Is Compliance Private? Peer Effects in Regulation and Product-Market Competition

Please click here for the latest version

Vishavdeep Sharma*

Abstract

Regulatory compliance is usually treated as a private firm-regulator decision: each firm chooses how much to comply, bears its own cost, and faces its own regulator. This paper asks whether that view is incomplete when firms' compliance decisions affect one another. I develop a model in which each firm chooses compliance effort that lowers its own operating cost, while rivals' effort affects that cost through a spillover of either sign. Embedding this decision in a differentiated-products oligopoly, I characterize how cooperative and rivalrous compliance spillovers shape equilibrium effort and market outcomes under price and quantity competition. A sufficiently strong cooperative spillover can make price competition more profitable than quantity competition, whereas a rivalrous spillover can lower total surplus under price competition. Consumers are better off under price competition regardless of the sign of the spillover. I then provide empirical evidence on the model's central mechanism using facility-level compliance records under the U.S. Clean Air Act and Clean Water Act merged with firm-level financial data from Compustat. For each firm-year, I construct the compliance of peer firms operating in the same state and three-digit NAICS industry. The evidence is consistent with a pattern in which, within firms, higher peer compliance is linked to higher own unit cost, while higher own compliance is linked to lower unit cost. This pattern remains after accounting for firm, year, and industry-by-year fixed effects and is concentrated among smaller firms. Taken together, the evidence suggests that environmental compliance is not purely private: peer firms' compliance behavior may affect a firm's costs and thereby shape equilibrium compliance, competition, and welfare.

Keywords: Regulatory Compliance; Peer Effects; Cost Spillovers; Environmental Regulation;

*Department of Agricultural, Environmental, and Development Economics (AEDE), The Ohio State University, Columbus, OH 43210, United States. Email: sharma.1094@osu.edu

Product-Market Competition

JEL Classification: D43, D62, L13, L51, Q58

"Too many regulations cause big headaches for small businesses."

Tom Sullivan, Vice President of Small Business Policy, U.S. Chamber of Commerce

1 Introduction

Regulatory compliance has become an increasingly visible part of doing business. Figure 1 documents this pattern for manufacturing firms: the share of firms reporting that the time or resources spent on licensing, compliance, and other government requirements has increased has trended upward over 2021–2025 and recently overtaken the share reporting no change. This evidence is consistent with recent work showing that regulatory compliance absorbs a meaningful share of firm resources (Trebbi et al., 2023). A growing number of firms therefore report that compliance weighs on them.

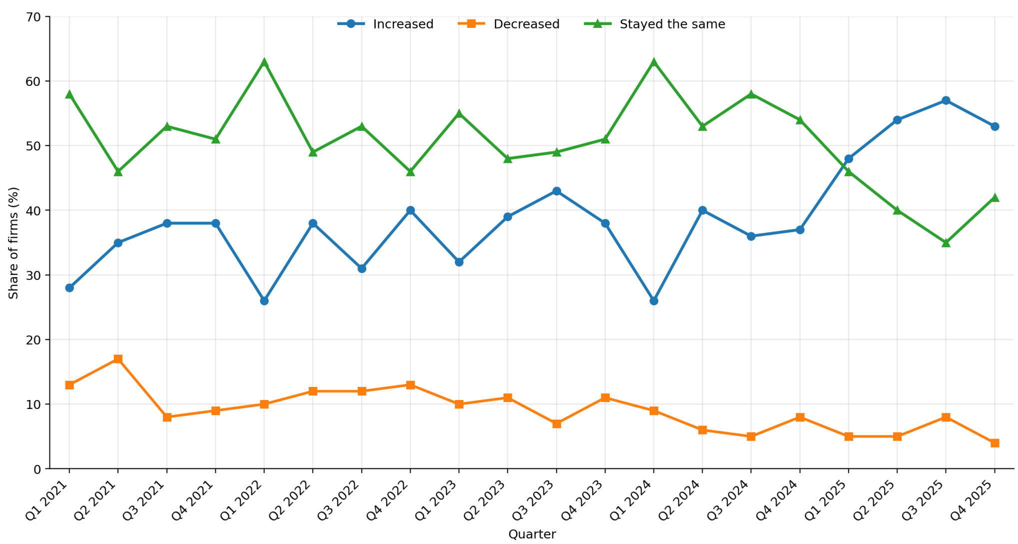

Figure 1. Compliance burden among manufacturing firms. The figure reports the share of manufacturing firms saying that the time or resources spent completing licensing, compliance, and other government requirements increased, decreased, or stayed the same relative to six months earlier.
Source: U.S. Chamber of Commerce, Small Business Index, Manufacturing sector.

The immediate explanation is straightforward: compliance effort is costly, and profit-maximizing firms dislike higher costs. This paper asks whether that explanation is complete. The cost a firm bears in complying need not depend only on its own choices. It may also depend on how its peers behave in the same regulatory setting. I therefore ask a simple question: is regulatory compliance effort a private decision?

The standard treatment assumes that it is. Each firm weighs its own cost of effort against its own benefit of avoiding sanctions, and the regulator is taken to face each firm one at a time (Becker, 1968). Recent episodes suggest a more connected picture. When Silicon Valley Bank failed in March 2023, supervisory scrutiny tightened across regional U.S. banks that had committed no failure of their own, raising compliance costs at peer institutions (Board of Governors of the Federal Reserve System, 2023). After the Sarbanes–Oxley Act imposed internal-control requirements, firms' compliance costs depended in part on the design choices, auditing practices, and implementation experience of other firms in the same regulatory environment (Coates and Srinivasan, 2014). In each case, the focal firm's compliance cost moved with the compliance behavior of its peers, not only with the regulator's action toward that firm alone.

A natural question follows: who bears this burden? It is not borne evenly. Survey evidence from the U.S. Chamber of Commerce and MetLife reports that many small businesses spend more per employee on compliance than their larger competitors. This asymmetry matters for market structure. If peer compliance raises the cost of the focal firm, and if smaller firms are less able to absorb that cost, then compliance spillovers may tilt an industry toward larger firms. The issue is not only that regulation is costly. It is that compliance behavior may be linked across firms in a way that affects costs, survival, concentration, and eventually welfare.

I study this question using compliance records that the U.S. Environmental Protection Agency maintains for regulated facilities. I merge facility-level compliance under the Clean Air Act and the Clean Water Act, drawn from the EPA's ECHO system, with firm-level financials from Compustat and audit data from Audit Analytics. For each firm-year, own compliance is the share of the firm's matched facilities with a clean record in that year. Peer compliance is the average compliance of other public firms in the same state, three-digit NAICS industry, and year. The state $\times$ industry cell is the natural peer group in this setting because environmental enforcement in the United States runs primarily through state agencies: in my data, 77,735 of 101,550 air violations are determined by state agencies. The estimation sample contains 6,628 firm-years across 700 firms, 36 states, and 40 industries over 2010–2025. The outcome is log unit cost, defined as the log of cost of goods sold over sales.

The empirical results suggest that compliance is not fully private. Within firm, higher peer compliance is associated with higher own unit cost. With firm and year fixed effects, own compliance, and firm-level controls, a ten-percentage-point increase in peer compliance is associated with roughly a one percent increase in own unit cost. Own compliance carries the opposite sign: years in which a firm's own facilities comply more are years of lower unit cost. The tightest specification adds NAICS3-by-year fixed effects, which absorb industry-level time-varying shocks and identify the association from cross-state variation within an industry-year. In that specification, a ten-percentage-point increase in peer compliance is associated with roughly a half-percent increase in own unit cost. I interpret this evidence cautiously. It is a conditional within-firm association, not a causal estimate.

The association is not borne evenly. When I estimate the same specification separately by terciles of firm size, the effect is concentrated among small firms. A ten-percentage-point increase in peer compliance is associated with roughly a 2.6 percent increase in unit cost for small firms, and essentially no change for medium or large firms. This finding connects with the broader evidence that regulatory costs vary across the firm-size distribution (Trebbi et al., 2023), but the focus here is different: I study how peers' compliance behavior, rather than the level of regulation alone, is associated with the focal firm's cost.

To interpret this evidence, I develop a model in which firms' compliance decisions are interdependent. Each firm chooses compliance effort that lowers its own marginal cost, but a rival's compliance effort can also affect that cost. The spillover can be cooperative, as in settings with shared compliance infrastructure, common reporting practices, or industry-level diffusion. It can also be rivalrous, as in settings with scarce regulatory attention, limited specialized consultants, or peer benchmarking by regulators. Once compliance decisions are linked in this way, a firm's cost is shaped not only by its own compliance behavior but also by the compliance behavior of its rivals.

I embed this compliance decision in a standard differentiated-products oligopoly (Singh and Vives, 1984; Häckner, 2000). The model delivers two main results. First, when compliance spillovers are cooperative, price competition becomes more profitable relative to quantity competition, because price competition induces more compliance effort and cooperative spillovers make that effort more valuable. Second, when spillovers are rivalrous, price competition can lower total surplus, because the additional effort induced by price competition is partly wasted when rivals' effort raises costs. Consumers, by contrast, are better off under price competition regardless of the sign of the spillover.

The paper contributes to three literatures. First, it contributes to the literature on peer effects, beginning with Manski (1993) and Bramoullé et al. (2009), by studying peer effects in regulatory compliance and costs rather than in behavior alone. Second, it contributes to the literature on regulation, from Becker (1968) to Stigler (1971) and Laffont and Tirole (1991), by moving beyond a purely firm-regulator view of compliance. Third, it contributes to the oligopoly literature following

Singh and Vives (1984), where firms typically enter the competitive stage with costs already fixed. My contribution is to show empirically that compliance behavior is associated with rivals' costs, and theoretically that such spillovers can change the ranking of market outcomes under price and quantity competition.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the institutional setting. Section 4 presents the empirical evidence. Section 5 introduces the model of compliance effort with cross-firm spillovers. Section 6 characterizes equilibrium compliance and market outcomes and compares profits and welfare. Section C shows that the results extend to a general framework. Section 7 concludes.

2 Literature Review

This paper connects three literatures: peer effects in firm behavior, the economics of regulation and compliance, and oligopoly models in which firms' cost choices shape market outcomes. The central question is not only how firms compete once costs are given, but whether regulatory compliance itself is a private decision. If one firm's compliance behavior affects the cost of another firm, then compliance becomes a channel of strategic interaction among firms, not only a bilateral interaction between a firm and a regulator.

The first related literature studies peer effects. Since Manski (1993), the central difficulty has been to distinguish true social or peer effects from correlated group shocks and common environments. Bramoullé et al. (2009) show how network structure can help identify peer effects when individual behavior depends on the behavior of others. In the regulatory setting, Liu et al. (2018) document peer effects in compliance among Mexican gasoline retailers. My paper is closest in spirit to this line of work, but differs in the outcome of interest. Rather than asking only whether firms' compliance behavior moves with peer compliance, I ask whether peers' compliance behavior is associated with the focal firm's cost. This shifts the object from behavioral imitation to cost spillovers.

The second related literature studies regulation and compliance. The classic approach treats compliance as a firm-regulator problem: firms choose whether to comply by comparing the private cost of compliance with the expected penalty from violation (Becker, 1968). The political-economy tradition emphasizes how regulation can be shaped by interest groups, capture, and political incentives (Stigler, 1971; Laffont and Tirole, 1991). These frameworks are central for understanding the private incentives to comply, evade, or influence regulation. My focus is different. I study whether the cost of compliance faced by one firm depends on the compliance behavior of other firms exposed to the same regulatory environment. In this sense, the paper adds a firm-firm channel to a literature that usually emphasizes the firm-regulator channel.

The third related literature studies spillovers across firms. A large body of work analyzes how firms' actions affect rivals through cost asymmetries, technology licensing, product differentiation, ownership structure, and innovation sharing. Acharyya and Marjit (1998) and Zanchettin (2006) study how cost asymmetries affect oligopoly outcomes; Mukherjee (2010) studies technology licensing; and Brander and Spencer (2022) study endogenous product differentiation. López and Vives (2019) and Vives (2020) analyze common and overlapping ownership with R&D spillovers, while Denicolò and Polo (2021) studies innovation-sharing in mergers. The rivalrous side of my mechanism also connects to rent-seeking and lobbying models, where firm-specific effort can impose negative spillovers on rivals through competition for shared political or regulatory resources (Tullock, 1980; Grossman and Helpman, 1994). My contribution relative to this literature is to take regulatory compliance effort as the object of strategic interaction and to allow the spillover to be either cooperative or rivalrous.

The paper also relates to the literature on R&D and process innovation in oligopoly. In two-stage models, firms often commit to R&D before market competition begins. Qiu (1997) shows that Cournot firms may over-invest in R&D relative to Bertrand firms because R&D commitment functions as a strategic substitute that softens rival behavior, generating a welfare reversal. Symeonidis (2003) obtains a related welfare result for sequential product R&D, while Bonanno and Haworth (1998), Delbono and Denicolò (1990), Reynolds and Isaac (1992), and Mukherjee (2011) analyze how competition shapes innovation incentives in related settings. This literature is useful because compliance effort, like process R&D, affects cost. However, the mechanism here is different. I do not treat compliance primarily as a pre-market commitment device. Instead, I model compliance as a cost-affecting decision whose consequences depend on how rival firms comply in the same regulatory environment.

A more recent strand studies simultaneous R&D rather than pre-committed R&D. This approach is motivated by evidence that firms may observe rivals' R&D only late in the development cycle (Cohen, 1995; Koh and Reeb, 2015) and by concerns about the strategic-commitment interpretation of R&D (Vives, 2008). Zhang et al. (2024) show that under positive knowledge spillovers and simultaneous R&D, price competition can become more profitable when products are sufficiently differentiated. I share the simultaneity assumption and the linear-quadratic framework with this analysis, and the cooperative-spillover case in my model is closely related to their positive-spillover mechanism. What is new here is the object and the sign of the spillover: I study compliance rather than R&D, and I allow the spillover to be cooperative or rivalrous.

Finally, I use the Bertrand–Cournot comparison as a framework for studying the market consequences of compliance spillovers. Singh and Vives (1984) show in a differentiated duopoly that quantity competition yields higher profits, while price competition yields higher consumer and total surplus when goods are substitutes, with the rankings reversed for complements. Häckner (2000) extends the comparison to oligopoly, and Amir and Jin (2001) generalizes the analysis beyond linear demand. Subsequent work shows that the standard rankings can change under cost asymmetry (Zanchettin, 2006), vertical bargaining (López, 2007; López and Naylor, 2004), outsourcing (Arya et al., 2008), licensing and vertical structures (Mukherjee et al., 2012; Alipranti et al., 2014; Basak and Mukherjee, 2017), and strategic entry-deterrence decisions (Sharma and Dastidar, 2026). I do not position the present paper as another Bertrand–Cournot ranking exercise. Rather, I use this comparison as a tractable way to show why compliance spillovers matter for market outcomes. Once compliance is not private, the ranking of profits, consumer surplus, and total surplus depends not only on whether firms compete in prices or quantities, but also on whether compliance spillovers are cooperative or rivalrous.

Taken together, these literatures motivate the paper's contribution. I provide evidence that peer compliance is associated with the focal firm's cost in U.S. environmental regulation, and I develop a simple oligopoly model to study the consequences. The main contribution is to bring regulatory compliance into the analysis of firm-firm spillovers: compliance is not only a private cost of regulation, but also a possible channel through which firms affect one another.

3 Institutional Setting

The empirical analysis in Section 4 uses U.S. regulatory compliance data covering tens of thousands of facilities across the universe of Environmental Protection Agency (EPA) regulated industries. The institutional features that may underlie the cross-firm compliance spillovers documented in that section are not obvious from the data alone. This section describes two institutional foundations the analysis rests on: the federal compliance-monitoring infrastructure that produces the underlying data, and two stylized regulatory regimes that provide a vocabulary for interpreting the empirical results.

3.1 Federal Compliance Monitoring

The U.S. Environmental Protection Agency (EPA) maintains compliance records for facilities regulated under the Clean Air Act (CAA), the Clean Water Act (CWA), and the Resource Conservation and Recovery Act, among other statutes. The agency consolidates facility-level enforcement and compliance information through the Environmental Compliance History Online (ECHO) system, which integrates records from program-specific databases including the Integrated Compliance Information System for Air (ICIS-Air), the National Pollutant Discharge Elimination System Integrated Compliance Information System (ICIS-NPDES), and the Facility Registry Service (FRS) that assigns each regulated facility a unique `REGISTRY_ID` usable across programs (U.S. Environmental Protection Agency, 2024b).

ICIS-Air records violations of CAA permits, distinguishing High Priority Violations (HPVs) from lesser violations and recording the date of first determination of each event. ICIS-NPDES records permit-level compliance for facilities discharging pollutants into U.S. waters; the Quarterly Noncompliance Report (QNCR) classifies each permit-quarter into one of several Historic Latest Reported Noncompliance (HLRNC) status codes. Table 1 lists the six HLRNC codes that EPA designates as either Significant Noncompliance (SNC, Category I) or Reportable Noncompliance (RNC, Category II) under 40 CFR §123.45 (Code of Federal Regulations, 2024; U.S. Environmental Protection Agency, 2024c). Each facility-year compliance indicator I construct (Appendix A) aggregates these primitive records into a binary measure of whether the facility experienced any Air violations or any reportable NPDES noncompliance quarters in the given year, with a facility-year coded as noncompliant if its HLRNC status carries any of the codes in Table 1. The construction respects the EPA's own two-tier classification of noncompliance — it excludes facility-quarters coded R (Resolved), V (Non-RNC violations below the reporting threshold), P and Q (Resolved Pending), W (legacy non-major DMR non-receipt), and U (Undetermined).

The reach of the monitoring system is broad. ECHO covers more than one million regulated facilities across the United States (U.S. Environmental Protection Agency, 2024a), and the FRS Program Links table I use to bridge program-specific identifiers contains 4,329,013 facility-program records (a single facility can appear under multiple program codes such as Air, Water, and Hazardous Waste). For the purposes of this paper, the breadth of coverage matters because it provides compliance records for the full universe of regulated facilities, from which I select and match the

Table 1. ICIS-NPDES HLRNC Codes for Significant and Reportable Noncompliance

Code	Type	Description
D	Failure to Report DMR	Reporting violation: non-receipt of the required Discharge Monitoring Report.
E	Effluent — Monthly Avg. Limit	Effluent violations of monthly average DMR limits or SNC-level single event violation.
S	Compliance Schedule — Violations	Enforcement action or permit compliance schedule event violation >90 days late.
T	Compliance Schedule — Reporting	Enforcement action or permit compliance reporting violation >30 days late.
X	Effluent — Non-monthly Avg. Limit	Effluent violations of non-monthly average DMR limits.
N	Reportable Noncompliance	Reportable noncompliance below the SNC threshold but tracked by the QNCR.

Codes D, E, S, T, and X are Category I (Significant Noncompliance, SNC); code N is Category II (Reportable Noncompliance, RNC). Both categories enter the noncompliance indicator in the baseline construction. A facility-year is coded noncompliant if its HLRNC status carries any of these codes in any quarter of the year. Codes R (Resolved), V (Non-RNC violations), P and Q (Resolved Pending), W (legacy non-major DMR non-receipt), and U (Undetermined) are excluded. Source: ECHO ICIS-NPDES Download Summary and Data Element Dictionary, U.S. EPA, available at echo.epa.gov/tools/data-downloads/icis-npdes-download-summary.

facilities of the publicly traded firms whose costs are observable in Compustat. The breadth also permits the construction of compliance measures for the broader regulated industry — small private firms, single-facility operators, conglomerate subsidiaries — which I leave to extensions of the present analysis.

3.2 Two Stylized Regulatory Regimes

I describe two stylized regulatory regimes that operate in the U.S. economy. The two regimes differ in how compliance costs are organized and how enforcement attention is allocated, and they provide a vocabulary for interpreting the empirical results I document in Section 4.2.

The first regime operates through *shared compliance infrastructure*. Industry associations, common reporting frameworks, voluntary certification programs such as ISO 14001 (Potoski and Prakash, 2005), and joint regulatory engagement allow the cost of developing compliance procedures to be partly borne by an institutional carrier rather than by each firm in isolation. In settings of this kind, compliance investment by one firm may reduce the marginal cost of compliance for other firms through the diffusion that shared infrastructure makes possible. The mechanism is closest in spirit to the input-spillover models of Hinloopen (2009) and the technology-diffusion literature, transposed from research and development to compliance.

The second regime operates under *capacity-constrained, peer-benchmarked enforcement*. EPA and its delegated state agencies operate under resource constraints: federal and state inspectors are limited in number, permit reviews and compliance evaluations are scheduled within annual budgets, and enforcement priorities are set in advance through National Enforcement and Compliance Initiatives that focus on specific sectors and pollutants (U.S. Environmental Protection Agency, 2024d). The allocation of inspector attention across the regulated universe appears to depend on several factors, including relative risk, prior violation history, and the capacity available to the inspecting agency in a given fiscal year (Stafford, 2002; Helland, 1998; Gray and Shimshack, 2011). Two features of this enforcement environment matter for interpreting the results in Section 4.2. First, risk-targeting protocols often use industry-relative metrics: facilities may be scored against peer facilities within the same industry classification, which suggests that improvements in peer compliance could raise the relative risk position of a facility whose own compliance does not move. Second, some compliance inputs — specialized regulatory consultants, environmental engineers, attorneys, and the regulator's own time — are likely in finite supply in the short run, so increased peer effort could plausibly raise the marginal cost of obtaining the same inputs for the focal firm.

The two regimes are not exhaustive, and real industries often combine features of both. I do not claim that each U.S. industry maps cleanly to one regime or the other, nor that the regime type is exogenous to the industries that exhibit it. The taxonomy is descriptive rather than predictive, and its usefulness lies in providing a vocabulary for interpreting the results I document in Section 4.2 rather than in forecasting which industries should exhibit which spillover sign.

4 Empirical Evidence

This section asks a single empirical question: is regulatory compliance effort a private decision? If yes, one firm's compliance should not be associated with another firm's costs once industry and time are controlled for. If no, the conditional association should be detectable, and its sign is informative about the spillover structure. Using a panel of U.S. public firms merged with EPA facility-level compliance records, I document that peer compliance is positively associated with own unit cost within firm, while own compliance is negatively associated with it. The association survives firm and year fixed effects and, at reduced magnitude, NAICS3-by-year fixed effects. It is concentrated among small firms (Section 4.3). This evidence motivates the theoretical analysis in Sections 5–C, which studies how a compliance spillover of either sign shapes firms' compliance decisions and market outcomes. The relevant empirical analog of the model's simultaneous-move structure is contemporaneous within-firm variation rather than dynamic adjustment; the evidence below is calibrated to that interpretation.

4.1 Data

I merge three sources. From the EPA's Enforcement and Compliance History Online (ECHO) system I take facility-level compliance under the Clean Air Act and the Clean Water Act; from WRDS Compustat Fundamentals Annual I take firm financials; and from WRDS Audit Analytics I take audit fees and SOX 404 internal-control opinions. I match EPA facilities to Compustat firms by fuzzy name comparison, blocking on three-digit NAICS industry and first letter and keeping pairs above a Jaro–Winkler similarity of 0.90; the match links 6,109 facilities to 1,121 firms. Appendix A documents each source and reports the sample-selection counts at every step.

For each firm-year I build a compliance indicator equal to the share of the firm's matched facilities with a clean record that year, coding a facility noncompliant if it had any Clean Air Act High Priority Violation or any NPDES permit-quarter in significant or reportable noncompliance (Section 3.1). Own compliance is this share. Peer compliance is the leave-one-out mean of own compliance over the other firms in the same state, three-digit NAICS industry, and year,

$$\text{peer\_compliance}_{it} = \frac{1}{N_{g(i),s(i),t} - 1} \sum_{j \in g(i) \cap s(i),\, j \neq i} \text{compliance}_{jt},$$

defined whenever the cell holds at least two firms. I use the state × industry cell because U.S. environmental enforcement runs mainly through state agencies: 77,735 of 101,550 air violations in my data are state-determined. Two plants in the same state and industry share a regulator, a pool of inspectors, and a common set of enforcement priorities; two plants in different states do not. The leave-one-out construction keeps a firm's own compliance out of its peer measure, and because the outcome (cost) differs from the regressor (peer compliance), the reflection problem of Manski (1993) does not apply directly.

The outcome is log unit cost, log(COGS/Sale), the standard average-variable-cost proxy in the empirical industrial-organization and corporate-finance literatures (Foster et al., 2008; De Loecker and Warzynski, 2012); I trim it at $[0.01, 5]$. I control for own compliance, log total assets, leverage, log audit fees, and a SOX 404 material-weakness indicator. Requiring non-missing values on all of these yields an estimation sample of 6,628 firm-years across 700 firms, 36 states, and 40 three-digit industries over 2010–2025. Table 2 reports summary statistics; peer and own compliance are both high on average (means near 0.92), with most of the identifying variation coming from the lower tail and from differences across states within an industry-year.

Table 2. Summary Statistics: Estimation Sample

Variable	N	Mean	SD	P25	Median	P75
log(unit cost)	6,628	−0.538	0.536	−0.644	−0.415	−0.247
peer compliance	6,628	0.920	0.155	0.899	1.000	1.000
own compliance	6,628	0.921	0.220	1.000	1.000	1.000
log AT	6,628	8.204	2.257	7.057	8.417	9.644
leverage	6,628	0.274	0.210	0.138	0.267	0.369
log(audit fees)	6,628	14.681	1.316	13.917	14.732	15.562
material weakness	6,628	0.031	0.173	0.000	0.000	0.000
R&D intensity	3,682	0.050	0.101	0.008	0.021	0.052
SG&A intensity	5,351	0.283	0.580	0.087	0.160	0.281
capital intensity	6,566	0.373	0.279	0.138	0.278	0.644
log(emp)	6,291	1.362	2.079	0.152	1.650	2.812

Source: I compute summary statistics over the 6,628 firm-years in the baseline estimation sample (700 unique firms, 2009–2025). Peer compliance is the state × NAICS3 × year leave-one-out mean of own compliance over firms in the same product market and same state. Both own and peer compliance are bounded in $[0, 1]$; the medians of one reflect that most firms and most state-industry cells are compliant in most years. The variables in the lower panel (R&D intensity, SG&A intensity, capital intensity, log employment) enter extended specifications; their sample sizes are smaller because of additional Compustat missingness on `xrd`, `xsga`, `ppent`, and `emp`.

4.2 Main Results

Table 3 reports the baseline regression of log unit cost on peer compliance, where peer compliance is the state × NAICS3 × year leave-one-out mean. Column (1) pools the data without fixed effects: the coefficient on peer compliance is −0.204 ($t = -3.15$), suggesting that firms in state-industry cells where peers comply more have lower unit costs. This pooled relationship, however, reflects between-cell differences rather than within-firm responses to peer behavior. Adding firm fixed effects in column (2) reverses the sign: peer compliance now has a coefficient of +0.101 ($t = 3.25$). The reversal indicates that the negative pooled correlation comes from systematic between-cell differences in both cost levels and compliance levels, while within a firm, increases in peer compliance are associated with increases in own unit cost.

Columns (3) through (5) successively add year fixed effects, own compliance, and the full set of firm-level controls. The estimated coefficient on peer compliance is stable across these specifications, ranging from 0.096 to 0.097 and remaining statistically significant at the one percent level. The point estimate in column (5) is 0.096, which says that a ten-percentage-point increase in state × NAICS3 peer compliance is associated with a 1.0 percent increase in own unit cost.

The coefficient on own compliance is consistently negative and significant at conventional levels (column (4): -0.030, $p < 0.10$; column (5): -0.034, $p < 0.05$). This pattern says that within a firm, years in which own facilities comply more carry lower unit cost, consistent with the direct cost benefit of one's own compliance technology. The contrast between the negative coefficient on own compliance and the positive coefficient on peer compliance is the central empirical finding of this section: own compliance is associated with lower own cost, while peer compliance is associated with higher own cost.

The control coefficients in column (5) behave as expected from the broader Compustat literature. Larger firms have lower unit costs ($\log \text{AT}$: -0.094, $p < 0.01$), consistent with scale economies. More leveraged firms have weakly higher unit costs ($+0.108$, $p = 0.10$). Firms paying higher audit fees show higher unit costs ($+0.078$, $p < 0.05$), reflecting that audit complexity correlates with cost-of-operations complexity. Firms flagged with SOX 404 material weaknesses have higher unit costs ($+0.059$, $p < 0.05$), consistent with the idea that internal-control deficiencies are correlated with operational inefficiency.

The within R^2 in column (5) is 0.019, indicating that peer compliance and the firm-level controls together explain about 1.9 percent of the within-firm variation in log unit cost beyond what firm and year fixed effects already absorb. This is a modest share but is consistent with the magnitudes reported in the peer effects literature, where peer regressors typically explain a small share of within-unit variation. The point of the regression is not that peer compliance is the dominant driver of unit cost, but that its association with cost is statistically detectable, economically meaningful, and robust to absorbing all common industry-time variation.

Specification (6) is the tightest specification the data support, and I anchor the main interpretation on it. It adds NAICS3-by-year fixed effects on top of firm and year fixed effects, absorbing any industry-level time-varying shock and forcing identification of the association off within-industry-year variation across states. The coefficient on peer compliance is 0.053 (SE 0.029, $p \approx 0.07$), significant at the ten percent level. A ten-percentage-point increase in state × NAICS3 peer compliance is associated with a 0.5 percent increase in own unit cost.

The drop from 0.096 in column (5) to 0.053 in column (6) is informative. Roughly half of the firm-and-year-FE coefficient reflects shared industry-time variation that I cannot cleanly attribute to peer behavior (an industry-wide regulatory shock would generate the same pattern).

The remaining half, identified across states within an industry-year, is the part that I can more plausibly attribute to peer compliance, because no industry-time shock can mechanically produce within-industry-year variation across geographic peer groups. I treat 0.053 as an upper bound on the association that survives standard identification concerns. Specification (5) remains useful as a comparison point and is consistent with the broader peer-effects literature, which typically does not absorb industry-time variation; I report it for completeness but anchor the main interpretation on specification (6).

The coefficient on own compliance is −0.042 ($p < 0.05$) in specification (6), reinforcing the asymmetry: within a firm and within an industry-year, increases in own compliance are associated with lower unit cost while increases in peer compliance are associated with higher unit cost. The control coefficients are stable across (5) and (6).

Table 3. Peer Compliance and Log Unit Cost

	(1) Pooled	(2) Firm FE	(3) +Year FE	(4) +Own	(5) +Controls	(6) +N3×Year
peer compliance	-0.204^{***}	0.101^{***}	0.096^{***}	0.097^{***}	0.096^{***}	0.053^{*}
	(0.065)	(0.031)	(0.029)	(0.029)	(0.029)	(0.029)
own compliance				-0.030^{*}	-0.034^{**}	-0.042^{**}
				(0.016)	(0.016)	(0.018)
log AT					-0.094^{***}	-0.097^{***}
					(0.034)	(0.028)
leverage					0.108	0.056
					(0.066)	(0.061)
log(audit fees)					0.078^{**}	0.075^{**}
					(0.031)	(0.030)
material weakness					0.059^{**}	0.050^{*}
					(0.029)	(0.028)
Firm FE	No	Yes	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	Yes	Yes	Yes
NAICS3 × Year FE	No	No	No	No	No	Yes
Observations	6,628	6,628	6,628	6,628	6,628	6,628
R^2	0.003	0.751	0.753	0.753	0.758	0.809
Within R^2		0.002	0.002	0.002	0.019	0.019

Source: I regress firm-year log unit cost on peer compliance and controls. Peer compliance is the state × NAICS3 × year leave-one-out mean of own compliance. Specifications (2)–(6) include firm and year fixed effects; specification (6) adds NAICS3 × year fixed effects, which absorb any industry-level time-varying shock and force identification off within-industry-year variation across states. Standard errors clustered by firm in parentheses. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

I also estimate the main specification separately within sectors. Under the tighter state × NAICS3 peer construction, most sector cells are too small for precise within-firm estimation, so I read the sector results cautiously. Only one two-digit sector rejects the null at conventional levels: Transportation, where peer compliance is associated with *lower* own cost ($\beta = -0.304$, $p = 0.02$), a cooperative pattern consistent with shared compliance protocols among carriers under common DOT and EPA rules. The remaining sectors show mostly positive but imprecise point estimates, in line with the positive pooled coefficient. I report the full sector composition and the two- and three-digit heterogeneity estimates in Appendix B (Tables 22–25).

4.3 Heterogeneity by Firm Size

To examine whether the association varies with firm size, I estimate the main specification (column (5) of Table 3) separately on three subsamples defined by global terciles of log total assets, computed over the full estimation sample. Firms at or below the 33rd percentile (roughly $1.2 billion in assets) form the small tercile; firms between the 33rd and 67th percentiles (roughly $1.2 billion to $8.5 billion) form the medium tercile; firms above the 67th percentile form the large tercile. The three groups contain approximately equal numbers of firm-years.

Figure 2 reports the heterogeneity results by firm size. The peer-compliance association is concentrated among small firms. For small firms, the coefficient on peer compliance is 0.263, implying that a ten-percentage-point increase in peer compliance is associated with roughly a 2.6 percent increase in unit cost. The corresponding estimates for medium and large firms are close to zero and statistically indistinguishable from zero. This pattern suggests that the peer-related cost of compliance is borne primarily by smaller firms. I interpret this result cautiously: the evidence is associational rather than causal, but it is consistent with the view that compliance spillovers may work in the direction of greater market concentration when smaller firms are less able to absorb peer-driven regulatory costs.

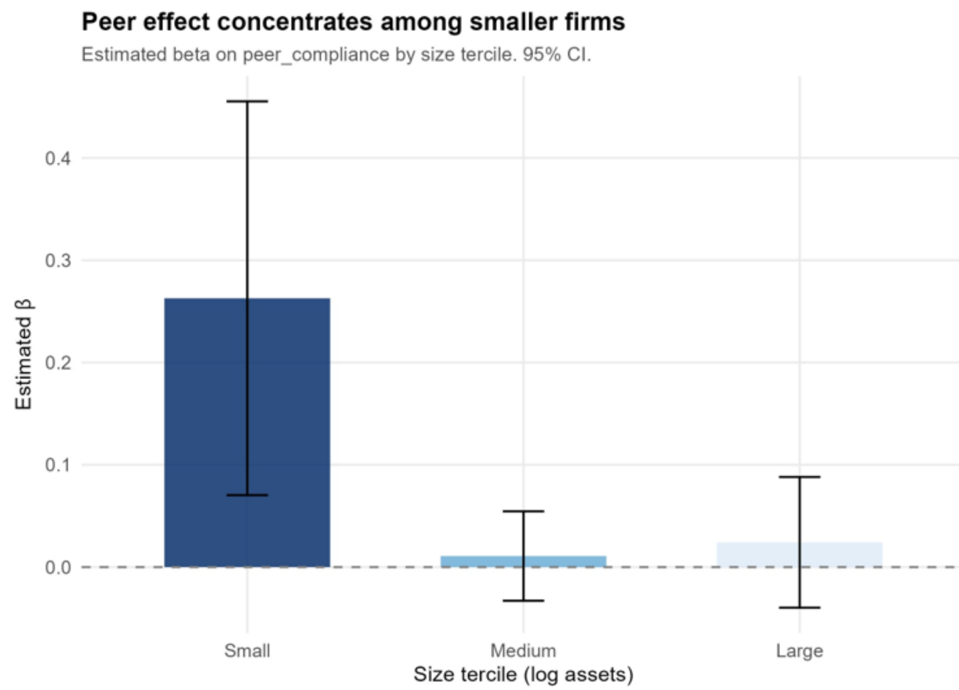

Figure 2. Peer-compliance association by firm size. The figure reports the estimated coefficient on peer compliance from the baseline specification, estimated separately by terciles of firm size based on log total assets. Whiskers denote 95 percent confidence intervals.

The pattern is broadly consistent with evidence that regulatory costs vary across the firm-size distribution (Trebbi et al., 2023).

5 A Model of Compliance Spillovers

The empirical evidence suggests that regulatory compliance is not fully private: a firm's cost is associated not only with its own compliance, but also with the compliance behavior of its peers. This section develops a simple model to interpret that pattern. The purpose of the model is not to describe the full regulatory process, but to isolate one mechanism: compliance effort lowers a firm's own cost, while rivals' compliance effort may either lower or raise that cost through a cross-firm spillover.

Consider an industry with $N \geq 2$ firms producing differentiated products. I use the standard differentiated-products environment of Singh and Vives (1984), extended to N firms as in Häckner (2000). A representative consumer chooses $\mathbf{q} = (q_1, \ldots, q_N)$ to maximize

$$U(\mathbf{q}) - \sum_{i=1}^{N} p_i q_i,$$

where

$$U(\mathbf{q}) = \sum_{i=1}^{N} \alpha q_i - \frac{1}{2}\beta \sum_{i=1}^{N} q_i^2 - \gamma \sum_{i<j} q_i q_j. \tag{1}$$

Here $\alpha > 0$, $\beta > 0$, and $\gamma \geq 0$. The parameter γ captures product substitutability: when $\gamma = 0$, products are independent; as γ rises, products become closer substitutes. I impose the usual concavity condition on utility so that demand is well behaved. The inverse demand system is

$$p_i = \alpha - \beta q_i - \gamma \sum_{j \neq i} q_j, \qquad i = 1, \ldots, N. \tag{2}$$

The new ingredient is compliance effort. Each firm has baseline marginal cost $c > 0$ and chooses effort $e_i \geq 0$. Effort reduces the firm's own effective marginal cost, but the effect of rivals' effort is allowed to go in either direction. Let $E_{-i} = \sum_{j \neq i} e_j$. Firm i's effective marginal cost is

$$\mathcal{C}_i = c - e_i - \phi E_{-i}, \tag{3}$$

where ϕ captures the sign and strength of the compliance spillover.

The sign of ϕ has a direct interpretation. When $\phi > 0$, rivals' compliance effort lowers firm i's cost. This is the cooperative case: firms may benefit from shared compliance infrastructure, common reporting templates, industry association guidance, or diffusion of regulatory know-how. When $\phi < 0$, rivals' compliance effort raises firm i's cost. This is the rivalrous case: firms may compete for scarce regulatory attention, specialized consultants, legal resources, or favorable treatment in a peer-benchmarked enforcement environment. When $\phi = 0$, compliance is purely private and the model collapses to the standard case in which a firm's cost depends only on its own action.

Effort is costly. Firm i incurs

$$\mathcal{K}(e_i) = \vartheta e_i^2, \tag{4}$$

where $\vartheta > 0$ governs the marginal cost of compliance effort. I impose sufficient curvature on this cost function to ensure an interior symmetric equilibrium:

$$\vartheta > \frac{N}{4c}. \tag{5}$$

This assumption plays the same technical role as the curvature restrictions commonly imposed in models of cost-reducing investment and R&D (Qiu, 1997). Economically, it rules out cases in which compliance effort is so cheap that firms would choose unbounded cost reductions.

Firm i's profit is

$$\Pi_i = (p_i - \mathcal{C}_i)q_i - \vartheta e_i^2. \tag{6}$$

Consumer surplus and total surplus are given by

$$CS = U(\mathbf{q}) - \sum_{i=1}^{N} p_i q_i, \tag{7}$$

and

$$TS = U(\mathbf{q}) - \sum_{i=1}^{N} \mathcal{C}_i q_i - \sum_{i=1}^{N} \vartheta e_i^2. \tag{8}$$

This setup differs from the standard Bertrand–Cournot comparison in one important respect. In the usual model, firms enter the market stage with costs already fixed. Here, costs are partly determined by compliance effort, and compliance effort is linked across firms. The comparison between price and quantity competition is therefore not an end in itself; it is a way to study how the market consequences of compliance spillovers depend on the competitive environment.

The model also differs from the sequential R&D literature, where firms commit to cost-reducing investment before market competition begins (Qiu, 1997; Symeonidis, 2003). In the present model, compliance effort is chosen together with the market action. This removes the strategic-commitment channel and focuses attention on the spillover channel: whether rivals' compliance effort lowers or raises a firm's cost.

The linear-quadratic structure is used for tractability. It delivers closed-form expressions for equilibrium compliance, costs, prices, profits, and welfare. Section C then shows that the main qualitative results do not depend on this particular functional form.

6 Equilibrium with Compliance Spillovers

I now characterize the symmetric equilibrium of the model with effort. Restrict attention to substitutes, $\gamma \in [0, 1)$, and impose symmetric intercepts $\alpha_i = a$. Under each regime, firm i jointly chooses its market action and effort level e_i. The externality parameter ϕ is admitted on the full range $[-1, 1]$ throughout.

Cournot. Firm i maximizes

$$\Pi_i^C = \left[a - c - q_i - \gamma \sum_{j \neq i} q_j + e_i + \phi \sum_{j \neq i} e_j \right] q_i - \vartheta e_i^2. \tag{9}$$

The first-order conditions in q_i and e_i, evaluated at the symmetric equilibrium, yield

$$q^C = \frac{2\vartheta(a-c)}{2\vartheta[2+(N-1)\gamma] - [1+\phi(N-1)]}, \qquad e^C = \frac{q^C}{2\vartheta}. \tag{10}$$

Bertrand. Define $A = 1+(N-1)\gamma$, $M = 1+(N-2)\gamma$, $\bar{N} = 2+(N-3)\gamma$, and $B = 1+\phi(N-1)$. Firm i's symmetric-equilibrium output and effort are

$$q^B = \frac{2\vartheta M(a-c)}{2\vartheta A\bar{N} - MB}, \qquad e^B = \frac{q^B}{2\vartheta}. \tag{11}$$

Interiority. The denominators in (10) and (11) must remain strictly positive for the symmetric equilibrium to be interior. The Cournot denominator $2\vartheta[2+(N-1)\gamma]-[1+\phi(N-1)]$ is decreasing in ϕ and strictly positive for ϕ on a neighborhood of zero under the curvature condition (5); the Bertrand denominator behaves analogously. Throughout I restrict attention to the interior region of the parameter space, defined as the set of $(\gamma, \vartheta, \phi, N)$ for which both denominators are positive.

Comparison. The benchmark Singh–Vives ranking on quantities and prices survives the addition of the effort layer, and effort itself inherits the same ranking.

Proposition 1 (Quantity, price, and effort ranking)**.** *For all $\gamma \in (0,1)$ and all $\phi \in [-1,1]$ on the interior region of the parameter space, the symmetric equilibria satisfy $q^B > q^C$, $e^B > e^C$, and $p^B < p^C$.*

Proof. From (10) and (11), $q^B > q^C$ reduces (after cross-multiplication and cancellation) to

$$M\,[2+(N-1)\gamma] > A\,[2+(N-3)\gamma],$$

which simplifies to $(N-1)\gamma^2 > 0$ and holds strictly for $\gamma > 0$. The argument is independent of the sign of ϕ, so the inequality holds across both positive- and negative-externality regimes. The effort ranking $e^B > e^C$ follows immediately from $e^k = q^k/(2\vartheta)$ in both regimes. The price ranking follows from $p = a - Aq$ at the symmetric equilibrium and $A > 0$. □

The benchmark Singh–Vives ranking on prices and quantities therefore extends unchanged to the model with effort, and effort itself inherits the same ranking: Bertrand firms exert more cost-affecting effort than Cournot firms in equilibrium. The reason is mechanical — effort is proportional to equilibrium output, and Bertrand output exceeds Cournot output — but it is precisely this output-scale advantage that interacts with the externality ϕ to drive the profit and welfare results in the following sections. While the rankings of q, e, and p do not depend on the sign of ϕ, the comparison of profits and total surplus does. I characterize how the compliance spillover shapes each comparison next.

6.1 Profit Comparison

Without compliance effort, Singh and Vives (1984) show that profits are higher under quantity competition than under price competition when goods are substitutes. I now show that a cooperative compliance spillover can overturn this: when ϕ is sufficiently positive, price competition becomes the more profitable mode. A rivalrous spillover ($\phi < 0$) never does, so this effect is specific to cooperative compliance spillovers.

Substituting the symmetric equilibria (10) and (11) into the profit expressions yields

$$\Pi^C = (q^C)^2 \left(1 - \tfrac{1}{4\vartheta}\right), \tag{12}$$

$$\Pi^B = (q^B)^2 \left[\tfrac{(1-\gamma)A}{M} - \tfrac{1}{4\vartheta}\right], \tag{13}$$

where A and M are as defined in Section 6. The first proposition of this section characterizes the parameter region in which the standard ranking reverses.

Proposition 2 (Profit reversal across the externality sign)**.** *Suppose $\vartheta > 1/4$, $\gamma > 0$, and the equilibria are interior. Define $K = 4\vartheta - 1$, $m = N - 1$, $M = 1 + (N-2)\gamma$, and let*

$$\phi_L = 2\vartheta\gamma + \frac{MK - \sqrt{MK(MK - 4m\vartheta\gamma^2)}}{2mM}, \quad \phi_U = 2\vartheta\gamma + \frac{MK + \sqrt{MK(MK - 4m\vartheta\gamma^2)}}{2mM}. \tag{14}$$

Then for admissible externalities $\phi \in [-1, 1]$,

$$\Pi^B > \Pi^C \iff \phi \in (\phi_L, \phi_U) \cap [-1, 1],$$

with $\Pi^B = \Pi^C$ at the threshold values and $\Pi^B \leq \Pi^C$ elsewhere on $[-1, 1]$. Both thresholds are *strictly positive, so the reversal arises only in the positive-externality regime; under negative externalities ($\phi < 0$), $\Pi^B < \Pi^C$ throughout the interior region.*

Proof. Define $X = 2\vartheta[2 + (N-1)\gamma] - [1 + \phi(N-1)]$, $s = 2\vartheta(N-1)\gamma^2$, and $d = 2\vartheta\gamma - \phi$. Direct algebra gives

$$\Pi^C - \Pi^B = \frac{\vartheta(a-c)^2\, s\,[2MX(X-K) + Ks]}{X^2(MX-s)^2}. \tag{15}$$

Under interiority the denominator is strictly positive. For $\gamma > 0$, $s > 0$, and the sign of $\Pi^C - \Pi^B$ is governed by the bracketed term. Using $X = K + md$, this term equals $2m\,H(d)$ with

$$H(d) = mMd^2 + MKd + \vartheta K\gamma^2.$$

H is a quadratic in d with positive leading coefficient, so it is negative precisely between its two roots. The discriminant $MK(MK - 4m\vartheta\gamma^2)$ is positive in the interior region. Solving and inverting the substitution $d = 2\vartheta\gamma - \phi$ gives the thresholds ϕ_L, ϕ_U. Since $MK > 0$ under $\vartheta > 1/4$, both thresholds satisfy $\phi_L, \phi_U > 0$, locating the reversal interval entirely within the positive-externality regime. Intersecting with the admissible interval $[-1, 1]$ yields the stated result. □

The headline implication of Proposition 2 is twofold. First, the reversal is a positive-externality phenomenon: rivalrous interactions ($\phi < 0$) preserve the standard Singh–Vives ranking, while cooperative or shared-benefit externalities ($\phi > 0$) can flip it. Second, when both thresholds fall inside the admissible range, the profit difference $\Pi^B - \Pi^C$ is non-monotone in ϕ: Cournot dominates at low and at high externality intensities, and Bertrand dominates only in an intermediate band.

Corollary 1 (Non-monotonic reversal under positive externalities)**.** *Suppose $\gamma > 0$, $\vartheta > 1/4$, and $0 < \phi_L < \phi_U < 1$. Then $\Pi^B - \Pi^C$ is non-monotone in ϕ on $[0, 1]$: Cournot profits dominate at low and high externality intensities ($\phi \in [0, \phi_L) \cup (\phi_U, 1]$), and Bertrand profits dominate in the intermediate range ($\phi \in (\phi_L, \phi_U)$).*

Remark 1 (Monotone-reversal case)**.** *When $\phi_U > 1$ but $\phi_L \in (0, 1)$, the standard Singh–Vives ranking still reverses on the admissible range, but monotonically: Cournot profits dominate for $\phi < \phi_L$ and Bertrand profits dominate for $\phi > \phi_L$. The non-monotonic regime of Corollary 1 is therefore the qualitatively distinctive feature of the model; the monotone-reversal case extends classical reversal results in the literature.*

Figure 3 illustrates the reversal at the calibration $\gamma = 0.55$, $\vartheta = 0.50$, $N = 2$ across the full externality range $\phi \in [-1, 1]$. Under negative externalities, $\Pi^B - \Pi^C$ remains negative: rivalrous interactions preserve Cournot's profit advantage. As ϕ becomes positive and rises past the threshold $\phi_L \approx 0.74$, the profit difference turns positive: cooperative externalities amplify Bertrand competition's output-scale advantage through the cost-reduction channel, and Bertrand profits exceed Cournot profits over the remainder of the interior region. The shape of the curve is consistent with Proposition 2: the profit difference is convex in ϕ and crosses zero exactly once within the interior region at this calibration.

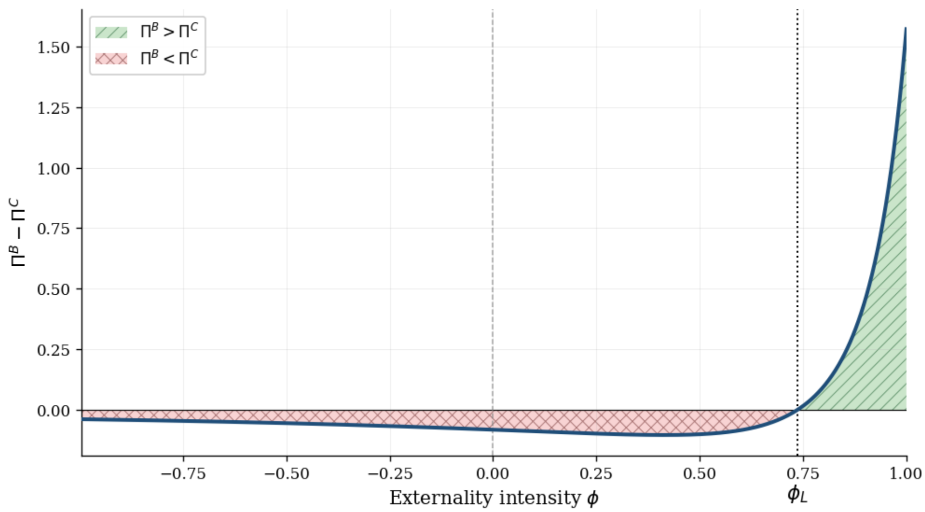

Figure 3. Profit difference $\Pi^B - \Pi^C$ as a function of externality intensity ϕ at $\gamma = 0.55$, $\vartheta = 0.50$, $N = 2$. Negative values of ϕ correspond to rivalrous interactions; positive values correspond to cooperative externalities, separated by the dashed vertical line at $\phi = 0$. The threshold $\phi_L \approx 0.74$ marks the externality intensity at which the standard Singh and Vives (1984) profit ranking reverses: for $\phi < \phi_L$, Cournot competition yields higher profits; for $\phi > \phi_L$, Bertrand competition does. The reversal lies entirely in the positive-externality regime; under negative externalities, the standard Singh–Vives ranking holds throughout. The dashed line at the right edge marks the upper boundary of the interior region of the parameter space at this calibration.

Mechanism. The asymmetry between positive and negative externalities reflects a structural feature of the equilibrium response. Bertrand firms produce more in equilibrium (Proposition 1), so each unit of effort delivers a larger inframarginal cost saving under Bertrand than under Cournot. When externalities are positive, rivals' effort lowers a firm's effective cost, amplifying the productivity of own effort and disproportionately benefiting the higher-output Bertrand regime. As ϕ rises, the output-scale advantage of Bertrand is magnified through the cost channel, and once ϕ crosses ϕ_L, this magnification outweighs the standard Singh–Vives margin advantage of Cournot. At sufficiently high ϕ, the cost reduction becomes nearly symmetric across firms; price competition under Bertrand intensifies and erodes margins faster than the output-scale advantage can compensate, restoring Cournot dominance and producing the non-monotone shape. Under negative externalities, by contrast, rivals' effort raises a firm's effective cost: equilibrium effort is suppressed under both regimes, and the asymmetric amplification mechanism that drives the reversal is absent. The profit difference therefore remains in the standard Singh–Vives direction throughout the negative-externality regime.

This mechanism is structurally distinct from the dynamic-versus-static efficiency trade-off identified by Qiu (1997) and Symeonidis (2003). In their sequential settings, Cournot firms over-invest in R&D because investment is a strategic substitute that softens rival behavior, and the resulting welfare reversal is monotone in spillover intensity. Here, effort is chosen jointly with the market action, eliminating the strategic-commitment channel; the reversal that remains is driven entirely by the interaction between equilibrium output scale and the externality-mediated cost channel, exhibits a genuinely non-monotone regime in ϕ, and arises asymmetrically across the externality sign.

Two-firm specialization. For $N = 2$, the structure of $H(d)$ collapses to a quadratic in γ with explicit roots, yielding sharper thresholds in the substitutability dimension.

Proposition 3 (Substitutability range for higher price-competition profits, $N = 2$)**.** *Suppose $N = 2$, $\vartheta > 1/4$, $\phi \in (0, 1]$, and the equilibria are interior. Then $\Pi^B > \Pi^C$ if and only if*

$$F(\gamma) = \vartheta(8\vartheta - 1)\gamma^2 + 2\vartheta(4\vartheta - 1 - 2\phi)\gamma + \phi(\phi + 1 - 4\vartheta) < 0.$$

When the discriminant $\Delta = 4\vartheta(4\vartheta - 1)\big(4\vartheta^2 + (4\phi - 1)\vartheta - \phi(\phi + 1)\big)$ is positive, the region in which price competition yields higher profits is $\gamma \in (\max\{0, \gamma_L\}, \min\{1, \gamma_U\})$, where γ_L, γ_U are the roots of F.

Proof. Setting $N = 2$ in (15) gives $\Pi^C - \Pi^B$ proportional to $4\vartheta^2(a-c)^2\gamma^2 F(\gamma)$ divided by a positive denominator. The leading coefficient $\vartheta(8\vartheta - 1)$ is positive under $\vartheta > 1/4$, so $F < 0$ holds strictly between the two real roots when the discriminant is positive. □

To make the magnitudes transparent, fix $N = 2$, $a = 2$, $c = 1$, $\vartheta = 1$, $\phi = 1/2$. Closed-form substitution yields $\Pi^B > \Pi^C$ if and only if $28\gamma^2 + 16\gamma - 5 < 0$, i.e. $\gamma \in (0, (\sqrt{51} - 4)/14) \approx (0, 0.224)$. In this calibration, Bertrand profits exceed Cournot profits only when products are sufficiently differentiated; once substitutability rises above the threshold, the standard Singh–Vives ranking reasserts itself.

6.2 Welfare Comparison

I now compare equilibrium consumer surplus and total surplus across the two regimes. The headline finding is that the externality sign affects the welfare comparison and the profit comparison in opposite directions: positive externalities can flip the producer-side ranking (Section 6.1), while negative externalities can flip the social-welfare ranking under specific parameter conditions. Consumer surplus, by contrast, ranks regimes unambiguously throughout the admissible parameter space.

In the symmetric equilibrium, consumer surplus simplifies to a function of equilibrium output:

$$CS = \frac{N(1+(N-1)\gamma)}{2}\, q^2. \tag{16}$$

Since $q^B > q^C$ for all $\gamma > 0$ and $\phi \in [-1, 1]$ on the interior region (Proposition 1), consumer surplus is unconditionally higher under Bertrand competition.

Proposition 4 (Consumer surplus ranking)**.** *For all $\gamma > 0$ and $\phi \in [-1, 1]$ on the interior region of the parameter space, $CS^B > CS^C$.*

Proof. Equation (16) expresses CS as a strictly increasing function of q^2 for $\gamma > 0$. By Proposition 1, $q^B > q^C$ throughout the interior region, hence $CS^B > CS^C$. $\square$

The total-surplus comparison is more subtle because it nets the producer-surplus difference against the consumer-surplus difference. Substituting the symmetric equilibria into $TS = U(\mathbf{q}) - \sum_i \mathcal{C}_i q_i - \sum_i \vartheta e_i^2$ and using the effort first-order condition $e = q/(2\vartheta)$ yields

$$TS = Nq(a-c) + Nq^2\,\Lambda(\phi), \qquad \Lambda(\phi) = \frac{1+2\phi(N-1)}{4\vartheta} - \frac{1+(N-1)\gamma}{2}. \tag{17}$$

The first term is linear in q and unambiguously favors Bertrand because $q^B > q^C$ and $a > c$. The second term has the sign of $\Lambda(\phi)$, which is increasing in ϕ. The total-surplus difference can therefore be written as

$$TS^B - TS^C = N(q^B - q^C)\left[(a-c) + \Lambda(\phi)(q^B + q^C)\right]. \tag{18}$$

Proposition 5 (Total surplus ranking)**.** *For all $\gamma > 0$ and $\phi \in [-1, 1]$ on the interior region:*

1. *If $\Lambda(\phi) \geq 0$, then $TS^B > TS^C$.*

2. *If* $\Lambda(\phi) < 0$*, then* $TS^B > TS^C$ *if and only if* $(a - c) > |\Lambda(\phi)|(q^B + q^C)$.

A sufficient condition for $TS^B > TS^C$ *throughout* $\phi \in [-1, 1]$ *is* $\vartheta \leq (1-2(N-1))/(2(1+(N-1)\gamma))$ *when* $\Lambda(-1) \geq 0$.

Proof. Both claims follow directly from (18). Since $q^B > q^C > 0$ and $a > c$, the first claim is immediate when $\Lambda(\phi) \geq 0$. The second claim follows by sign comparison of the linear and quadratic terms. The sufficient condition holds because Λ is increasing in ϕ, so $\Lambda(-1)$ is the binding case. □

Corollary 2 (Total surplus reversal under negative externalities)**.** *For sufficiently large* γ*,* ϑ*, and* N*, there exists a threshold* $\phi_T < 0$ *such that* $TS^B < TS^C$ *for all* $\phi \in [-1, \phi_T)$ *on the interior region. The reversal arises only in the negative-externality regime; under positive externalities,* $TS^B > TS^C$ *throughout.*

The corollary identifies a parameter region in which strongly rivalrous externalities flip the social-welfare ranking. The driving mechanism is structural: equilibrium effort under Bertrand competition exceeds equilibrium effort under Cournot, and the resource cost of effort enters total surplus negatively. When $\phi < 0$, rivals' effort raises a firm's effective cost; the asymmetry in effort levels means Bertrand firms incur a disproportionately larger gap between gross effort exerted and net cost reduction obtained. At extreme parameters — close substitutes, high effort cost, and many firms — this efficiency loss outweighs Bertrand's consumer-surplus advantage, and Cournot competition delivers higher total surplus.

Combining Propositions 2, 4, and 5 yields the following structure of the Bertrand–Cournot comparison across externality sign: (i) Consumers always prefer Bertrand competition, regardless of externality sign or magnitude. (ii) Firms prefer Cournot competition under negative or weakly positive externalities ($\phi < \phi_L$) and Bertrand competition under sufficiently strong positive externalities ($\phi > \phi_L$). (iii) Society (measured by total surplus) prefers Bertrand competition throughout the positive-externality regime and the moderate-negative-externality regime, but may prefer Cournot competition under strongly rivalrous externalities at extreme parameters.

The two regime-flipping results — profit reversal under positive externalities and total-surplus reversal under negative externalities — arise from distinct mechanisms operating in opposite externality regimes. Positive externalities amplify Bertrand's output-scale advantage through cost reduction; negative externalities amplify Bertrand's effort-cost disadvantage through rivalrous cancellation. Consumer surplus, which depends only on output and not on the resource cost of effort, is insensitive to either mechanism and ranks regimes unambiguously.

Policy implications. Two implications follow. First, the standard Singh–Vives prescription that Bertrand competition is socially preferred survives in the consumer-surplus dimension under all admissible parameters but fails in the total-surplus dimension under strongly rivalrous externalities. Industries with rivalrous lobbying, advertising crowding, or competition for scarce shared resources may therefore present cases where consumer-surplus and total-surplus criteria diverge. Second, the profit reversal under positive externalities aligns firm and consumer preferences in the high-ϕ regime: in industries with strong cooperative externalities, firms voluntarily prefer the consumer-friendly competitive mode, reducing the need for regulatory intervention to push toward Bertrand pricing.

Figure 4 illustrates the total-surplus reversal at parameter values that activate Corollary 2: $\gamma = 0.9$, $\vartheta = 1.5$, $N = 5$. The threshold $\phi_T \approx -0.40$ partitions the admissible range. For $\phi > \phi_T$, including the entire positive-externality regime, Bertrand competition delivers higher total surplus. For $\phi < \phi_T$, rivalrous externalities at this extreme parameter combination flip the social-welfare ranking toward Cournot. The figure also confirms the structural prediction of Proposition 5: the total-surplus difference is monotonically increasing in ϕ, crossing zero exactly once at ϕ_T. Combined with the unconditional consumer-surplus ranking of Proposition 4, the figure shows that strongly rivalrous externalities create a wedge between consumer welfare (which still favors Bertrand) and social welfare (which favors Cournot below ϕ_T).

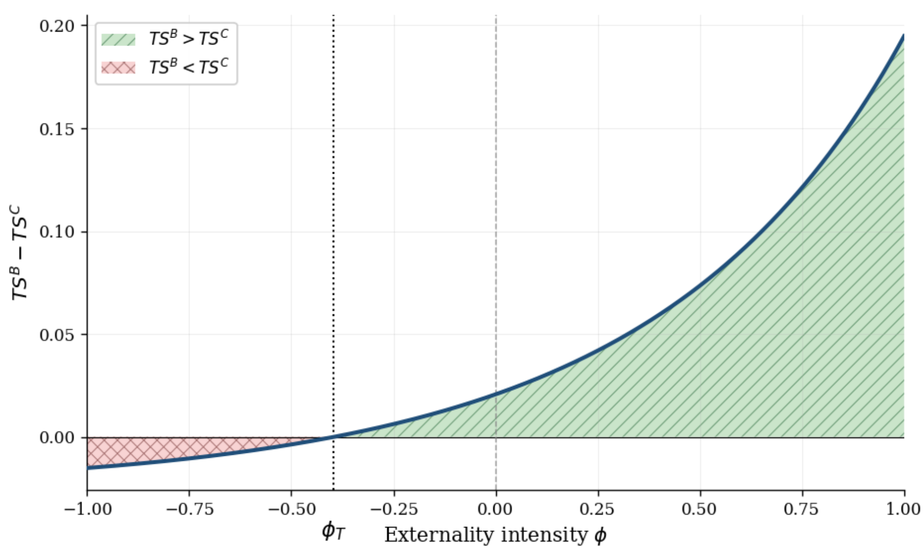

Figure 4. Total surplus difference $TS^B - TS^C$ as a function of the compliance spillover ϕ at $\gamma = 0.9$, $\vartheta = 1.5$, $N = 5$. For $\phi > \phi_T \approx -0.40$, including the entire cooperative regime, price competition delivers higher total surplus; for $\phi < \phi_T$, quantity competition does. The crossing occurs only under strongly rivalrous spillovers. The dashed vertical line at $\phi = 0$ separates the rivalrous and cooperative regimes.

The results above rely on the explicit closed forms of the linear-quadratic model. In Section C I show that the same qualitative pattern — a profit comparison that turns on cooperative spillovers and a total-surplus comparison that turns on rivalrous spillovers — holds under general conditions on demand and cost.

7 Conclusion

This paper asked whether regulatory compliance effort is a private decision. In a panel of U.S. public firms merged with EPA facility-level compliance records, I document a within-firm association between peer compliance and own unit cost. Higher peer compliance is associated with higher own cost, while higher own compliance is associated with lower own cost. The association survives firm, year, and industry-by-year fixed effects, and is concentrated among small firms. I interpret this evidence as suggesting that compliance behavior is not fully private, while acknowledging that the estimates are conditional associations rather than causal effects.

To interpret this evidence, I develop a model in which firms' compliance decisions are linked. Each firm's cost depends on its own compliance effort and, through a spillover of either sign, on the effort of its rivals. Embedding this decision in a standard oligopoly, I show that the sign of the spillover shapes market outcomes. When the spillover is cooperative, price competition becomes more profitable relative to quantity competition; when it is rivalrous, price competition can lower total surplus. Consumers are better off under price competition regardless of the spillover. These results do not claim to capture every feature of regulatory compliance. Their value is to show that once compliance is not private, the sign of the spillover matters for profits, welfare, and the comparison of market regimes.

Several extensions remain for future work. On the empirical side, a stronger parent–subsidiary match using EPA ownership records, and a research design that isolates exogenous variation in peer compliance, would strengthen the evidence beyond the conditional association reported here. On the theory side, endogenizing the regulator's response to firm behavior and calibrating the spillover parameter to observed compliance or lobbying activity would connect the model more tightly to the institutional settings that motivate it. The central message is that compliance should not always be viewed as a purely private cost of regulation: it can also be a channel through which firms affect one another.

References

Acharyya, R. and Marjit, S. (1998). To liberalize or not to liberalize an LDC-market with an inefficient incumbent. *International Review of Economics & Finance*, 7(3):277–296.

Alipranti, M., Milliou, C., and Petrakis, E. (2014). Price vs. quantity competition in a vertically related market. *Economics Letters*, 124(1):122–126.

Amir, R. and Jin, J. Y. (2001). Cournot and Bertrand equilibria compared: Substitutability, complementarity and concavity. *International Journal of Industrial Organization*, 19(3-4):303–317.

Arya, A., Mittendorf, B., and Sappington, D. E. M. (2008). Outsourcing, vertical integration, and price vs. quantity competition. *International Journal of Industrial Organization*, 26(1):1–16.

Ashbaugh-Skaife, H., Collins, D. W., Kinney, W. R., and LaFond, R. (2008). The effect of sox internal control deficiencies and their remediation on accrual quality. *The Accounting Review*, 83(1):217–250.

Basak, D. and Mukherjee, A. (2017). Cournot vs. Bertrand under centralised bargaining. *Economics Letters*, 154:124–127.

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2):169–217.

Board of Governors of the Federal Reserve System (2023). Review of the Federal Reserve's supervision and regulation of Silicon Valley Bank.

Bonanno, G. and Haworth, B. (1998). Intensity of competition and the choice between product and process innovation. *International Journal of Industrial Organization*, 16(4):495–510.

Bramoullé, Y., Djebbari, H., and Fortin, B. (2009). Identification of peer effects through social networks. *Journal of Econometrics*, 150(1):41–55.

Brander, J. A. and Spencer, B. J. (2022). Differentiated entry or "me-too" entry in Bertrand and Cournot oligopoly. *Review of Industrial Organization*, 60:1–27.

Coates, J. C. and Srinivasan, S. (2014). SOX after ten years: A multidisciplinary review. *Accounting Horizons*, 28(3):627–671.

Code of Federal Regulations (2024). 40 CFR section 123.45 — noncompliance and program reporting by the Director. Title 40 – Protection of Environment, Part 123 – State Program Requirements.

Cohen, W. M. (1995). Empirical studies of innovative activity.

De Loecker, J. and Warzynski, F. (2012). Markups and firm-level export status. *American Economic Review*, 102(6):2437–2471.

Delbono, F. and Denicolò, V. (1990). R&D investment in a symmetric and homogeneous oligopoly. *International Journal of Industrial Organization*, 8(2):297–313.

Denicolò, V. and Polo, M. (2021). Mergers and innovation sharing. *Economics Letters*, 202:109841.

Doyle, J., Ge, W., and McVay, S. (2007). Determinants of weaknesses in internal control over financial reporting. *Journal of Accounting and Economics*, 44(1–2):193–223.

Foster, L., Haltiwanger, J., and Syverson, C. (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, 98(1):394–425.

Gray, W. B. and Shimshack, J. P. (2011). The effectiveness of environmental monitoring and enforcement: A review of the empirical evidence. *Review of Environmental Economics and Policy*, 5(1):3–24.

Grossman, G. M. and Helpman, E. (1994). Protection for sale. *The American Economic Review*, 84(4):833–850.

Häckner, J. (2000). A note on price and quantity competition in differentiated oligopolies. *Journal of Economic Theory*, 93(2):233–239.

Helland, E. (1998). The enforcement of pollution control laws: Inspections, violations, and self-reporting. *The Review of Economics and Statistics*, 80(1):141–153.

Hinloopen, J. (2009). Comparing alternative research joint ventures. *Journal of Economics*, 98(3):183–202.

Koh, P.-S. and Reeb, D. M. (2015). Missing R&D. *Journal of Accounting and Economics*, 60(1):73–94.

Laffont, J.-J. and Tirole, J. (1991). The politics of government decision-making: A theory of regulatory capture. *The Quarterly Journal of Economics*, 106(4):1089–1127.

Leary, M. T. and Roberts, M. R. (2014). Do peer firms affect corporate financial policy? *Journal of Finance*, 69(1):139–178.

Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, 70(2):317–341.

Liu, X., Kirwan, B., and Martens, A. (2018). Regulatory compliance, information disclosure and peer effects: Evidence from the Mexican gasoline market. *Journal of Regulatory Economics*, 54(1):53–80.

López, Á. L. and Vives, X. (2019). Overlapping ownership, R&D spillovers, and antitrust policy. *Journal of Political Economy*, 127(5):2394–2437.

López, M. C. (2007). Price and quantity competition in a differentiated duopoly with upstream suppliers. *Journal of Economics and Management Strategy*, 16(2):469–505.

López, M. C. and Naylor, R. A. (2004). The Cournot–Bertrand profit differential: A reversal result in a differentiated duopoly with wage bargaining. *European Economic Review*, 48(3):681–696.

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *Review of Economic Studies*, 60(3):531–542.

Mukherjee, A. (2010). Competition and welfare: The implications of licensing. *The Manchester School*, 78(1):20–40.

Mukherjee, A. (2011). Competition, innovation and welfare. *The Manchester School*, 79(6):1045–1057.

Mukherjee, A., Broll, U., and Mukherjee, S. (2012). Bertrand versus Cournot competition in a vertical structure: A note. *The Manchester School*, 80(5):545–559.

Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297.

Potoski, M. and Prakash, A. (2005). Covenants with weak swords: Iso 14001 and facilities' environmental performance. *Journal of Policy Analysis and Management*, 24(4):745–769.

Qiu, L. D. (1997). On the dynamic efficiency of Bertrand and Cournot equilibria. *Journal of Economic Theory*, 75(1):213–229.

Reynolds, S. S. and Isaac, R. M. (1992). Stochastic innovation and product market organization. *Economic Theory*, 2(4):525–545.

Sharma, V. and Dastidar, K. G. (2026). The optimal bribe: Price versus quantity competition in oligopolies. *The Manchester School*. Forthcoming.

Singh, N. and Vives, X. (1984). Price and quantity competition in a differentiated duopoly. *The RAND Journal of Economics*, 15(4):546–554.

Stafford, S. L. (2002). The effect of punishment on firm compliance with hazardous waste regulations. *Journal of Environmental Economics and Management*, 44(2):290–308.

Stigler, G. J. (1971). The theory of economic regulation. *The Bell Journal of Economics and Management Science*, 2(1):3–21.

Symeonidis, G. (2003). Comparing Cournot and Bertrand equilibria in a differentiated duopoly with product R&D. *International Journal of Industrial Organization*, 21(1):39–55.

Trebbi, F., Zhang, M. B., and Simkovic, M. (2023). The cost of regulatory compliance in the United States. NBER Working Paper w30691, National Bureau of Economic Research.

Tullock, G. (1980). Efficient rent seeking. *Toward a Theory of the Rent-Seeking Society*, pages 97–112.

U.S. Environmental Protection Agency (2024a). ECHO quick start guide. Enforcement and Compliance History Online.

U.S. Environmental Protection Agency (2024b). Enforcement and Compliance History Online (ECHO). Accessed 2025.

U.S. Environmental Protection Agency (2024c). ICIS-NPDES download summary and data element dictionary. Enforcement and Compliance History Online (ECHO).

U.S. Environmental Protection Agency (2024d). National enforcement and compliance initiatives.

Vives, X. (2008). Innovation and competitive pressure. *The Journal of Industrial Economics*, 56(3):419–469.

Vives, X. (2020). Common ownership, market power, and innovation. *International Journal of Industrial Organization*, 70:102528.

Zanchettin, P. (2006). Differentiated duopoly with asymmetric costs. *Journal of Economics and Management Strategy*, 15(4):999–1015.

Zhang, Q., Wang, L. F. S., and Mukherjee, A. (2024). Bertrand–Cournot profit reversal under non-commitment process innovation. *The Manchester School*, 92(4):371–382.

APPENDIX

A Data and Sample Construction

Throughout the paper I distinguish between facilities (physical sites regulated under EPA programs and tracked in ECHO) and firms (legal entities reporting financial information in Compustat). A firm owns one or more facilities; the empirical pipeline aggregates facility-level compliance to firm-year before merging with firm-level financial data.

I construct a firm-year panel that merges three sources. From the EPA Environmental Compliance History Online (ECHO) Exporter and its constituent files (ICIS-Air, ICIS-NPDES, and the Facility Registry Service), I obtain facility-level compliance records under the Clean Air Act and the Clean Water Act. From WRDS Compustat Fundamentals Annual, I obtain firm-level financials including cost of goods sold and net sales. From WRDS Audit Analytics, I obtain audit fees and SOX 404 internal-control opinions used as firm-level controls.

A.1 EPA Compliance Data (ECHO)

I draw facility-level compliance data from the U.S. EPA's Enforcement and Compliance History Online (ECHO) system. ECHO consolidates records from program-specific databases into a common framework keyed on the Facility Registry Service identifier `REGISTRY_ID`. I use six files from the ECHO download distribution; the file selection reflects the goal of constructing a single binary compliance indicator per facility-year from both Clean Air Act and Clean Water Act sources.

Files used and rationale. The six files I load are summarized below.

1. `ECHO_EXPORTER.csv` — the facility master file. It contains 3,141,165 facilities with universal identifier `REGISTRY_ID`, facility name, location, NAICS code, and active-status flag. I use this file as the base table for all facility-level operations.

2. `FRS_PROGRAM_LINKS.csv` — the bridge table mapping `REGISTRY_ID` to program-specific facility identifiers. A single facility may carry distinct identifiers in the Air program and in NPDES; this table reconciles them.

3. `ICIS-AIR_FACILITIES.csv` and `ICIS-AIR_VIOLATION_HISTORY.csv` — the Air-program facility list and violation history, including the High Priority Violator flag and date of first determination. These are the Clean Air Act input to the compliance indicator.

4. ICIS_FACILITIES.csv (within the NPDES download) and NPDES_QNCR_HISTORY.csv — the NPDES facility list and the Quarterly Noncompliance Report. The QNCR contains the Historic Latest Reported Noncompliance (HLRNC) status code that EPA assigns to each permit-quarter and is the Clean Water Act input to the compliance indicator.

Sample windows. I retain ECHO compliance records with year between 2009 and 2025, the window for which I have complete Compustat coverage. The NPDES QNCR history covers 7,855,727 permit-quarters prior to filtering and 6,066,926 permit-quarters after filtering to the sample window (Table 12).

Table 4 summarizes the ECHO facility universe. The ECHO_EXPORTER file contains 3,141,165 facilities with 3,095,510 unique REGISTRY_ID values, of which 1,527,789 carry a non-blank NAICS code. A small share of records (9,852) carry the placeholder name "UNKNOWN" and are filtered out at the firm-match step in Section A.2. The 24,911 facilities flagged as major NPDES dischargers correspond to the subset of facilities subject to the most intensive water-quality reporting requirements; I retain both major and non-major facilities in the baseline construction.

Table 4. ECHO Facility Universe

Statistic	Count
Total facilities in ECHO_EXPORTER	3,141,165
Facilities with non-blank NAICS code	1,527,789
Facilities with placeholder name ("UNKNOWN")	9,852
Facilities flagged active (FAC_ACTIVE_FLAG = Y)	1,351,898
Facilities flagged major NPDES discharger	24,911
Unique REGISTRY_ID values	3,095,510
Unique states represented	90

Source: U.S. EPA Enforcement and Compliance History Online (ECHO) Exporter file. Counts reflect all facilities in the ECHO universe regardless of activity status. The substantial share of facilities without a NAICS code is not consequential for the analysis: matched facilities inherit NAICS classification from the corresponding Compustat firm. The 90 state codes include the 50 states, the District of Columbia, U.S. territories, and a small number of facility records with non-standard state codes used internally by EPA.

The FRS_PROGRAM_LINKS table bridges program-specific facility identifiers to the universal REGISTRY_ID, allowing me to consolidate Air and Water compliance records for the same physical facility. Table 5 shows the distribution of links across EPA programs together with a short description of each. Of the 4,329,013 program-facility records covering 3,362,563 unique REGISTRY_ID values, the largest shares belong to hazardous-waste tracking under RCRAINFO (1,578,620 records), water discharge permits under NPDES (1,162,783 records), and public water systems under the Safe Drinking

Water Act (676,905 records). For the baseline analysis I retain links for the NPDES, AIR, and ICIS programs — the Clean Water Act and Clean Air Act records — which together account for 1,672,012 program-facility links.

Table 5. FRS Program Coverage

Program	Description	Records	Unique Facilities
RCRAINFO	Resource Conservation and Recovery Act (hazardous waste tracking)	1,578,620	1,476,646
NPDES	National Pollutant Discharge Elimination System (water dischargers)	1,162,783	994,147
SFDW	Safe Drinking Water Information System (public water systems)	676,905	452,826
AIR	Clean Air Act stationary sources (current ICIS-AIR program)	279,292	265,828
ICIS	Integrated Compliance Information System (legacy NPDES records)	229,937	207,791
EIS	Emissions Inventory System (criteria air pollutant emissions)	211,197	199,026
TRIS	Toxic Release Inventory (toxic chemical releases)	74,769	73,731
CEDRI	Compliance and Emissions Data Reporting Interface (industry reports)	37,165	35,702
TSCA	Toxic Substances Control Act (industrial chemicals)	27,002	18,190
RMP	Risk Management Plan (accidental chemical release prevention)	21,466	21,183
SEMS	Superfund Enterprise Management System (contaminated sites)	16,155	16,030
E-GGRT	Electronic Greenhouse Gas Reporting Tool (large GHG emitters)	11,540	11,125
CAMDBS	Clean Air Markets Database System (acid rain and cap-and-trade)	2,182	2,136
Total		**4,329,013**	**3,362,563**

Source: U.S. EPA Facility Registry Service (FRS) Program Links table. Each record links a program-specific facility identifier (`PGM_SYS_ID`) to a universal `REGISTRY_ID`, allowing facility-level records across programs to be consolidated. The baseline analysis retains records for the NPDES, AIR, and ICIS programs (which together cover Clean Water Act and Clean Air Act compliance), totaling 1,672,012 program-facility records.

Within the program-facility records retained for the baseline analysis, Table 6 reports the overlap of facilities across the NPDES, AIR, and ICIS programs. Of the 1,399,138 facilities registered in at least one of these programs, 58,376 appear in both an Air and a Water program. This dual-signal universe is the subset where the compliance indicator I construct in Section A.1 draws on both Clean Air Act and Clean Water Act records for the same physical facility. Missingness on join keys is minimal (zero on `REGISTRY_ID` and `PGM_SYS_ID`, 0.06 percent on `PRIMARY_NAME`, and 10 percent on `STATE_CODE`); the state code is recovered downstream from `ECHO_EXPORTER`.

Table 6. FRS Air-Water Program Overlap and Data Quality

Statistic	Count
Facilities in NPDES, AIR, or ICIS (any one)	1,399,138
In only 1 of these programs	1,340,762
In 2 of these programs	48,124
In all 3 (NPDES + AIR + ICIS)	10,252
Facilities in both Air and Water (dual-signal universe)	58,376
Records with missing PRIMARY_NAME	2,506
Records with missing STATE_CODE	436,557

Source: U.S. EPA Facility Registry Service (FRS) Program Links, restricted to NPDES, AIR, and ICIS records. ICIS is the legacy identifier for NPDES records prior to the system migration; AIR is the Clean Air Act stationary-source program. The dual-signal universe (58,376 facilities) is the subset where both Clean Air Act and Clean Water Act compliance can be observed for the same physical facility. Missing STATE_CODE values are recovered from the corresponding ECHO_EXPORTER record in Section A.2.

Table 7 summarizes the ICIS-Air facility universe, which contains 279,124 facilities subject to Clean Air Act stationary-source regulation. Two-thirds of facilities are currently operating, with the remainder permanently or temporarily closed; I retain closed facilities because they may carry historical violation records relevant to the panel. The pollutant-class distribution is heavily weighted toward minor emitters (196,302), with 19,081 Major emitters and 37,414 Synthetic Minor emitters. The Major and Synthetic Minor categories together contain most of the facilities I later match to publicly listed firms in Compustat. Missingness on the universal facility identifier (`registry_id`) is minimal at 268 records (0.1 percent), and missingness on NAICS codes is negligible.

Table 8 reports the annual distribution of ICIS-Air violations within the 2009–2025 sample window. The raw file contains 101,550 violation events, of which 56,761 fall within the sample window and 38,308 cannot be dated because the `earliest_frv_determ_date` field is empty. The annual count of dated violations rises from approximately 1,400 in the early years of the sample to 4,500–5,500 in the most recent years; I do not interpret this trend as a change in compliance behavior because reporting completeness in ECHO has improved substantially over time. Table 9 summarizes the composition of violation records. Of the 101,550 events, 57,032 are federally reportable (FRV) and 44,518 are classified as High Priority (HPV). Lead-agency designations show that 77,735 violations are determined by state agencies, 18,998 by local authorities, and 4,805 directly by the U.S. EPA. The 35,364 unique facilities with violation history correspond to roughly 13 percent of the ICIS-Air facility universe.

Table 10 summarizes the NPDES facility universe. The file contains 1,187,641 permit records

Table 7. ICIS-Air Facility Universe

Statistic	Count
Total facilities in ICIS-AIR	279,124
Operating	184,001
Permanently Closed	77,044
Temporarily Closed	4,017
Major emitters	19,081
Synthetic Minor emitters	37,414
Minor emitters	196,302
Currently flagged High Priority Violator (any state)	3,020
Unique `pgm_sys_id`	279,124
Unique `registry_id`	265,416
Unique states represented	55
Records with missing `registry_id`	268
Records with missing `naics_codes`	4

Source: U.S. EPA ICIS-Air Facilities download. Operating status reflects the facility's current state as of the data release; permanently and temporarily closed facilities are retained because they may have historical violation records relevant to the panel. Pollutant class follows EPA's emissions classification under the Clean Air Act: Major facilities exceed source-specific thresholds for criteria pollutants, Synthetic Minor facilities have accepted federally enforceable limits to remain below the Major threshold, and Minor facilities fall below both. The Major and Synthetic Minor categories together contain most facilities matched to publicly listed firms. The 55 state codes include the 50 states, the District of Columbia, U.S. territories, and a small number of EPA-internal codes.

covering 979,486 unique physical facilities, indicating that some facilities hold multiple permits. The Privately Owned Facility (POF, 422,320) and Corporation (COR, 214,826) categories together account for 637,146 records and contain most facilities I later match to publicly listed firms; municipal, state, and federal facilities are not part of the Compustat universe. Missingness on the universal facility identifier is 3.7 percent (43,454 records), and missingness on the state code is negligible.

Table 11 reports the distribution of HLRNC status codes across the 7,855,727 permit-quarter records in the QNCR history. The Significant Noncompliance codes (D, E, S, T, X) and the Reportable Noncompliance code (N) together account for 2,233,830 permit-quarters, or 28.4 percent of the total. The remaining records are coded as undetermined (U, 31.5 percent), blank (17.3 percent), resolved or pending (R, P, Q, 6.3 percent combined), non-RNC violations (V, 6.5 percent), or legacy non-major DMR non-receipt (W, 9.9 percent); all are excluded from the noncompliance indicator. Table 12 reports the annual distribution of permit-quarters within the 2009–2025 sample window. The volume rises through the mid-2010s as ECHO's reporting completeness improved, then stabilizes in recent years. Within the sample window, 1,520,375 of 6,066,926 permit-quarters

Table 8. ICIS-Air Violations by Year (Sample Window)

Year	Violations
2009	1,417
2010	1,403
2011	1,366
2012	1,119
2013	1,187
2014	1,864
2015	3,205
2016	3,981
2017	3,815
2018	4,350
2019	4,495
2020	3,766
2021	4,979
2022	4,855
2023	5,546
2024	4,941
2025	4,472
Total within sample window (2009–2025)	56,761
Raw violation records (all years)	101,550
Records with missing `earliest_frv_determ_date`	38,308

Source: U.S. EPA ICIS-Air Violation History. Each row represents one violation event at one facility. Violation year is derived from `earliest_frv_determ_date`; the `hpv_dayzero_date` field is empty in the current data release. The 38,308 records with missing determination date cannot be assigned to a year and do not enter the year-aggregated compliance indicator; I retain them in the raw file for completeness.

(25.1 percent) carry a noncompliance code.

Compustat Fundamentals Annual. I retain the variables required for the unit-cost regression (`sale`, `cogs`), size and capital structure controls (`at`, `dltt`, `ppent`, `emp`, `xrd`, `xsga`), and identifiers needed for downstream merges (`gvkey`, `cik`, `conm`, `naics`, `sic`). I do not pull current-asset, dividend, or aggregate profit measures (`ni`, `ebit`, `oibdp`), which are not used in the baseline specification. I also do not use the Compustat Quarterly file or the Segments file, since the analysis operates at the firm-year level and the segments structure would multiply observations within firm. The WRDS export is restricted to the standard industrial universe (`indfmt=INDL`), which excludes financial-services-format records; this explains why a substantial share of firm-years carries missing sales values that are dropped at the positive-sales filter in Table 13.

Table 13 traces the Compustat sample selection. I begin with 156,631 firm-year observations in the WRDS export, which is pre-filtered to the standard industrial universe (`datafmt=STD`,

Table 9. ICIS-Air Violation Records: Response Policy and Lead Agency

Statistic	Count
Total violation records	101,550
FRV (Federally Reportable Violation)	57,032
HPV (High Priority Violation)	44,518
State-led violations	77,735
Local-led violations	18,998
U.S. EPA-led violations	4,805
Other agency	12
Unique facilities with violations	35,364
Records with missing `earliest_frv_determ_date`	38,304

Source: U.S. EPA ICIS-Air Violation History. FRV records correspond to violations that are federally reportable; HPV records correspond to violations classified as High Priority. The lead agency reflects whether the determination was made by a state agency, a local authority, or the U.S. EPA directly. The 35,364 unique facilities are the subset of the 279,124 ICIS-Air facilities with at least one recorded violation in the available history.

`indfmt=INDL`, `consol=C`). Restricting to U.S.-incorporated firms removes 23,045 firm-years and yields 133,586 observations across 17,492 firms. All observations fall within the 2009–2025 sample window. The requirement of positive sales removes 52,275 firm-years — predominantly financial-services firms whose reporting does not include an industrial sales construct — and the requirement of positive cost of goods sold removes a further 711 observations. All remaining firm-years carry a non-blank NAICS code. The final Compustat base contains 80,600 firm-years across 9,474 unique firms.

Table 14 reports missingness on the variables retained for the analysis, computed before any filtering. The 29–30 percent missingness on `sale`, `cogs`, `at`, and `dltt` reflects financial-services-format records carried in the export that lack an industrial reporting construct; these records are removed by the positive-sales filter in Table 13. The 63.9 percent missingness on `xrd` reflects that most firms do not disclose research-and-development expense, consistent with standard Compustat treatment of this variable. Missingness on NAICS is minimal at 4.8 percent. The 24.8 percent missingness on `cik` matters only for the downstream merge to Audit Analytics in Section A.3, where it determines the coverage of audit-fee controls in the final firm-year panel.

Audit Analytics: audit fees. Audit Analytics provides audit-fee disclosures for SEC-registered firms, hand-collected from proxy statements and 10-K filings. Table 15 traces the audit fees sample selection. The raw fees file contains 178,329 records across 21,713 unique CIKs spanning fiscal years

Table 10. NPDES Facility Universe

Statistic	Count
Total NPDES facility records (one per permit)	1,187,641
Unique npdes_id	1,187,641
Unique facility_uin (universal facility ID)	979,486
Unique states represented	69
Privately Owned Facilities (POF)	422,320
Corporations (COR)	214,826
Municipal / Water Districts (MWD)	35,800
Municipalities (CTG)	29,623
State Government (STF)	19,396
Federal Facilities (FDF)	6,770
Other / unknown facility type	458,906
Records with missing facility_uin	43,454
Records with missing state_code	23

Source: U.S. EPA NPDES ICIS_FACILITIES download. Each row is one NPDES permit; some physical facilities hold multiple permits, which is why the count of unique facility_uin (universal facility identifiers) is smaller than the count of unique npdes_id. The Corporation (COR) and Privately Owned Facility (POF) categories together contain most facilities matched to publicly listed firms; municipal, state, and federal facilities are not part of the Compustat universe.

2000–2026. Restricting to the 2009–2025 window removes 1,012 records, and dropping observations with zero audit fees removes a further 1,823 records. After deduplication to one row per CIK-fiscal-year, the audit fees base contains 169,730 records across 21,646 unique CIKs.

Table 16 reports the zero-value distribution for the fee components. Audit Analytics encodes "no fees of this type" as zero rather than missing, so the table separates NA counts (which are zero throughout) from zero counts (which carry economic meaning). The 1,865 records with audit_fees of zero are removed at the positive-fees filter. The other fee components show substantial zero rates: 29 percent of firm-years report no non-audit fees, 63 percent report no audit-related fees, 46 percent report no tax fees, and 77 percent report no other fees. These rates are consistent with the audit-fee literature and reflect that most firms do not engage their auditor for additional service lines.

I retain the fee components needed for the analysis (audit_fees, non_audit_fees, total_fees, audit_related_fees, tax_fees, other_fees), the auditor identity (auditor_name) for descriptive checks, and the restatement flag (restatement). The baseline regression uses audit_fees alone as a control because it is the cleanest measure of audit complexity; non_audit_fees can be endogenous to firm characteristics that also drive compliance behavior, and total_fees mixes the two.

I do not use other Audit Analytics files. The Auditor Changes file records the timing of auditor switches but is not relevant to the cost regression; the Going Concern Opinions file is rare in

Table 11. NPDES QNCR Records by HLRNC Code

Code	Description	Records	In NC Set
U	Undetermined	2,476,567	
(blank)	No code assigned	1,358,678	
N	Reportable Noncompliance	1,274,665	✓
W	Legacy non-major DMR non-receipt	775,289	
D	Failure to Report DMR	529,234	✓
V	Non-RNC Violations	514,542	
R	Resolved	347,473	
E	Effluent — Monthly Avg. Limit	184,370	✓
S	Compliance Schedule — Violations	163,968	✓
P	Resolved Pending	148,869	
X	Effluent — Non-monthly Avg. Limit	59,554	✓
T	Compliance Schedule — Reporting	22,039	✓
Q	Resolved Pending — Manual	479	
Total permit-quarters (all years)		**7,855,727**	
Permit-quarters in NC set (D, E, S, T, X, N)		2,233,830	
Noncompliance share		28.4%	

Source: U.S. EPA NPDES Quarterly Noncompliance Report (QNCR) History. Codes marked with ✓ are included in the noncompliance indicator: Category I Significant Noncompliance (D, E, S, T, X) and Category II Reportable Noncompliance (N). Codes R, V, P, Q, W, U, and blank are excluded because they correspond to resolved status, non-RNC-level violations, or undetermined records. See Table 1 in Section 3.1 for full descriptions and EPA classifications.

occurrence and not a primary outcome; the Non-Reliance Restatements file is more granular than the `restatement` flag carried in the audit fees file, which is sufficient for the analysis. The Auditor Compensation Detail file disaggregates fees into engagement-partner-level components and adds no incremental information for firm-year regressions. These files are available for robustness checks but do not enter the baseline panel.

Audit Analytics: SOX 404 internal-control opinions. Audit Analytics provides internal-control opinions filed under Section 404 of the Sarbanes–Oxley Act. Each firm-year carries two opinions: one issued by the external auditor and one self-reported by management. Table 17 traces the SOX 404 sample selection. The raw file contains 190,533 records across 17,366 unique CIKs spanning fiscal years 2000–2026. I retain only auditor opinions (63,039 records, 33 percent of the raw file) because they are externally verified and more reliable than the self-reported management assessment. The remaining 127,494 management opinions are dropped. Restricting to the 2009–2025 sample window yields 62,906 records. Of these, 59,182 (94 percent) report effective internal controls, and 3,724 (6 percent) carry a material weakness. After deduplication to one record per CIK-fiscal-year (keeping the material-weakness row when duplicates arise), the SOX 404 base

Table 12. NPDES QNCR Permit-Quarters by Year (Sample Window)

Year	Permit-Quarters
2009	218,217
2010	185,839
2011	215,811
2012	299,351
2013	278,393
2014	311,409
2015	343,992
2016	420,305
2017	485,950
2018	509,115
2019	495,108
2020	389,678
2021	410,774
2022	381,517
2023	361,787
2024	383,775
2025	375,905
Total 2009–2025	**6,066,926**
Noncompliance permit-quarters (sample window)	1,520,375
Noncompliance share (sample window)	25.1%

Source: U.S. EPA NPDES Quarterly Noncompliance Report (QNCR) History, restricted to fiscal years 2009–2025. Each permit-quarter is one row; a single facility may have multiple permits and each permit appears in four quarters per fiscal year. The noncompliance share is the fraction of permit-quarters with HLRNC status in {D, E, S, T, X, N}.

contains 61,575 records across 7,528 unique CIKs.

Table 18 breaks down the 3,724 material weakness records by reason. In this data release the granular reason flags for financial fraud, errors, and other causes are not separately populated; the accounting-rule flag captures 99 percent of cases. Because the granular reasons are not reliably distinguishable, the baseline regression uses the binary material weakness indicator rather than reason-specific flags. I construct the material weakness indicator as the logical-or of `ic_is_effective = N` and `count_weak` > 0; the second condition captures the rare case where an auditor reports the opinion as nominally effective while flagging one or more weaknesses, though no such case appears in this sample.

I do not use other variables in the SOX 404 file. The auditor-disagreement flag (`auditor_agrees`) records whether the auditor agrees with management's assessment and is mechanically redundant with the auditor opinion. The exemption flag (`exemption`) identifies smaller-reporting-company exemptions from SOX 404(b) auditor attestation; these filings carry no auditor opinion and are already excluded by the `ic_op_type = a` restriction. The going-concern modifier and the granular

Table 13. Compustat Sample Selection

Step	Firm-Years	Unique Firms
Raw Compustat Fundamentals Annual	156,631	20,510
Standard universe (`datafmt=STD, indfmt=INDL, consol=C`)	156,631	20,510
U.S.-incorporated (`loc=USA`)	133,586	17,492
Sample window (`fyear` 2009–2025)	133,586	17,492
Positive sales (`sale` > 0)	81,311	9,532
Positive cost of goods sold (`cogs` > 0)	80,600	9,474
Non-blank NAICS	80,600	9,474
Final Compustat base	**80,600**	**9,474**

Source: WRDS Compustat Fundamentals Annual. The standard-universe filter (`datafmt=STD`, `indfmt=INDL`, `consol=C`) does not drop observations because the WRDS export was already restricted to this universe. The positive-sales filter removes 52,275 firm-years, predominantly financial-services records that lack an industrial sales construct. The final base contains 80,600 firm-years across 9,474 unique firms.

reason flags add limited variation in this sample. I also do not use other Audit Analytics files such as the Auditor Changes file, the Going Concern Opinions file, or the Non-Reliance Restatements file; the binary material weakness indicator captures the internal-control quality dimension that the baseline specification requires, and the additional files are available for robustness checks but do not enter the main panel.

The ECHO system tracks compliance under several EPA programs beyond Air and Water, but the baseline compliance indicator uses only the Clean Air Act and the Clean Water Act (NPDES). I make this choice deliberately. First, these are the two programs for which ECHO publishes facility-level compliance signals at the resolution required for a year-aggregated indicator: ICIS-Air records violations as dated events, and NPDES records noncompliance status quarterly. Other major programs track different dimensions: the Resource Conservation and Recovery Act tracking system (RCRAINFO) records generator status and hazardous-waste handling rather than violation events at comparable resolution; the Toxic Release Inventory (TRIS) records emissions levels rather than violations; the Safe Drinking Water Information System (SDWIS) applies almost exclusively to municipal utilities outside the Compustat universe. Second, the Air and Water programs together generate the largest dual-signal facility universe relevant to publicly listed firms: 58,376 facilities appear in both an Air and a Water program (Table 6), and the publicly listed firms in this overlap typically hold both Air and Water permits. Restricting to these two programs therefore yields a compliance indicator that is comparable across facilities, industries, and time, while leaving RCRAINFO, TRIS, and SDWIS available for robustness checks involving alternative compliance

Table 14. Missingness on Compustat Variables (Raw, Before Filtering)

Variable	Missing	Percent
`sale`	46,201	29.5
`cogs`	46,200	29.5
`at`	45,698	29.2
`dltt`	46,040	29.4
`ppent`	48,667	31.1
`xrd`	100,038	63.9
`xsga`	64,062	40.9
`emp`	57,369	36.6
`naics`	7,471	4.8
`cik`	38,759	24.8

Source: WRDS Compustat Fundamentals Annual, raw export before any filtering. Missingness counts include records with blank values for character fields. The high missingness on `xrd` (63.9 percent) reflects that most firms do not disclose research-and-development expense, which is standard in the Compustat literature. The 29–30 percent missingness on `sale`, `cogs`, `at`, and `dltt` reflects financial-services records carried in the export that lack an industrial reporting construct; these are removed by the positive-sales filter in Table 13. The 24.8 percent missingness on `cik` affects only the downstream merge to Audit Analytics, which I document in Section A.3.

dimensions.

Constructing the facility-year compliance indicator. Having documented the raw inputs, I now aggregate them into a single facility-year panel. For each facility identified by its universal `REGISTRY_ID`, I construct a binary compliance indicator that equals zero in a given year if the facility experienced either any Clean Air Act High Priority Violation determined during the year or any NPDES permit-quarter coded as significant or reportable noncompliance, and one otherwise. The mapping from program-specific identifiers (`pgm_sys_id` for Air and `npdes_id` for Water) to the universal `REGISTRY_ID` uses the FRS Program Links bridge documented in Table 5. Table 19 summarizes the resulting facility-year panel.

Table 19 summarizes the facility-year compliance panel. The panel contains 6,367,185 facility-year observations covering 811,224 unique facilities over fiscal years 2009–2025. Of these, 4,330,742 facility-years are observed under the Air program only, 1,855,130 under the Water program only, and 181,313 under both. The compliance indicator equals one for 5,819,498 facility-years (91.4 percent) and zero for 547,687 facility-years (8.6 percent).

A facility that is regulated under only one of the two programs carries an NA for the other program's signal. I treat such NAs as "not observed under that program" rather than "no violations under that program," meaning the compliance indicator is computed only from the signals

Table 15. Audit Fees Sample Selection

Step	Records	Unique CIKs
Raw Audit Analytics fees records	178,329	21,713
Sample window (fiscal_year 2009–2025)	177,317	21,690
Positive audit fees (audit_fees > 0)	175,494	21,646
Non-blank CIK	175,494	21,646
After dedup (one row per CIK-year)	**169,730**	**21,646**

Source: WRDS Audit Analytics audit fees file. Restriction to the 2009–2025 sample window removes 1,012 records. The positive-audit-fees filter removes 1,823 additional records with zero or missing audit fees, typically reflecting amended or restated filings. The deduplication step retains one record per CIK-fiscal-year, keeping the largest audit fee value when duplicates arise. Each record links to Compustat through the company_fkey field, which corresponds to the SEC Central Index Key.

Table 16. Audit Fees Components: Zero-Value Distribution

Variable	NAs	Zero Records	% Zero
audit_fees	0	1,865	1.0
non_audit_fees	0	52,284	29.3
total_fees	0	2	0.0
audit_related_fees	0	112,591	63.1
tax_fees	0	81,141	45.5
other_fees	0	136,751	76.7

Source: WRDS Audit Analytics audit fees file, raw export before filtering. The file contains no true missing values for fee components; zero values represent firms that did not pay for that fee type in a given fiscal year. The 1,865 records with audit_fees of zero are dropped at the positive-audit-fees filter in Table 15; these typically reflect amended or restated filings. High zero rates on audit_related_fees, tax_fees, and other_fees reflect that most firms do not engage their auditor for these additional service lines, which is standard in the audit-fee literature.

actually observed. A facility-year is therefore coded noncompliant if any observed signal indicates noncompliance, and compliant otherwise. This treatment preserves the distinction between facilities outside a regulatory program and facilities inside it with a clean record, and it ensures the indicator does not mechanically assign compliance to facilities for which no compliance information exists.

For robustness, I also construct two alternative indicators: a dual-signal indicator restricted to the 181,313 facility-years observed under both Air and Water (where both signals are non-missing), and program-specific indicators (compliance_air and compliance_water) that separate the two channels. The baseline regression uses the combined indicator described above.

Table 17. SOX 404 Sample Selection

Step	Records
Raw SOX 404 records (auditor + management opinions)	190,533
Auditor opinions only (`ic_op_type = a`)	63,039
Sample window (`fy_ic_op` 2009–2025)	62,906
Effective internal controls	59,182
Material weakness flagged	3,724
After dedup (one row per CIK-year)	**61,575**
Unique CIKs	7,528

Source: WRDS Audit Analytics SOX 404 file. The raw file contains two rows per firm-year: one auditor opinion (`ic_op_type = a`) and one self-reported management opinion (`ic_op_type = m`). I retain only the auditor opinion because it is externally verified and more reliable than the self-reported management assessment. The material weakness indicator equals one if `ic_is_effective = N` or `count_weak` > 0. Material weaknesses affect 5.9 percent of auditor-opinion records in the sample window.

Table 18. Material Weakness Reason Breakdown

Reason	Records
Internal controls flagged ineffective (`ic_is_effective = N`)	3,724
due to accounting rule (`noteff_acc_rule = 1`)	3,692
due to financial fraud (`noteff_fin_fraud = 1`)	0
due to errors (`notefferrors = 1`)	0
other reasons (`noteff_other = 1`)	0
Effective opinion but `count_weak` > 0	0
Total material weakness records	**3,724**

Source: WRDS Audit Analytics SOX 404 file, restricted to auditor opinions within the 2009–2025 sample window. The granular reason flags (`noteff_fin_fraud`, `notefferrors`, `noteff_other`) are not separately populated in this data release; `noteff_acc_rule` captures 99 percent of material weakness records. The baseline regression uses the binary material weakness indicator rather than the granular reason flags.

A.2 Facility–Firm Match

I match EPA facilities to Compustat firms by fuzzy name comparison. The procedure cleans both facility names (from `ECHO_EXPORTER` (`FAC_NAME`), with `ICIS-AIR_FACILITIES` (`facility_name`) as fallback) and Compustat firm names (`conm`) by converting to uppercase, replacing punctuation with spaces, and stripping common corporate suffixes (Inc., Corp., LLC, LP, Ltd., Limited, Company, Co., Holdings, Group, International, PLC, SA, AG, NV). I block candidate pairs by three-digit NAICS industry code and the first letter of the cleaned name, which reduces the number of pairwise comparisons by roughly three orders of magnitude relative to the naive product. Within each block I compute the Jaro–Winkler similarity between every facility name and every firm name, and I

Table 19. Facility-Year Compliance Panel

Statistic	Value
Total facility-year observations	6,367,185
Unique facilities (`registry_id`)	811,224
Year range	2009–2025
In Air program only	4,330,742
In Water program only	1,855,130
In both Air and Water programs	181,313
Compliant facility-years (`compliance` = 1)	5,819,498
Noncompliant facility-years (`compliance` = 0)	547,687
Noncompliance share	8.6%

Source: I construct this panel from the EPA ECHO files (ICIS-Air violations, NPDES QNCR, ICIS-Air facilities, NPDES facilities). Each row is one facility-year observation, where facility is identified by the universal `REGISTRY_ID`. The compliance indicator equals zero if the facility had any Clean Air Act High Priority Violation determined during the year or any NPDES permit-quarter coded as significant or reportable noncompliance, and one otherwise. Facilities that appear in only one program (Air or Water) carry an NA for the other program's signal; I interpret these NAs as "not observed under that program" rather than "no violations under that program," so I compute the compliance indicator only from the signals actually observed.

retain the best firm match for each facility.

I keep matches with similarity at least 0.90 in the baseline specification. The threshold choice is principled rather than tuned: in matched pairs across all candidate blocks, mean Jaro–Winkler similarity clusters around 0.896, which suggests that 0.90 captures the natural threshold separating genuine matches from coincidental string similarity. A stricter threshold (e.g., 0.92) discards matches near the natural boundary, while a looser threshold (e.g., 0.85) admits matches that the cleaning procedure cannot reliably distinguish from coincidences. Robustness to thresholds 0.85, 0.92, and 0.95 is available on request.

Table 20 reports the outcome. The procedure matches 6,109 facilities to 1,121 unique Compustat firms (`gvkey`), with 1,759 of these (29 percent) at near-exact similarity. Mean similarity in the matched sample is 0.945 and median similarity is 0.929. The matched sample concentrates in industries where publicly listed firms have substantial physical operations.

I prefer the conservative side of this trade-off where it matters most. Misassigning a facility to the wrong firm contaminates that firm's compliance signal directly; missing a true match merely reduces statistical power. The blocking by NAICS3 protects against the most common type of false positive (facilities with similar names but operating in entirely different industries).

Table 20. Facility–Firm Fuzzy Match

Statistic	Value
Facilities in compliance panel (universe before match)	811,224
Facilities with valid NAICS3 and name	418,608
Compustat firms with valid name and NAICS3	9,434
Match candidates above Jaro–Winkler threshold (0.90)	6,109
of which exact matches (similarity $\geq$ 0.999)	1,759
Unique facilities matched	6,109
Unique Compustat firms matched (`gvkey`)	1,121
Mean Jaro–Winkler similarity (matched pairs)	0.945
Median similarity	0.929

Source: I match EPA facilities to Compustat firms by fuzzy name comparison. Facility names come from `ECHO_EXPORTER` (`FAC_NAME`), with `ICIS-AIR_FACILITIES` (`facility_name`) as fallback when the master file has no name. NAICS classifications come from `ECHO_EXPORTER` (`FAC_NAICS_CODES`), with `ICIS-AIR_FACILITIES` (`naics_codes`) as fallback. Firm names come from Compustat (`conm`). I keep the best Compustat match per facility. The threshold of 0.90 is the baseline, chosen because the mean Jaro–Winkler similarity in matched pairs at lower thresholds clusters around 0.89; robustness to thresholds 0.85, 0.92, and 0.95 is available on request.

A.3 Final Firm-Year Panel

Constructing the firm-year panel. I merge the facility-year compliance panel, Compustat Fundamentals Annual, and Audit Analytics into a single firm-year panel. The merge proceeds in three steps. First, I aggregate facility-year compliance to firm-year by taking the mean of facility compliance across all facilities matched to a given firm-year. For a firm i with F_{it} matched facilities in year t, the own compliance measure is

$$\text{compliance}_{it} = \frac{1}{F_{it}} \sum_{f \in \mathcal{F}_{it}} \mathbb{1}\{\text{facility } f \text{ compliant in year } t\},$$

which lies in $[0, 1]$ and admits a natural interpretation as the share of the firm's regulated facilities that complied in year t. Second, I merge this firm-year compliance variable into the Compustat base on `(gvkey, fyear)`. Third, I merge audit fees and SOX 404 internal-control indicators on `(cik, fiscal_year)`, padding the SEC CIK identifier to ten digits on both sides to handle differences in leading-zero conventions between WRDS exports.

Peer compliance. The main regressor of interest is the leave-one-out mean of compliance over all other firms in the same three-digit NAICS industry, same U.S. state, and same year:

$$\text{peer\_compliance}_{it} = \frac{1}{N_{g(i),s(i),t} - 1} \sum_{j \in g(i) \cap s(i),\, j \neq i} \text{compliance}_{jt},$$

where $g(i)$ denotes firm i's NAICS3 industry, $s(i)$ denotes firm i's headquarters state, and $N_{g(i),s(i),t}$ is the number of firms with non-missing compliance in cell $(g(i), s(i), t)$. Peer compliance is defined whenever the cell contains at least two firms with observed compliance.

This construction operationalizes the cost-coupling primitive of the theoretical model in Section 5, in which rivals' compliance effort enters the focal firm's effective marginal cost. The spillover requires two conditions: firms i and j compete in the same product market, and their compliance costs are linked through a shared regulatory environment. The state $\times$ NAICS3 cell captures both: NAICS3 captures product-market overlap at the level used in most peer-effects work (Leary and Roberts, 2014), and state captures the shared regulator that operates the cost-coupling channel. Environmental compliance in the United States runs primarily through state agencies: state environmental departments issue NPDES permits, conduct most inspections, and lead the majority of enforcement actions (77,735 of 101,550 air violation records in my sample are state-led, against 18,998 local and 4,805 federal). Two chemical plants in the same state therefore share a common pool of inspectors, a common set of enforcement priorities, and a common stock of regulatory learning, while two chemical plants in different states do not.

The leave-one-out construction ensures that a firm's own compliance does not enter its peer compliance (Manski, 1993; Bramoullé et al., 2009; Liu et al., 2018). The reflection problem of Manski (1993) does not apply directly because the dependent variable (log unit cost) differs from the regressor (peer compliance). I address residual identification concerns through firm and year fixed effects; a coarser NAICS3-only peer construction, a lagged version in the spirit of Liu et al. (2018), and a TNIC-3 product-market peer construction are available on request.

Dependent variable. The dependent variable is the log unit cost,

$$\log(\text{unit cost})_{it} = \log\left(\frac{\text{COGS}_{it}}{\text{Sale}_{it}}\right),$$

the log of cost of goods sold scaled by net sales. This measure is the standard average variable cost proxy in the empirical industrial organization and corporate finance literatures (Foster et al., 2008;

De Loecker and Warzynski, 2012). The ratio normalizes for firm size (a flow divided by a flow), the log transformation handles the right-skewed distribution of cost ratios, and the resulting coefficient on peer compliance reads as a semi-elasticity. I trim unit cost at $[0.01, 5]$ to remove extreme values from accounting anomalies; the trim affects fewer than one percent of firm-years.

I considered three alternative cost measures and rejected each for the baseline specification. Cost of goods sold alone (`COGS`) is not size-normalized: larger firms have mechanically larger cost of goods sold, which would conflate scale with cost efficiency. Operating margin, $(\text{Sale} - \text{COGS})/\text{Sale}$, is a linear transformation of unit cost but lacks the variance-stabilizing log transformation. Total factor productivity estimated via Olley–Pakes (Olley and Pakes, 1996) or Levinsohn–Petrin (Levinsohn and Petrin, 2003) measures productivity rather than cost; productivity and cost are related but distinct objects, and my theoretical framework is about cost-affecting effort rather than productivity. Marginal cost estimated via the production-function approach of De Loecker and Warzynski (2012) is structurally cleaner than log unit cost but requires industry-level production function estimation and additional identification assumptions; results using this measure are available on request.

Controls. The baseline specification includes own compliance as a control variable. Without this control, the coefficient on peer compliance would capture both the direct spillover on own cost and any indirect effect operating through the firm's own compliance response. Including own compliance separates these two channels: the coefficient on peer compliance identifies the pure spillover association, net of any influence operating through the firm's own behavior. I include log total assets ($\log \text{AT}$), leverage (DLTT/AT), log audit fees ($\log \text{audit\_fees}$), and a material weakness indicator from SOX 404 as the baseline firm-level controls. The audit-fee and material-weakness controls absorb firm-level audit complexity and internal-control quality, which are correlated with both compliance behavior and operating cost; without them, the peer compliance coefficient could partially reflect omitted variation in firm-level governance (Doyle et al., 2007; Ashbaugh-Skaife et al., 2008). Firm-years without SOX 404 data carry a zero on the material weakness indicator, reflecting that smaller firms not subject to 404(b) auditor attestation are by default not flagged. Additional controls (R&D intensity, SG&A intensity, capital intensity, log employees) enter extended specifications available on request.

Final estimation sample. Table 21 summarizes the resulting panel. Of 80,600 Compustat firm-years in the 2009–2025 window, 11,476 are matched to at least one EPA-regulated facility and carry own compliance. Of these, 7,438 belong to a state × NAICS3 × year cell with at least one other observed firm and therefore have a defined peer compliance. The median cell contains four firms. Audit fees merge to 70,105 firm-years (87 percent of Compustat) and SOX 404 records to 44,206 firm-years (55 percent), with 2,484 firm-years flagged as having a material weakness. Imposing non-missing values on all variables in the baseline specification yields an estimation sample of 6,628 firm-years across 700 unique firms.

The tighter peer definition reduces the estimation sample relative to a coarser NAICS3-only construction. I accept this trade-off because the tighter definition matches the cost-coupling primitive of the theoretical model more closely, identifies the spillover off within-industry-state variation rather than between industries, and survives the inclusion of NAICS3 × year fixed effects (Table 3, column (6)). The coarser NAICS3-only construction is available on request.

Table 21. Final Firm-Year Panel

Statistic	Value
Compustat firm-years (base)	80,600
with own compliance (matched facilities)	11,476
with peer compliance (state × NAICS3-year leave-one-out mean)	7,438
with audit fees	70,105
with SOX 404 data merged	44,206
with material weakness flagged (`material_weakness = 1`)	2,484
with non-missing log unit cost	77,404
Unique firms in estimation sample	700
Median firms per state × NAICS3 × year cell	4
Year range	2009–2025
Estimation sample (all key variables non-missing)	**6,628**

Source: I construct this panel by merging the facility-year compliance panel (aggregated to firm-year via the facility–firm match at Jaro–Winkler threshold 0.90), Compustat Fundamentals Annual, and Audit Analytics. Peer compliance is the leave-one-out mean of own compliance over all other firms in the same NAICS3 industry, same U.S. state, and same year, defined when the cell contains at least two firms with observed compliance. The estimation sample requires non-missing values on log unit cost, peer compliance, own compliance, log assets, leverage, log audit fees, and the material weakness indicator. SOX 404 coverage is lower than audit fees coverage because the auditor attestation requirement under Section 404(b) applies only to accelerated filers; firm-years without SOX 404 data carry zero on the material weakness indicator.

B Additional Heterogeneity Results

This appendix reports the sector composition of the estimation sample and the heterogeneity estimates summarized in Section 4.2.

Table 22. Sample Composition by NAICS2 Sector

NAICS2	Firm-yrs	Firms	States	Years	Peer comp.	Own comp.	log(u.c.)
33 (Mfg. III)	2,410	255	24	2010–2025	0.936	0.937	−0.501
32 (Mfg. II)	1,412	163	19	2010–2025	0.870	0.867	−0.608
22 (Utilities)	980	82	25	2010–2025	0.916	0.926	−0.399
21 (Mining/Oil & Gas)	713	79	6	2010–2025	0.913	0.913	−0.658
31 (Mfg. I)	332	33	9	2010–2025	0.914	0.907	−0.412
48 (Transportation)	197	25	4	2010–2025	0.969	0.975	−0.574
52 (Finance)	187	16	5	2010–2025	0.995	0.995	−1.283
42 (Wholesale Trade)	129	13	4	2010–2025	0.965	0.965	−0.133
Sectors below regression threshold (firm-years $<$ 100):							
53 (Real Estate)	74	5	2	2010–2025	0.993	0.993	−0.489
81 (Other Services)	37	3	1	2010–2025	1.000	1.000	−0.510
51 (Information)	29	6	2	2010–2019	1.000	1.000	−0.880
54 (Prof. Services)	29	4	2	2014–2025	1.000	1.000	−0.436
56 (Admin. Services)	28	3	2	2010–2022	0.938	0.941	−0.465
72 (Accommodation/Food)	28	4	2	2010–2023	0.929	0.929	−0.563
62 (Health Care)	23	3	1	2010–2025	0.986	0.986	−0.215
99 (Unclassified)	14	2	1	2010–2016	1.000	1.000	−0.424
44 (Retail Trade)	6	4	1	2010–2015	1.000	1.000	−0.316
Total	**6,628**	**700**		2010–2025	0.920	0.921	−0.538

Source: I tabulate the baseline estimation sample by 2-digit NAICS sector. Sectors are sorted by firm-year count. “States” counts the number of distinct U.S. states represented in each sector. I apply a 100 firm-year threshold for the sector-specific regressions reported in Table 23 to ensure stable firm and year fixed-effect estimation. The first eight sectors clear this threshold; the remaining nine are reported here for transparency but excluded from the heterogeneity analysis. Mean values for peer compliance and own compliance equal 1.000 in sectors where all observed firms comply in every observed year.

Table 23 reports the main specification estimated separately within each NAICS2 sector with at least 100 firm-years. Only one sector rejects the null at conventional levels: Transportation (NAICS2 = 48) shows a coefficient of -0.304 ($p = 0.02$), a cooperative spillover in which peer compliance is associated with lower own unit cost. The remaining seven sectors show point estimates between 0.01 and 0.39 that are statistically indistinguishable from zero.

I read this pattern conservatively. The tighter state $\times$ NAICS3 peer construction generates within-industry-year variation at the cost of cell size: most NAICS2 sectors retain fewer than 1,000 firm-years and fewer than 25 firms in the sub-sample, which limits power for sector-specific within-firm estimation. The pooled coefficient of $+0.096$ in Table 3 therefore reflects an average over heterogeneous sectors in which most individual estimates are imprecise, with one well-identified cooperative sector (Transportation) and a tendency toward positive but insignificant coefficients elsewhere.

The theoretical framework in Section 5 accommodates both signs of the spillover: cooperative when rivals' effort lowers own effective cost, and rivalrous when it raises it. The Transportation finding is consistent with shared infrastructure or shared compliance protocols among carriers operating under common DOT and EPA rules, where one firm's compliance lowers the cost of all peers' compliance. I do not claim the pooled positive coefficient identifies a specific rivalrous channel; the sector-level evidence is too thin to do so. What the data support is a positive average association with at least one identifiable cooperative subsector, and a theoretical framework that accommodates both signs.

Table 23. Heterogeneity by NAICS2 Sector

NAICS2	N	Firms	States	β_{peer} (SE)	Sign
48 (Transportation)	197	25	4	-0.304^{**} (0.123)	cooperative
31 (Mfg. I)	332	33	9	0.011 (0.027)	—
32 (Mfg. II)	1,412	163	19	0.019 (0.050)	—
42 (Wholesale Trade)	129	13	4	0.021 (0.022)	—
22 (Utilities)	980	82	25	0.022 (0.019)	—
21 (Mining/Oil & Gas)	713	79	6	0.029 (0.227)	—
33 (Mfg. III)	2,410	255	24	0.072 (0.047)	—
52 (Finance)	187	16	5	0.390 (0.238)	—

Source: I estimate the main specification (column (5) of Table 3) separately within each NAICS2 sector with at least 100 firm-years. Each row reports the coefficient on peer compliance, standard error clustered by firm in parentheses, and the implied sign of the spillover ("cooperative" if $\beta < 0$ and significant, "rivalrous" if $\beta > 0$ and significant). Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. The eight sectors are ordered by β_{peer}. Only Transportation rejects the null at conventional levels under the tighter state × NAICS3 peer construction.

Table 24 reports the sample composition by NAICS3 industry. The estimation panel spans 40 NAICS3 industries; the thirteen that clear the 100 firm-year threshold for the heterogeneity regression account for roughly 83 percent of the estimation sample. Chemical Manufacturing (NAICS3 = 325), Utilities (221), and Computer & Electronic Product Manufacturing (334) are the three largest industries by firm-years, each contributing more than 750 observations. Machinery Manufacturing (333), Oil & Gas Extraction (211), and Transportation Equipment Manufacturing (336) also contribute meaningfully to the panel.

Mean peer compliance varies across the above-threshold industries from 0.841 in Primary Metal Manufacturing (331) to 0.993 in Credit Intermediation (522). This cross-industry dispersion is the source of variation that the heterogeneity regression in Table 25 exploits. Mean log unit cost varies from −1.703 in Credit Intermediation (522)—reflecting that financial firms have a different cost structure than industrial firms—to −0.133 in Merchant Wholesalers (424). The heterogeneity regression in the next subsection runs separately within each NAICS3 industry, so this cross-industry dispersion in levels does not contaminate the within-industry identification.

Table 24. Sample Composition by NAICS3 Industry

NAICS3	Firm-yrs	Firms	States	Years	Peer comp.	Own comp.	log(u.c.)
325 (Chemicals)	1,064	119	16	2010–2025	0.891	0.885	−0.725
221 (Utilities)	980	82	25	2010–2025	0.916	0.926	−0.399
334 (Computer/Electronic Mfg.)	780	88	18	2010–2025	0.963	0.968	−0.638
333 (Machinery Mfg.)	543	52	12	2010–2025	0.927	0.931	−0.461
211 (Oil & Gas Extraction)	523	54	4	2010–2025	0.965	0.967	−0.787
336 (Transportation Equip.)	449	42	12	2010–2025	0.929	0.928	−0.260
311 (Food Mfg.)	286	29	7	2010–2025	0.913	0.906	−0.401
339 (Misc. Mfg.)	198	23	7	2010–2025	0.971	0.968	−1.006
331 (Primary Metal Mfg.)	150	18	6	2010–2025	0.841	0.829	−0.191
332 (Fabricated Metal Mfg.)	138	13	5	2010–2025	0.898	0.898	−0.352
486 (Pipeline Transportation)	137	17	2	2010–2025	0.988	0.986	−0.678
522 (Credit Intermediation)	135	11	4	2010–2025	0.993	0.993	−1.703
424 (Merchant Wholesalers)	129	13	4	2010–2025	0.965	0.965	−0.133
27 industries with fewer than 100 firm-years (omitted from regression): 213, 322, 324, 335, 212, 531, 326, 337, 327, 524, 481, 812, 314, 517, 541, 321, 562, 622, 721, 482, 315, 999, 722, 323, 445, 484, 441 (total: 1,116 firm-years)							
Above-threshold total	**5,512**						

Source: I tabulate the baseline estimation sample by 3-digit NAICS industry. Industries are sorted by firm-year count. "States" counts the number of distinct U.S. states represented. The thirteen industries clearing the 100 firm-year threshold enter the heterogeneity regressions in Table 25. The remaining 27 industries are too thin for stable within-firm estimation and are reported here only for transparency. Mean values for peer compliance and own compliance equal 1.000 in industries where all observed firms comply in every observed year.

Table 25 reports the same regression run separately within each NAICS3 industry. No single industry rejects the null at conventional levels under the tighter state × NAICS3 peer construction. Within-industry sample sizes range from 111 to 1,064 firm-years, with firm and year fixed effects absorbing most of the within-firm variation; the remaining identifying variation in any single industry is too thin to support sector-by-sector inference.

I read the table as supportive but not confirmatory. Eleven of the fourteen coefficients are positive, consistent with the positive pooled coefficient in Table 3. Three of the positive coefficients (Chemicals, Computer & Electronic Manufacturing, Machinery Manufacturing) have point estimates between 0.09 and 0.15, in line with the pooled estimate, but with standard errors too large to declare significance. The pattern is consistent with a stable positive average association that the data cannot localize to a single industry under the tighter peer construction. The NAICS2 heterogeneity in Table 23, which pools across NAICS3 sub-industries, provides better-powered sector-level inference and identifies Transportation as a cooperative outlier.

Table 25. Heterogeneity by NAICS3 Industry

NAICS3	N	Firms	States	β_{peer} (SE)
331 (Primary Metal Mfg.)	150	18	6	−0.088 (0.089)
311 (Food Mfg.)	286	29	7	−0.039 (0.034)
336 (Transportation Equip.)	449	42	12	0.012 (0.038)
332 (Fabricated Metal Mfg.)	138	13	5	0.012 (0.035)
424 (Merchant Wholesalers)	129	13	4	0.021 (0.022)
221 (Utilities)	980	82	25	0.022 (0.019)
486 (Pipeline Transportation)	137	17	2	0.065 (1.037)
339 (Misc. Mfg.)	198	23	7	0.078 (0.079)
325 (Chemicals)	1,064	119	16	0.088 (0.107)
334 (Computer/Electronic Mfg.)	780	88	18	0.141 (0.147)
333 (Machinery Mfg.)	543	52	12	0.153 (0.139)
522 (Credit Intermediation)	135	11	4	0.189 (0.297)
213 (Mining Support)	111	14	1	0.430 (0.277)
211 (Oil & Gas Extraction)	523	54	4	0.594 (0.823)

Source: I estimate the main specification (column (5) of Table 3) separately within each NAICS3 industry with at least 100 firm-years. Standard errors clustered by firm in parentheses. No coefficient rejects the null at the ten percent level. The sign distribution is consistent with the positive pooled coefficient in Table 3: eleven of fourteen coefficients are positive and only two coefficients are negative.

C General Analysis

The following assumptions play the same role as the stability and reaction-ordering restrictions in Singh and Vives (1984): they ensure that the symmetric Bertrand and Cournot equilibria are well behaved and can be ranked.

C.1 Assumptions

(A1) *Demand.* $U(\mathbf{q}) : \mathbb{R}_+^N \to \mathbb{R}$ is twice continuously differentiable and strictly concave. Inverse demand $p_i(\mathbf{q}) = \partial U/\partial q_i$ satisfies $\partial p_i/\partial q_i < 0$, $\text{sign}(\partial p_i/\partial q_j) = -\text{sign}(\gamma)$ for $j \neq i$, and $|\partial p_i/\partial q_i| > (N-1)|\partial p_i/\partial q_j|$.

(A2) *Cost function.* $\mathcal{C}(e_i, E_{-i}; \phi)$ is twice continuously differentiable with: (i) $\partial \mathcal{C}/\partial e_i < 0$; (ii) $\text{sign}(\partial \mathcal{C}/\partial E_{-i}) = -\text{sign}(\phi)$; (iii) $\partial^2 \mathcal{C}/\partial e_i \partial \phi = 0$; (iv) $\partial^2 \mathcal{C}/\partial E_{-i}\partial \phi < 0$; (v) $|\partial \mathcal{C}/\partial e_i| \geq |\partial \mathcal{C}/\partial E_{-i}|$; (vi) $\partial^2 \mathcal{C}/\partial e_i \partial e_j = 0$ for $j \neq i$; (vii) $\Psi(e; \phi) \equiv -\partial \mathcal{C}/\partial e_i\big|_{(e,(N-1)e;\phi)}$ is bounded above on $\mathbb{R}_+ \times [-1, 1]$; (viii) $\partial \mathcal{C}/\partial \phi\big|_{E_{-i}=0} = 0$.

(A3) *Effort cost.* $\mathcal{K} : \mathbb{R}_+ \to \mathbb{R}_+$ is twice continuously differentiable, strictly increasing, strictly convex, with $\mathcal{K}(0) = \mathcal{K}'(0) = 0$ and $\mathcal{K}''(e) \geq \kappa > 0$ for all $e \geq 0$.

(A4) *Hessian diagonal dominance.* For both regimes $k \in \{B, C\}$, at every candidate symmetric equilibrium:

(i) $\partial^2 \Pi_i^k / \partial x_i^2 < 0$ and $\partial^2 \Pi_i^k / \partial e_i^2 < 0$, where $x_i \in \{q_i, p_i\}$;

(ii) $\det(H_i^k) > 0$;

(iii) $|\partial^2 \Pi_i^k / \partial x_i^2| > \sum_{j \neq i} |\partial^2 \Pi_i^k / \partial x_i \partial x_j|$ and $|\partial^2 \Pi_i^k / \partial e_i^2| > \sum_{j \neq i} |\partial^2 \Pi_i^k / \partial e_i \partial e_j|$.

(A5) *Market viability.* For all $\phi \in [-1, 1]$, $p(0) - \mathcal{C}(0, 0; \phi) > 0$.

(A6) *Stability* (Friedman, 1977). The symmetric reduced FOC under regime $k \in \{B, C\}$,

$$G^k(q; \phi) \equiv p(q) - \mathcal{C}(h(q), (N-1)h(q); \phi) + q\sigma_k,$$

with $p(q) \equiv p_i(q\mathbf{1})$, $h(q)$ from Lemma 1(a), $\sigma_C \equiv \partial p_i / \partial q_i\big|_{\text{symm}}$, $\sigma_B \equiv 1/(\partial D_i / \partial p_i)\big|_{\text{symm}}$, satisfies $\partial G^k / \partial q < 0$ for all admissible $q > 0$.

(A7) *Demand-side dominance.* The symmetric equilibrium markup $\mu(q; \phi) \equiv p(q) - \mathcal{C}(h(q), (N-1)h(q); \phi)$ is strictly decreasing in q at every symmetric interior equilibrium.

(A8) *No-spillover benchmark.* At $\phi = 0$, $\Pi^C > \Pi^B$ at the symmetric interior equilibrium.

(A9) *Welfare benchmark.* For all $\phi \in [0, 1]$, $CS^B - CS^C > N(\Pi^C - \Pi^B)$ at the symmetric interior equilibrium.

C.2 Equilibrium Orderings

Lemma 1 (Equilibrium orderings)**.** *Under (A1)–(A7), at any symmetric interior equilibrium with* $\gamma > 0$*:*

(a) $e^k = h(q^k)$ *for* $k \in \{B, C\}$*, where* $h : \mathbb{R}_+ \to \mathbb{R}_+$ *is strictly increasing and independent of* ϕ*;*

(b) $q^B > q^C > 0$*;*

(c) $e^B > e^C > 0$*;*

(d) $\mu^C > \mu^B > 0$*.*

Proof. **Part (a): Effort FOC defines h.**

The effort FOC at the symmetric profile $(e_i, E_{-i}) = (e, (N-1)e)$ is

$$-\frac{\partial \mathcal{C}}{\partial e_i}(e, (N-1)e; \phi) \cdot q = \mathcal{K}'(e).$$

Define $\Psi(e; \phi) \equiv -\partial \mathcal{C}/\partial e_i\big|_{(e,(N-1)e;\phi)}$, which is strictly positive by (A2)(i). The FOC becomes

$$\Psi(e; \phi) \cdot q = \mathcal{K}'(e). \tag{19}$$

Let $L(e; q, \phi) \equiv \mathcal{K}'(e) - \Psi(e; \phi)q$.

Boundary values. At $e = 0$: $L(0; q, \phi) = -\Psi(0; \phi)q < 0$ for $q > 0$ by $\mathcal{K}'(0) = 0$ from (A3) and $\Psi > 0$. As $e \to \infty$: $\mathcal{K}'(e) \geq \kappa e \to \infty$ by (A3), while Ψ is bounded above by (A2)(vii); hence $L(e; q, \phi) \to \infty$.

Monotonicity. Differentiating,

$$\frac{\partial L}{\partial e} = \mathcal{K}''(e) - q \cdot \frac{\partial \Psi}{\partial e}.$$

By (A2)(vi), $\partial \Psi/\partial e = -\partial^2 \mathcal{C}/\partial e_i^2$. By (A4)(i), $\partial^2 \Pi/\partial e_i^2 = -q \cdot \partial^2 \mathcal{C}/\partial e_i^2 - \mathcal{K}''(e) < 0$, equivalently

$$\mathcal{K}''(e) + q \cdot \frac{\partial^2 \mathcal{C}}{\partial e_i^2} > 0.$$

Therefore $\partial L/\partial e > 0$ on $\mathbb{R}_+$.

Existence and uniqueness. L is continuous, strictly increasing on $\mathbb{R}_+$, with $L(0) < 0$ and $L(e) \to \infty$. By IVT, there exists a unique $e \in (0, \infty)$ solving $L = 0$. Define $h(q; \phi)$ as this solution.

Smoothness. By IFT,

$$h_q(q; \phi) = \frac{\Psi(h; \phi)}{\mathcal{K}''(h) + q \cdot \partial^2 \mathcal{C}/\partial e_i^2} > 0.$$

Independence of ϕ. By (A2)(iii), $\partial \Psi/\partial \phi = 0$, so equation (19) is ϕ-independent, and so is its unique solution $h(q; \phi) = h(q)$. ■

Part (b): $q^B > q^C$.

By Lemma 1(a), $e^k = h(q^k)$ with h strictly increasing and ϕ-independent. Substituting into the symmetric quantity FOC yields the reduced FOC $G^k(q; \phi)$ defined in (A6).

Cournot. Firm i's FOC in q_i is $p_i + q_i \partial p_i/\partial q_i - \mathcal{C}_i = 0$. At the symmetric profile,

$$G^C(q;\phi) = p(q) - \mathcal{C}(h(q),(N-1)h(q);\phi) + q\sigma_C, \qquad \sigma_C = \alpha,$$

where $\alpha \equiv \partial p_i/\partial q_i\big|_{\text{symm}} < 0$ by (A1).

Bertrand. Firm i's FOC in p_i is $D_i + (p_i - \mathcal{C}_i)\partial D_i/\partial p_i = 0$, giving the Lerner expression $p_i - \mathcal{C}_i = -D_i/(\partial D_i/\partial p_i)$. At the symmetric profile $D_i = q$, so

$$G^B(q;\phi) = p(q) - \mathcal{C}(h(q),(N-1)h(q);\phi) + q\sigma_B, \qquad \sigma_B = \frac{1}{\partial D_i/\partial p_i\big|_{\text{symm}}}.$$

Symmetric demand inversion. The Jacobian of inverse demand at the symmetric profile has diagonal entries α and off-diagonal entries $\beta \equiv \partial p_i/\partial q_j\big|_{\text{symm}}$. By (A1) with $\gamma > 0$, $\beta < 0$ and $|\alpha| > (N-1)|\beta|$, which implies $\det(J) = (\alpha-\beta)^{N-1}(\alpha + (N-1)\beta) \neq 0$, so J is invertible. Direct inversion gives

$$\frac{\partial D_i}{\partial p_i}\bigg|_{\text{symm}} = \frac{\alpha + (N-2)\beta}{(\alpha-\beta)(\alpha+(N-1)\beta)}, \qquad \sigma_B = \frac{(\alpha-\beta)(\alpha+(N-1)\beta)}{\alpha+(N-2)\beta}.$$

Sign chain. By (A1), $|\alpha| > (N-1)|\beta| \geq |\beta|$, so $\alpha - \beta = -|\alpha| + |\beta| < 0$. Also $\alpha + (N-1)\beta = -|\alpha| - (N-1)|\beta| < 0$ and $\alpha + (N-2)\beta < 0$ (sums of negatives). Hence $\sigma_B < 0$ and $\partial D_i/\partial p_i < 0$.

Strategic-channel comparison.

$$\sigma_B - \sigma_C = \frac{(\alpha-\beta)(\alpha+(N-1)\beta) - \alpha(\alpha+(N-2)\beta)}{\alpha+(N-2)\beta}.$$

Numerator: $\alpha^2 + (N-1)\alpha\beta - \alpha\beta - (N-1)\beta^2 - \alpha^2 - (N-2)\alpha\beta = -(N-1)\beta^2 < 0$. Denominator < 0. Hence

$$\sigma_B - \sigma_C = \frac{-(N-1)\beta^2}{\alpha+(N-2)\beta} > 0,$$

with both $\sigma_B, \sigma_C < 0$, so $|\sigma_B| < |\sigma_C|$.

Sign of $G^B(q^C;\phi)$. For any q,

$$G^B(q;\phi) - G^C(q;\phi) = q(\sigma_B - \sigma_C).$$

At $q = q^C$, $G^C(q^C;\phi) = 0$, so $G^B(q^C;\phi) = q^C(\sigma_B - \sigma_C) > 0$.

Ordering. Since a symmetric interior Bertrand equilibrium exists by hypothesis and $G^B(q^C;\phi) >$

0, while $\partial G^B/\partial q < 0$ for all admissible $q > 0$ by (A6), the unique zero q^B of G^B must satisfy $q^B > q^C$.

Positivity of q^C. Since a symmetric interior Cournot equilibrium exists by hypothesis, $q^C > 0$. Uniqueness follows from (A6). Moreover, $G^C(0;\phi) = p(0) - \mathcal{C}(0,0;\phi) > 0$ by (A5), so the interior zero is consistent with market viability. ■

Part (c): $e^B > e^C$.

By Lemma 1(a), h is strictly increasing on $\mathbb{R}_+$. By (b), $q^B > q^C > 0$. Therefore

$$e^B = h(q^B) > h(q^C) = e^C > 0. \qquad ■$$

Part (d): $\mu^C > \mu^B > 0$.

By (A7), $\mu(q;\phi)$ is strictly decreasing in q. By Lemma 1(b), $q^B > q^C$. Hence

$$\mu^B = \mu(q^B;\phi) < \mu(q^C;\phi) = \mu^C.$$

Positivity. From $G^C(q^C;\phi) = 0$:

$$\mu^C = \mu(q^C;\phi) = -q^C\sigma_C = q^C|\alpha| > 0$$

since $q^C > 0$ by (b) and $|\alpha| > 0$ by (A1). From $G^B(q^B;\phi) = 0$:

$$\mu^B = -q^B\sigma_B = q^B|\sigma_B| > 0$$

since $q^B > 0$ by (b) and $|\sigma_B| > 0$ from the formula in (b). ■ □

C.3 Profit-Difference Decomposition

Lemma 2 (Profit-difference decomposition)**.** *At any symmetric interior equilibrium pair $(q^k, e^k, \mu^k)_{k\in\{B,C\}}$,*

$$\Pi^B - \Pi^C = \mu^B(q^B - q^C) - (\mu^C - \mu^B)q^C - [\mathcal{K}(e^B) - \mathcal{K}(e^C)].$$

Proof. By definition, $\Pi^k = \mu^k q^k - \mathcal{K}(e^k)$ for $k \in \{B, C\}$. Hence

$$\Pi^B - \Pi^C = \mu^B q^B - \mu^C q^C - [\mathcal{K}(e^B) - \mathcal{K}(e^C)].$$

Adding and subtracting $\mu^B q^C$:

$$\mu^B q^B - \mu^C q^C = \mu^B(q^B - q^C) + \mu^B q^C - \mu^C q^C = \mu^B(q^B - q^C) - (\mu^C - \mu^B)q^C.$$

Substituting yields the stated decomposition. □

C.4 Comparative Statics in ϕ

Lemma 3 (Comparative statics)**.** *Under (A1)–(A8):*

(a) $q^k(\phi), e^k(\phi), \mu^k(\phi), \Pi^k(\phi)$ *are continuous on* $[-1, 1]$ *for* $k \in \{B, C\}$*;*

(b) $\partial q^k/\partial\phi > 0$ *at every symmetric interior equilibrium;*

(c) $F(\phi) \equiv \Pi^B - \Pi^C$ *is strictly increasing on* $[-1, 1]$*;*

(d) $F(0) < 0$*;*

(e) $F(\phi) < 0$ *for all* $\phi \leq 0$*.*

Proof. (a) By (A6), $\partial G^k/\partial q < 0$ for all admissible $q > 0$. Applying IFT to $G^k(q^k(\phi); \phi) = 0$ yields q^k continuously differentiable in ϕ. Continuity of $e^k = h(q^k)$, μ^k, and Π^k follows from continuity of h (Lemma 1(a)) and of $\mathcal{C}, \mathcal{K}, p$ in their arguments.

(b) By Lemma 1(c), $e^k > 0$ at any interior equilibrium, hence $E^k_{-i} = (N-1)e^k > 0$ for $N \geq 2$. Differentiating $G^k(q^k(\phi); \phi) = 0$:

$$\frac{\partial q^k}{\partial\phi} = -\frac{\partial G^k/\partial\phi}{\partial G^k/\partial q}.$$

The denominator is strictly negative by (A6). For the numerator,

$$\frac{\partial G^k}{\partial\phi} = -\frac{\partial\mathcal{C}}{\partial\phi}\bigg|_{(h(q^k),(N-1)h(q^k);\phi)}.$$

By (A2)(viii), $\partial\mathcal{C}/\partial\phi\big|_{E_{-i}=0} = 0$. By (A2)(iv), $\partial^2\mathcal{C}/\partial E_{-i}\partial\phi < 0$. Integrating from 0 to E^k_{-i}:

$$\frac{\partial\mathcal{C}}{\partial\phi}\bigg|_{E^k_{-i}} = \int_0^{E^k_{-i}} \frac{\partial^2\mathcal{C}}{\partial E_{-i}\partial\phi}\, dE_{-i} < 0.$$

Hence $\partial G^k/\partial\phi > 0$, and $\partial q^k/\partial\phi > 0$.

(c) By (A4), the joint (x_i, e_i) optimum is interior with negative-definite Hessian, so the envelope theorem applies. Since $\partial\mathcal{C}/\partial\phi < 0$ for every interior equilibrium with $E_{-i} > 0$ (by the integration

argument in (b)),

$$\frac{d\Pi^k}{d\phi} = -\frac{\partial \mathcal{C}}{\partial \phi}\bigg|_{(h(q^k),(N-1)h(q^k);\phi)} \cdot q^k = \left|\frac{\partial \mathcal{C}}{\partial \phi}\right|^k q^k > 0.$$

Then

$$F'(\phi) = \left|\frac{\partial \mathcal{C}}{\partial \phi}\right|^B q^B - \left|\frac{\partial \mathcal{C}}{\partial \phi}\right|^C q^C.$$

By Lemma 1(c), $e^B > e^C$, hence $E_{-i}^B = (N-1)e^B > (N-1)e^C = E_{-i}^C > 0$. Since the integrand $\partial^2\mathcal{C}/\partial E_{-i}\partial\phi$ is strictly negative everywhere by (A2)(iv), the magnitude

$$\left|\frac{\partial \mathcal{C}}{\partial \phi}\right|_{E_{-i}} = \int_0^{E_{-i}} \left|\frac{\partial^2 \mathcal{C}}{\partial E_{-i}\partial \phi}\right| dE_{-i}$$

is strictly increasing in E_{-i}. Hence $|\partial\mathcal{C}/\partial\phi|^B > |\partial\mathcal{C}/\partial\phi|^C$. By Lemma 1(b), $q^B > q^C > 0$. Both factors strictly larger under Bertrand, so $F'(\phi) > 0$ on $[-1, 1]$.

(d) Direct from (A8).

(e) By (d), $F(0) < 0$. By (c), F is strictly increasing on $[-1, 0]$, so $F(\phi) < F(0) < 0$ for $\phi < 0$. □

C.5 Asymmetric Profit Reversal

Proposition 6 (Asymmetric profit reversal)**.** *Under (A1)–(A8), if there exists $\phi^* \in (0, 1]$ with $F(\phi^*) > 0$, then there is a unique threshold $\bar{\phi} \in (0, \phi^*)$ such that:*

(i) $\Pi^B > \Pi^C$ *for* $\phi > \bar{\phi}$*;*

(ii) $\Pi^C > \Pi^B$ *for* $\phi < \bar{\phi}$*.*

The reversal can occur only in the cooperative regime ($\phi > 0$).

Proof. By Lemma 2, $\Pi^B > \Pi^C \iff F(\phi) > 0$.

Existence and uniqueness of $\bar{\phi}$. By Lemma 3(d), $F(0) < 0$. The hypothesis $F(\phi^*) > 0$ together with continuity (Lemma 3(a)) and IVT yields some $\bar{\phi} \in (0, \phi^*)$ with $F(\bar{\phi}) = 0$. Strict monotonicity (Lemma 3(c)) gives uniqueness, with $F(\phi) > 0$ for $\phi > \bar{\phi}$ and $F(\phi) < 0$ for $\phi < \bar{\phi}$ on $[0, 1]$.

No reversal on $\phi \leq 0$. By Lemma 3(e), $F(\phi) < 0$ for all $\phi \leq 0$, hence $\Pi^C > \Pi^B$. The reversal cannot occur in the rivalrous regime. □

C.6 Welfare Comparison

Lemma 4 (Consumer surplus ranking)**.** *Under (A1)–(A8), $CS^B > CS^C$ for all $\phi \in [-1, 1]$ at any symmetric interior equilibrium.*

Proof. Consumer surplus at the symmetric profile is

$$CS(q) = U(q\mathbf{1}) - N \cdot p(q) \cdot q.$$

By (A1), $p_i = \partial U/\partial q_i$, so $dU(q\mathbf{1})/dq = \sum_i \partial U/\partial q_i = Np(q)$. Differentiating CS:

$$\frac{dCS}{dq} = Np(q) - Np(q) - Nq \cdot p'(q) = -Nq \cdot p'(q) > 0$$

for $q > 0$, since $p'(q) < 0$ by (A1), where $p'(q)$ denotes the derivative of the symmetric inverse demand $p(q) \equiv p_i(q\mathbf{1})$.

By Lemma 1(b), $q^B > q^C > 0$. Hence $CS^B = CS(q^B) > CS(q^C) = CS^C$. □

Proposition 7 (Asymmetric welfare reversal)**.** *Under (A1)–(A9):*

(i) $CS^B > CS^C$ for all $\phi \in [-1, 1]$;

(ii) $TS^B > TS^C$ for all $\phi \geq 0$;

(iii) If there exists $\phi^{**} \in [-1, 0)$ with $N(\Pi^C - \Pi^B) > CS^B - CS^C$ at ϕ^{**}, then there exists $\underline{\phi} \in (\phi^{**}, 0)$ such that $TS^B(\underline{\phi}) = TS^C(\underline{\phi})$.

The total-surplus reversal can occur only in the rivalrous regime ($\phi < 0$).

Proof. Decomposition: $TS^B - TS^C = (CS^B - CS^C) + N(\Pi^B - \Pi^C)$.

(i) Lemma 4.

(ii) Direct from (A9): for all $\phi \in [0, 1]$, $CS^B - CS^C > N(\Pi^C - \Pi^B)$, equivalently $TS^B - TS^C > 0$.

(iii) The hypothesis gives $TS^B - TS^C < 0$ at ϕ^{**}. By (A9) at $\phi = 0$, $TS^B - TS^C > 0$. By continuity (Lemma 3(a)) and IVT, there exists at least one $\underline{\phi} \in (\phi^{**}, 0)$ such that

$$TS^B(\underline{\phi}) = TS^C(\underline{\phi}).$$

Thus a total-surplus reversal occurs somewhere in the rivalrous regime. □

Corollary 3 (Asymmetric ranking map)**.** *Under (A1)–(A9), profit reversal ($\Pi^B > \Pi^C$) can arise only for sufficiently positive spillovers, while total-surplus reversal ($TS^B < TS^C$) can arise only for sufficiently negative spillovers.*

Proof. Immediate from Propositions 6 and 7. □