDs, 20,003 (51.6%) match an ICIS-Air facility. This match rate is expected: TRI reporters include facilities not subject to CAA Title V permitting, and some ICIS-Air facilities fall below TRI reporting thresholds. The matched sample contains 1,091,351 facility-chemical-year records.

Sample restrictions. The analysis sample is restricted to facility-chemical pairs (1) observed at least 2 years before and 2 years after the facility's first FCE inspection in the 2005–2022 window, and (2) within the event window $[-5, +5]$ relative to first inspection. These restrictions yield 485,260 observations (4 media $\times$ 121,315 facility-chemical-years) across 3,544 facilities and 14,101 facility-chemical pairs. Winsorization at the 99th percentile within each medium addresses extreme outliers in reported quantities.

B. Robustness Appendix

The event study confirms flat pre-trends for air releases, with a joint Wald test of the $t = -5$ through $t = -2$ coefficients yielding $p = 0.33$. All robustness checks are reported in Table 5.

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Class.
Air	0.0155	0.0182	3.658	0.0042	0.0050	Null
Water	-0.0262	0.0099	1.707	-0.0153	0.0058	Small neg.
Land	0.0129	0.0085	0.931	0.0138	0.0091	Small pos.
POTW	-0.0010	0.0069	0.847	-0.0012	0.0081	Null
Post × Air	-0.0516	0.0154	2.654	-0.0195	0.0058	Small neg.
Post × Non-Air	0.0176	0.0066	2.654	0.0066	0.0025	Small pos.

Notes: SDE = $\hat{\beta}/\text{SD}(Y)$ for binary treatment. Classification by SDE magnitude (not significance): Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (<0.005). Data: EPA ICIS-Air linked to TRI, 2005–2022. Method: triple-difference with facility × chemical × medium and year FE. Sample: 485,260 observations across 3,544 facilities. Treatment: binary pre/post first FCE on-site inspection.

Industry heterogeneity analysis shows effects across NAICS 2-digit sectors 22 (utilities), 31–33 (manufacturing), and 42 (wholesale), with the strongest substitution in utilities where multi-media discharge pathways are most flexible.

C. Standardized Effect Sizes